

Coopetition Against an Amazon

Ronen Gradwohl

*Ariel University
Ariel 40700, Israel*

RONENG@ARIEL.AC.IL

Moshe Tennenholtz

*The Technion – Israel Institute of Technology
Haifa 3200003, Israel*

MOSHET@IE.TECHNION.AC.IL

Abstract

This paper analyzes cooperative data-sharing between competitors vying to predict a consumer's tastes. We design optimal data-sharing schemes both for when they compete only with each other, and for when they additionally compete with an Amazon—a company with more, better data. We show that simple schemes—threshold rules that probabilistically induce either full data-sharing between competitors, or the full transfer of data from one competitor to another—are either optimal or approximately optimal, depending on properties of the information structure. We also provide conditions under which firms share more data when they face stronger outside competition, and describe situations in which this conclusion is reversed.

1. Introduction

A key challenge to firms competing in today's electronic marketplace is competition against Big Tech companies that have considerably more data and so better predictive models. One way in which smaller firms may overcome this hurdle and survive or even thrive in such a market is to engage in coopetitive strategies—namely, to cooperate with other small firms that are its competitors. Such coopetitive strategies have increasingly become a field of study by both academics and practitioners (Brandenburger & Nalebuff, 2011; Bengtsson & Kock, 2000), but they largely focus on industrial applications such as healthcare, IT, and service industries. In this paper, we study coopetition between e-commerce companies, and focus on the possibility of data sharing as a way to deal with their data imbalance vis-à-vis Big Tech.

Such coopetitive data-sharing can be undertaken by the firms themselves or by external service-providers. To facilitate the former, the newly burgeoning area of federated machine learning has as its goal the design of mechanisms that generate predictive models for firms based on such shared data (Yang et al., 2019). The assumption underlying federated learning is that these firms actually want predictive models based on all shared data, but in competitive scenarios this need not be the case. Although each firm certainly desires its own model to be as predictive as possible, it most likely does not wish the same for its competitors. So when is full data-sharing, leading to maximally predictive models for *all* firms, optimal? When is partial or no sharing better? And do these answers differ when firms compete against an Amazon?

In addition to using federated machine learning, coopetitive data-sharing can also be facilitated by an external service-provider. While tools for data-sharing are prevalent—

and include platforms such as Google Merchant and Azure Data Share—they are currently undergoing even further development. For example, in its European Strategy for Data, the European Commission plans the following: “In the period 2021-2027, the Commission will invest in a High Impact Project on European data spaces and federated cloud infrastructures. The project will fund infrastructures, data-sharing tools, architectures and governance mechanisms for thriving data-sharing and Artificial Intelligence ecosystems” (European Commission, 2020). But just as with federated learning, the potential benefits from data-sharing may arise under partial, rather than full, sharing. For instance, one industry white paper urges service providers to

“...analyze the combined, crowdsourced data and generate benchmark analyses and comparative performance reports. Each participating client gains insights that it could not otherwise access, and each benefits from the service-host provider’s ability to *slice and dice* the aggregated data and share the results that are relevant to each client” (Loshen, 2014).

What is the optimal way for such service providers to combine, slice and dice, and share e-commerce firms’ data in order to facilitate successful co-competition against an Amazon?

Regardless of how co-competitive data-sharing is achieved—via federated learning or through a service provider—the answer to whether or not firms desire maximally predictive models for all depends on how these predictions are eventually used. If the firms involved act in completely unrelated markets, then of course full data-sharing is optimal—it yields a maximally predictive model for a firm, and that firm is not harmed by others’ better predictions. This is no longer true if the firms are competitors, since then a firm may be harmed by an improvement in its competitors’ predictions. In this case, whether or not the tradeoff is worthwhile depends on the details of the market in which firms compete.

In this paper we focus on one such market, which is motivated by recommendation systems for online advertising and Long Tail retail. In our model, there are many types of consumers, and firms use their data to infer and take a tailored action for each consumer type. In the online advertising market, for example, firms use their data to personalize advertisements in order to maximize the probability that a consumer clicks on their ad. Data-sharing is common in this market, and platforms such as Google Merchant and Azure Data Share specifically facilitate such sharing across clients in order to improve personalization and increase clickthrough rates.

For another concrete application, consider retailers in a Long Tail market—a market with an enormous number of low-demand goods that collectively make up substantial market share—interested in finding consumers for their products. Success in such markets crucially relies on retailers’ abilities to predict consumers’ tastes in order to match them to relevant products (Anderson, 2006; Pathak et al., 2010). On their own, smaller firms can be destroyed by giants with vastly greater amounts of data and thus a clear predictive advantage. In cooperation with other small firms, however, this data imbalance may be mitigated, possibly granting smaller firms a fighting chance.¹

Finally, a major motivation of our work is to address a critical obstacle faced by recent highly successful AI systems, systems that involve creating huge models through neural-

1. We do not discuss here antitrust issues surrounding data-sharing; the reader is referred to Martens et al. (2020) for a discussion thereof in the context of the EU’s European Strategy for Data.

network architectures exploiting vast amounts of data. While the so-called pre-training of such models typically utilizes publicly available data (Brown et al., 2020), making the models useful through various techniques such as fine-tuning or prompt-tuning (Lester et al., 2021) likely involves proprietary data. And because data relevant to accomplishing nearly all economic tasks is often distributed among competitors (or “coopetitors”), the incentives for data-sharing become a major obstacle. Indeed, this calls for a wide exploration of data-sharing design, dealing with the actual mechanisms/algorithms that determine what and how much data is to be shared, in order to overcome this fundamental obstacle. Our paper essentially initiates work on this topic, focusing on simple recommendation systems. Although simple, the framework is applicable to realistic scenarios of coopetition, e.g., among retailers (such as in Azure for Retail).

In particular, in this paper we study such cooperation between competing firms, such as retailers, and design data-sharing schemes that are optimal—ones that maximize total firm profits subject to each firm having the proper incentives to participate. More specifically, we study two settings, one without and the other with an Amazon. In the former, in which smaller firms compete only with one another, we show that coopetition is beneficial to the extent that firms can share data to simultaneously improve their respective predictions. In the optimal data-sharing scheme firms benefit from sharing because it allows them to better predict the tastes of consumers for whom no firm predicts well on its own.

In the latter setting, in which smaller firms compete with one another but also with an Amazon, sharing data about consumers for whom no firm predicts well can be beneficial, but is not optimal. We show that, in the optimal scheme, firms share data with others and weaken their own market positions on some market segments, in exchange for receiving data from others and strengthening their market positions on other segments.

The following simple example illustrates our model and previews some of our results. The example is framed as retailers offering goods to consumers who may make a purchase, but can also be interpreted as advertisers displaying advertisements to consumers who may click on the ads. In either case, suppose the market consists of two goods, g_0 and g_1 , that consumers may desire. Each consumer is described by a feature vector in $\{0, 1\} \times \{0, 1\}$ that determines that consumer’s taste: Consumers of types 00 and 11 are only interested in g_0 , whereas consumers of types 01 and 10 are only interested in g_1 . A priori, suppose the distribution of consumers in the market is such that an α fraction are of type 00, a β fraction of type 01, a γ fraction of type 10, and a δ fraction of type 11.

When a consumer shows up to the market, a retailer may offer him one of the goods (as a convention, we use masculine pronouns for consumers and feminine pronouns for retailers/players). If the retailer has data about the consumer’s type, she will offer the good desired by that consumer. Assume that when a consumer is offered the correct good he makes the purchase, and that this leads to a profit of 1 to the retailer. If there are two retailers offering this good to the consumer he chooses one at random from whom to make the purchase. A retailer whose good is not chosen gets profit 0.

Suppose now that there are two retailers pursuing the same consumer, but that they do not fully know the consumer’s type. Instead, the first retailer only knows the first bit of the consumer’s type, and the second retailer only knows the second bit. This situation is summarized in Figure 1, where the first retailer knows the row bit and the second knows the column bit. Which goods should the competing retailers offer the consumer? For this

	0	1
0	g_0	g_1
1	g_1	g_0

(a) Desired good

	0	1
0	α	β
1	γ	δ

(b) Prior probability

Figure 1: Data-sharing: An example

example, suppose for simplicity that $\alpha > 2\beta \geq 2\gamma > 4\delta$. Then each of the retailers has a dominant strategy: If the bit they know is 0, offer g_0 , and if the bit they know is 1, offer g_1 .

To see that this is a dominant strategy, consider for example the row retailer, and suppose she learns that a consumer’s row bit is 0. She thus knows that the correct good is g_0 with probability $\alpha/(\alpha + \beta)$ and g_1 with probability $\beta/(\alpha + \beta)$. Her utility, however, depends also on the good offered by the column retailer. The row retailer’s worst-case utility from offering g_0 is $\alpha/(2\alpha + 2\beta)$, which occurs when the column player also offers g_0 . Her best case utility from offering g_1 is $\beta/(\alpha + \beta)$. Since we assumed $\alpha > 2\beta$, offering g_0 is best for the row player regardless of the other retailer’s offer, and is thus dominant. The analyses for the case in which the row retailer’s bit is 1, as well as for the column retailer’s strategy, are similar.

The result from deployment of these dominant strategies is that consumers of type 00 will be offered the correct good by both retailers, and will thus choose one at random. Consumers of types 01 and 10 will be offered the correct good by only one of the retailers, whereas consumers of type 11 will not make a purchase as they will not be offered the correct good. The expected profit of the first retailer will thus be $\alpha/2 + \gamma$, and of the second retailer $\alpha/2 + \beta$.

Does data-sharing improve retailers’ profits? Suppose retailers could share their respective information with one another, so that both always knew the consumer’s type. This would lead to both always offering the correct good, and hence to expected profits of $1/2$ for each. However, this would be detrimental to the second retailer whenever $\alpha/2 + \beta > 1/2$, since her profits with data-sharing would be lower than without. Such cooperation is thus not *individually rational* – the second retailer will not want to cooperate with the first.

Instead, suppose there is a *mediator*—a trusted third-party that may represent either a service provider or a cryptographic protocol (such as a federated learning algorithm) run by the two retailers. The European Commission calls mediators *common data spaces* (European Commission, 2020), and describes them as “third-parties [who] may act as intermediaries and apply new technologies and ways of organising markets in order to...enable transactions that were previously not feasible. They can be private, public or community organisations that are neutral with respect to data uses...” (Martens et al., 2020). Given such a mediator, consider the following data-sharing scheme: both retailers share their data with the mediator, who then passes along the data to all retailers *only* if the consumer’s type is 11. If the consumer is not of type 11 then each retailer has her original data, and additionally infers that the consumer’s type is not 11 (for otherwise she would have learned that the type is 11).

How does this scheme affect retailers’ strategies? Clearly, when retailers learn that the consumer’s type is 11, both offer g_0 . What happens if a retailer learns that the consumer’s

type is not 11? Consider the row retailer. If her bit is 0, then learning that the consumer's type is not 11 does not provide any new information, and so she still has the same dominant strategy of offering g_0 . If her bit is 1, however, then she learns that the consumer's type must be 10. But note that in this case, her best strategy is to offer g_1 , which is the same as her dominant strategy absent a mediator. Thus, the mediator changes the row retailer's behavior *only* when the consumer is of type 11. A similar analysis and conclusion hold for the column retailer.

Since retailers offer the same goods with and without data-sharing in all cases except when the consumer is of type 11, this scheme changes retailer's profits only in this latter case. In particular, it leads to an additional (ex ante) profit of $\delta/2$ for each retailer beyond her original profit—with probability δ the consumer is of type 11, in which case the retailers split the additional surplus of 1—and is thus beneficial to both. Observation 1 shows that such a scheme is, in fact, not only individually rational but also optimal.

Suppose now that the two retailers are also competing against an Amazon for the consumer. Since the Amazon is a giant it has more data, and, in particular, we assume that it has complete information about each consumer's type. In this three-way competition, the original dominant strategies of our two smaller retailers do not perform as well: They lead to profits of $\alpha/3 + \gamma/2$ and $\alpha/3 + \beta/2$ to the first and second retailer, respectively, since when they offer the correct good the consumer now chooses amongst up to 3 retailers. The mediator described above, which reveals information when the type is 11, leads to higher profits, since now there is an additional $\delta/3$ to each of the smaller retailers. However, that mediator is no longer optimal, and the retailers can actually do better.

To see this, observe that, conditional on consumer type 01, the total profit to the small retailers is $\beta/2$, since only the second makes the correct offer g_1 and then competes with the Amazon. In contrast, if both retailers were to know the type and offer the correct good, then their total profit would be $2\beta/3$, namely $\beta/3$ each. This is harmful to the second retailer as it involves a loss of profit, but a gain for the first and the sum. But the second retailer can be compensated by getting data from the first elsewhere, for instance on consumer type 10.

For simplicity of this example, suppose that $\beta = \gamma$, and consider a mediator that facilitates full data-sharing, in which both retailers learn each other's data and thus have complete information about consumers' types. Here, the profit of each is $1/3 > \alpha/3 + \beta/2 = \alpha/3 + \gamma/2$, so both gain from this data-sharing. Furthermore, the total profit of the retailers is $2/3$, which is the maximal utility they can obtain when competing against an Amazon. Thus, this data-sharing scheme is individually rational and optimal. Observe that it leads to higher retailer welfare than the optimal scheme absent a mediator.

Full data-sharing is not always individually rational, however. If β is much greater than γ , then the second retailer is not sufficiently compensated by the first for sharing data about consumer 01. The second retailer will consequently be harmed by such data-sharing, and so the scheme will not be individually rational. However, in Theorem 4 we show that there is a different data-sharing scheme that is individually rational and optimal, a scheme in which the first retailer shares all her data and the second shares some of her data. Overall, we show that such data-sharing coopetition against an Amazon is beneficial to the small retailers.

Contribution Our model is more general than the simple example above, and involves many goods and many types of consumers. For the example it was useful to think of each consumer as having a feature vector describing his type, but in the model we take a more general approach that allows for a wider class of information structures. Furthermore, in general the retailers will not have dominant strategies, either with or without a mediator. Instead, we suppose that, absent a mediator, players play an arbitrary Bayesian Nash equilibrium, and construct mediators that lead to higher expected utilities in equilibrium.

In Section 3, we design optimal mediators for cooperative data-sharing. In addition to maximizing total retailer profits, all our mediators are individually rational—retailers attain higher utilities when they use the mediators—and incentive compatible—in equilibrium, each retailer’s best strategy is to follow the mediators’ recommendations. One of our main results is that the optimal schemes are simple. In particular, instead of relying on the details of retailers’ data or on their strategies in equilibrium sans data-sharing (as in the introductory example above), our mediators consist of threshold rules that probabilistically induce either full data-sharing between retailers, or the full transfer of data from one retailer to another.

The analysis in Section 3 retains one of the assumptions present in the example above—that the joint data of the small retailers is sufficient to uniquely identify each consumer’s type. In Section 4 we drop this assumption, leading to a new set of challenges. This is because, when the assumption holds, there is no conflict between welfare maximization and equilibrium, and so the main hurdle in designing optimal mediators is the individual rationality constraint. Without the assumption, however, participants’ incentive compatibility constraints impose a limit on the total welfare that can be achieved. Nonetheless, our main results in this section are that variants of the simple mediators from Section 3 are approximately optimal here as well. We also provide examples showing that our approximation factor is tight.

Finally, in Section 5 we delve into the intriguing question of whether players share more data in the presence of an Amazon or in its absence. First, under the assumption that the joint data of the small retailers is sufficient to uniquely identify each consumer’s type, we show that if data-sharing is strictly beneficial in the absence of an Amazon, then it is also strictly beneficial in its presence. This confirms the intuition that players share more data when facing stronger outside competition. However, we also show that this conclusion may be reversed when the assumption does not hold. In this latter case, we show that sometimes data-sharing can be strictly beneficial in the absence of an Amazon, but *not* in its presence. We show that the reason this may happen depends on whether or not players’ equilibrium considerations conflict with welfare maximization.

The rest of the paper proceeds as follows. First we survey the related literature, and then, in Section 2, develop the formal model. This is followed by our main analyses in Sections 3, 4, and 5. Concluding notes appear in Section 6, where we also discuss extensions such as monetary transfers and more than two players.

Related literature Most broadly, our paper contributes to a burgeoning literature on competition in prediction and machine learning. Within this literature, papers such as Mansour et al. (2018), Ben-Porat and Tennenholtz (2019), and Feng et al. (2022) study different models of learning embedded in settings where participants compete with others in

their predictions. They take participants' data as given, and point to the effect competition has on optimal learning algorithms. Our focus, in contrast, is on the effect of data-sharing on competitive prediction.

Because of its focus on data-sharing, our paper is also related to the more established literature on strategic information-sharing, a literature that focuses on a number of distinct applications: oligopolistic competition (Clarke, 1983; Raith, 1996), financial intermediation (Pagano & Jappelli, 1993; Jappelli & Pagano, 2002; Gehrig & Stenbacka, 2007), supply chain management (Ha & Tong, 2008; Shamir & Shin, 2016), price discrimination (Liu & Serfes, 2006; Jentzsch et al., 2013), and competition between data brokers (Gu et al., 2019; Ichihashi, 2020). Much of this literature revolves around the question of whether it is beneficial for participating firms to pool all their data, or whether they would prefer to pool only some or none at all. As Bergemann and Morris (2013) demonstrate, however, more finely tuned data-sharing can be beneficial even when full or partial pooling is not, and so the existing schemes do not exhaust the potential benefits of data-sharing. In this paper we analyze precisely such finely tuned sharing.

A different paper that does consider more finely tuned sharing is that of de Zegher and Lo (2020), who study an infinitely repeated setting in which firms compete for market share. Their main result is the construction of a data-sharing scheme that Pareto improves all participants' welfare. Like our results, their scheme also only reveals partial information to participants. Unlike our approach, however, de Zegher and Lo's scheme relies on the repeated nature of the interaction, and uses folk-theorem-type arguments to show that cooperation can be sustained. In addition, their focus is on data-sharing between many small firms, and they do not consider the possible presence of an Amazon.

In terms of modeling and techniques, our paper falls into the literature on information design. Information design, recently surveyed by Bergemann and Morris (2019), is the study of how the allocation of information affects incentives and hence behavior, with a focus on the extent and limits of purely informational manipulation. The information design problem encompasses work on communication in games and on Bayesian persuasion. When the mediator is assumed to have only the information held by the players, or only the information they are willing to share with him, the problem maps to the one studied in the literature on communication in games. The goal of these studies is to characterize the equilibrium outcomes achievable when players are allowed to communicate prior to playing a fixed game, and where communication is captured by players' interaction with a mediator. Myerson (1991) and Forges (1993) provide useful overviews. Our paper builds on this model by endowing the mediator with the specific aim of maximizing (some of) the players' utilities.

A different setting, called Bayesian persuasion (Kamenica & Gentzkow, 2011), is one in which only the mediator (here called the sender) has payoff-relevant information, and in which he can commit to a particular information structure prior to observing that information. Initial work on Bayesian persuasion focused on the case of a single sender and a single player (Kamenica & Gentzkow, 2011), but more recent research also consists of settings with multiple senders (Gentzkow & Kamenica, 2016) and multiple players. Our paper is most closely related to the latter, and specifically to Bergemann and Morris (2016) and Galperti and Perego (2018), who develop a linear programming approach to study the effect of different information structures on the Bayes correlated equilibria of the subsequent game,

as well as Mathevet et al. (2020), who extend the belief-based approach of Kamenica and Gentzkow (2011) to study optimal information structures under different solution concepts.

A closely related strand of the literature, recently surveyed by Bergemann and Bonatti (2019), focuses on the sale of information by a data broker. Most of this work differs from our paper in that it focuses on a market with one-sided information flow, with information going only from one party to another, whereas we study a market where information flow is bidirectional across firms. Nonetheless, the insights arising from that literature are related to ones we develop in this paper. For example, Bimpikis et al. (2019) consider a monopolistic data broker who sells data to firms who compete in a downstream market. A major insight in that paper is that the amount of information optimally sold by the data broker depends on the nature of downstream competition, and, in particular, that more (resp., less) data is sold if that competition features strategic complements (resp., substitutes). While our paper studies data-sharing as opposed to data sale, our results also shed light on the effect the nature of competition has on the amount of data shared. In our model firms' actions are strategic substitutes—when more firms offer the correct good to a particular consumer (e.g., due to the presence of an Amazon), the value of offering the correct good decreases. Our results indicate that sometimes there is more data-sharing in the absence of an Amazon and hence when strategic substitutes are weaker, which is consistent with the insight of Bimpikis et al. (2019). However, we also show that sometimes the opposite holds.

There is also a relationship between information sale and information sharing, which is pointed out by Bergemann and Bonatti (2019). These authors relate their general model of a market for information to incentives for information sharing, and point to the study of finely tuned sharing schemes as an open problem.

Our paper is also related to a set of papers that focuses on designing mediators to achieve various goals, such as to improve the incentives of players, make equilibria robust to collusion, or implement correlated equilibria while guaranteeing privacy (Monderer & Tennenholtz, 2004, 2009; Kearns et al., 2014). The first two differ from our work in that they make stronger assumptions about the mediator's capabilities, such as changing payoffs or limiting player actions, and the third focuses on a setting with many players that is quite different from our own.

Finally, data-sharing among agents has been addressed by the multi-agent systems community (Dutta et al., 2005), and some of this work also considered the parties' incentives (Rochlin & Sarne, 2015; Chen et al., 2015). We continue this trend by addressing the fundamental data-sharing design setting in which users' data is distributed among competitors, an essential barrier we should cross for successful AI systems. Within this literature, our work is related to recent research on how competition affects predictions and data purchase (Ginart et al., 2021; Feng et al., 2022; Kwon et al., 2022). It is also conceptually related to research on federated learning, and in particular on cross-silo federated learning (Kairouz et al., 2019). This framework consists of a set of agents with individual data whose goal is to jointly compute a predictive model. A recent emphasis within federated learning is on incentivizing the agents to participate, sometimes through monetary transfers (Zhan et al., 2021) and sometimes by providing different models to different agents (Lyu et al., 2020). The focus of these papers is on fairness—agents that provide more data should be compensated more generously. Our paper differs, in that agents are assumed to utilize the resulting models in some competition, and this subsequent competition drives agents' incentives to

participate. Our work is thus orthogonal to that surveyed by Zhan et al. (2021): we do not focus on the computational aspects of deriving a joint model, but rather on the incentives and joint benefits of participating even when other agents are competitors.

2. Model and Preliminaries

There is a set G of goods and a population of consumers interested in obtaining one of them. Each consumer has one of a finite set of types, $\omega \in \Omega$, that describe the good $g_\omega \in G$ in which he is interested. There are three players who compete for consumers: two regular players indexed 1 and 2, and an Amazon, a player indexed 0. We will separate the analysis to two settings: first, when the Amazon player 0 is not present and players 1 and 2 compete only with one another; and second, when they additionally compete with Amazon. The model and definitions here apply to both settings.

Player 0 (if present) has complete information of the consumer’s type. In Section 3 we assume that players 1 and 2 have *jointly complete information*—that, when combined, their respective data uniquely identifies each consumer’s type—but that each player on her own may only be partially informed. In Section 4 we remove this assumption.

To model this informational setting, we represent players’ data using the information partition model of Aumann (1976): Each player i is endowed with a partition Π_i of Ω , where Π_i is a set of disjoint, nonempty sets whose union is Ω . For each $\omega \in \Omega$ we denote by $P_i(\omega)$ the unique element of Π_i that contains ω , with the interpretation that if the realized type of consumer is ω , each player i only learns that the type belongs to the set $P_i(\omega)$.

Framing the example from the introduction within this model would associate Ω with $\{0, 1\}^2$ and the partitions $P_1(00) = P_1(01) = \{00, 01\}$, $P_1(10) = P_1(11) = \{10, 11\}$, $P_2(00) = P_2(10) = \{00, 10\}$, and $P_2(01) = P_2(11) = \{01, 11\}$.

In this model, player 0’s complete information means that $P_0(\omega) = \{\omega\}$ for all $\omega \in \Omega$, and players 1 and 2’s jointly complete information means that $P_1(\omega) \cap P_2(\omega) = \{\omega\}$ for all $\omega \in \Omega$. We further assume that, before obtaining any information, all players have a common prior π over Ω .

To model data-sharing between players 1 and 2 we suppose there is a mediator that gathers each player’s information and shares it with the other in some way. Formally, a mediator is a function $\mathcal{M} : 2^\Omega \times 2^\Omega \mapsto \Delta(M^2)$, where M is an arbitrary message space. The range is a distribution over pairs of messages, where the first (respectively, second) is the message sent to player 1 (respectively, player 2).

We begin with an informal description of the game: Players offer consumers a good, and consumers choose a player from whom to acquire the good. Consumers are *single-minded*: For each ω there is a unique $g_\omega \in G$ such that the consumer will only choose a player who offers good g_ω . If there is more than one such player, the consumer chooses uniformly at random between them. One interpretation of this consumer behavior is that he chooses by “satisficing”—making a random choice among options that are “good enough” (Simon, 1956), where g_ω represents these good-enough options. See Ben-Porat and Tennenholtz (2019) for a similar approach.

We assume that prices and costs are fixed, and normalize a player’s utility to 1 if she is chosen and to 0 otherwise. So, for example, in the online advertising application of our model, consumers are assumed to click on a random ad amongst those ads that are most

relevant; once they click, the expected profit to the advertiser (conversion rate times profit from a sale) is 1. Given that players' payoffs are symmetric, the normalization to 1 is without loss of generality. If we had a different profit p_ω for each consumer type ω , then we could change every p_ω to 1, modify the prior probability $\pi(\omega)$ of type ω to $p_\omega \cdot \pi(\omega)$, and then renormalize the prior to be a proper distribution. This would leave all of our analysis intact.

The following is the order of events, given a fixed mediator \mathcal{M} . It applies both to the case in which there are only two players $\{1, 2\}$ and to the case in which there is an additional Amazon player, indexed 0.

1. Each of players 1 and 2 chooses whether or not to opt into using the mediator.
2. If one or both players $\{1, 2\}$ did not opt into using the mediator, then each player i obtains her *base value*, the utility v_i (described below).
3. If both players $\{1, 2\}$ opted into using the mediator, then:
 - (a) A consumer of type ω is chosen from Ω with prior distribution π .
 - (b) If present, player 0 learns $P_0(\omega)$.
 - (c) Messages (M_1, M_2) are chosen from the distribution $\mathcal{M}(P_1(\omega), P_2(\omega))$, and each $i \in \{1, 2\}$ learns their own $P_i(\omega)$ and M_i .
 - (d) Each player simultaneously chooses a good to offer the consumer, and then the consumer chooses a player from whom to obtain the good.

There are several ways to interpret the base values v_1 and v_2 . Our main interpretation is that these are the expected utilities of the players in the game without data-sharing. To formalize this, consider the *unmediated* game $\Gamma = (\Omega, \mathcal{P}, G^{|\mathcal{P}|}, (P_i(\cdot))_{i \in \mathcal{P}}, (u_i)_{i \in \mathcal{P}}, \pi)$, where \mathcal{P} is the set of participating players, and is either $\{1, 2\}$ or $\{0, 1, 2\}$, G is the set of actions, $P_i(\cdot)$ is the information of player i (which consists of the partition element of the realized type), u_i is i 's utility function (described below), and π is the common prior over Ω . Given this Bayesian game, let $v = (v_1, v_2)$ be the expected utilities in some *Bayesian Nash equilibrium (BNE)* of Γ —a profile (s_1, s_2) in which s_i is the best strategy for each player i conditional on her information and the assumption that the other player plays the strategy s_j .

Although this is the main interpretation of the base values, our model and results permit additional and more general interpretations as well. One additional interpretation is to suppose each v_i is the minimax value of player i in the unmediated game. In this second interpretation, we could imagine player j “punishing” player i if the latter does not opt into using the mediator, by playing a strategy that minimizes the latter's payoff (off the equilibrium path). A third, more general interpretation is that the base values also depend on factors outside of the specific game, such as firm size, customer base, and so on, in addition to the primitives of the game Γ .

In our constructions of mediators in Sections 3 and 4, the interpretation of v will not matter—our mediators will be optimal given any such values, regardless of whether they are derived endogenously as the equilibrium utilities of the game absent a mediator, or whether they arise exogenously from factors outside of the specific game. However, in

Section 5, in which we compare data-sharing across two different environments—with and without a mediator—we will stick with the main interpretation of base values as equilibrium payoffs absent a mediator. This is because, when we compare two different environments, the payoffs players could attain in equilibrium without a mediator will differ across these environments.

Now, as noted above, data-sharing is modeled by a mediator $\mathcal{M} : 2^\Omega \times 2^\Omega \mapsto \Delta(M^2)$. Without loss of generality we invoke the revelation principle and assume that $M = G$, the set of possible goods. We interpret the messages of the mediator as recommended actions to the players, one recommendation for each player.

Formally, if both players 1 and 2 opt into using the mediator \mathcal{M} , then all play the mediated Bayesian game $\Gamma^{\mathcal{M}} = (\Omega, \mathcal{P}, G^{|\mathcal{P}|}, \mathcal{I}^{\mathcal{M}}, (u_i)_{i \in \mathcal{P}}, \pi)$. \mathcal{P} is the set of participating players, and is either $\{1, 2\}$ or $\{0, 1, 2\}$. The function $\mathcal{I}^{\mathcal{M}} : \Omega \mapsto \Delta(T^{|\mathcal{P}|})$ denotes the information of players in the game for each state, where $T = 2^\Omega \times M$ is the set of possible pieces of information a player may have—a partition element and a message from the mediator. For player 0, this information consists of the realized type of consumer only, and so $\mathcal{I}^{\mathcal{M}}(\cdot)_0 = (P_0(\cdot), \emptyset)$ (where the \emptyset means player 0 gets no message from the mediator). For players 1 and 2, the information consists of both the partition element of the realized type of consumer and the action recommended to her by the mediator, and so $\mathcal{I}^{\mathcal{M}}(\cdot)_i = (P_i(\cdot), \mathcal{M}(P_1(\cdot), P_2(\cdot))_i)$.

Next, each player’s set of actions in the mediated game is G , and her utility function $u_i = u_i : \Omega \times G^{|\mathcal{P}|} \mapsto \mathbb{R}$. The latter is equal to 0 if player i ’s action $g \neq g_\omega$, and otherwise it is equal to $1/k$, where k is the total number of players who play action g_ω . The nonzero utility corresponds to utility 1 if a consumer chooses the player’s good, which occurs if a player offers the consumer the correct good and the consumer chooses uniformly amongst all players that do so. We often write $E[u_i(\cdot)]$ when the expectation is over the choice of ω , in which case we omit the dependence of u_i on ω for brevity. Finally, (mixed) strategies of players in Γ are functions $s_i : T \mapsto \Delta(G)$. We also denote by $s_i(\omega) = s_i(\mathcal{I}_i(\omega))$.

An important note about the mediator is in order. We assume that when players opt into participating, they *truthfully* reveal all their data $P_i(\omega)$ to him. Players’ strategic behavior is relevant in their choice of opting in or not, and then in whether or not they follow the mediator’s recommendation in their interaction with the consumer. While truthful reporting is clearly restrictive, it is a natural assumption in our context—for instance, when service providers act as mediators, they are typically also the ones who host retailers’ data in the cloud, and so already have the (true) data on their servers. Similarly, federated machine learning algorithms work based on the assumption that participants share their true data. Furthermore, truthful reporting can be justified by the repeated nature of the interaction. Observe that if a firm is not truthful in its reporting, the other firm will discover this at some point, as the mediator’s recommendations will not pan out as well as they should. Thus, if firms have the ability to opt out of using the mediator at some point in the future, then each firm can use the threat of opting out to incentivize the other to report truthfully.²

Next, we define three important properties of mediators, namely, individual rationality, incentive compatibility, and optimality. A first observation is that player 0, if present, has a dominant strategy:

2. It is also an interesting question to determine the minimal powers a mediator requires in order to incentivize truthful reporting. We discuss this briefly in the conclusion.

Lemma 1 *In any game $\Gamma^{\mathcal{M}}$, player 0 has the unique dominant strategy $s_0(\omega) = g_\omega$.*

The proof is straightforward: playing action g_ω leads to positive utility to player 0, regardless of the other firms' actions. Playing any other action leads to utility zero. This implies that always offering the correct good is the unique dominant strategy for player 0. We will thus take that player's strategy as fixed throughout.

We will be interested in designing mediators that players actually wish to utilize, and so would like them to satisfy two requirements: first, that players want to follow the mediator's recommendations, and so that following these recommendations forms an equilibrium; and second, that players prefer their payoffs with the mediator over their payoffs without (the base values). Formally, we say that:

- \mathcal{M} is *incentive compatible (IC)* if the strategy profile \bar{s} in which players 1 and 2 always follow \mathcal{M} 's recommendation is a BNE of $\Gamma^{\mathcal{M}}$.
- \mathcal{M} is *individually rational (IR)* if the expected utility of each player $i \in \{1, 2\}$ under \bar{s} in $\Gamma^{\mathcal{M}}$ is (weakly) greater than v_i .

Observe that if a mediator is both IC and IR, then in equilibrium players will opt in and then follow the mediator's recommendations.

Finally, as we are interested in the extent to which data-sharing benefits players 1 and 2, we will study *optimal* mediators, namely, ones in which the sum of these two players' utilities is maximal subject to the IC and IR constraints. To this end, denote by $W(\mathcal{M})$ the sum of players 1 and 2's expected utilities under \bar{s} in $\Gamma^{\mathcal{M}}$. Then:

Definition 1 *Mediator \mathcal{M} is optimal if $W(\mathcal{M}) \geq W(\mathcal{M}')$ for any other \mathcal{M}' that is IC and IR.*

3. Jointly Complete Information

3.1 Competition Without an Amazon

We begin our analysis with the case of jointly complete information and without an Amazon. We first consider the following question: Under what base values v does there exist an IR, IC mediator? We then turn to our main result for the section, the construction of an optimal mediator.

3.1.1 WHEN DOES AN IR, IC MEDIATOR EXIST?

We begin with a simple observation. Since for any mediator \mathcal{M} and conditional on any realized type ω the total payoff to players is at most 1, it must be the case that $W(\mathcal{M}) \leq 1$. Thus, if $v_1 + v_2 > 1$, then there does not exist any IR mediator.

However, even if $v_1 + v_2 \leq 1$, there may not exist a mediator that is both IR and IC. In the example from the introduction, for instance, in the simple case without an Amazon and where $\alpha > 2\beta \geq 2\gamma > 4\delta$, the row retailer can guarantee herself a payoff of $(\alpha + \gamma)/2$ by ignoring the mediator and playing her dominant actions. This implies that the most the column retailer can obtain is $1 - (\alpha + \gamma)/2$. Any mediator that leads to a higher payoff to the column retailer must therefore fail to satisfy the IC constraints.

Theorem 1 below generalizes the bound in this example. But first, some notation: Consider the strategy s'_j of player j that, for every $P_j(\omega)$, chooses one of the goods that is most likely to be correct: $s'_j(\omega) \in \arg \max_{g \in G} \Pr[g = g_\omega | P_j(\omega)]$. Note that s'_j is a best-response to a player $i \neq j$ who always chooses the correct good g_ω . Furthermore, let α_j be the overall probability that s'_j chooses the correct good: $\alpha_j = \Pr[s'_j(\omega) = g_\omega]$.

Theorem 1 *Suppose $v_i \geq v_j$. Then an IR, IC mediator exists only if $v_1 + v_2 \leq 1$ and $v_i \leq E[u_i(g_\omega, s'_j(\omega))]$.*

In words, the theorem states that an IR, IC mediator exists only if each player's base value is bounded above by her utility in the hypothetical situation in which she always offers the correct good, and the other player best-responds conditional only on $P_j(\omega)$ (in addition to the obvious constraint $v_1 + v_2 \leq 1$). In Theorem 2 below we will see that these conditions are also sufficient.

We note that if the base values v are the equilibrium expected utilities of the game without a mediator, then the conditions in Theorem 1 are always satisfied. To see this, observe that the mediator that recommends to each player her equilibrium action (and nothing more) is IC, by the definition of a BNE. It is also IR, since players' utilities are the same whether or not they opt into using this mediator.

Proof: The condition $v_1 + v_2 \leq 1$ must hold for any IR mediator \mathcal{M} , since $W(\mathcal{M}) \leq 1$. Now, suppose towards a contradiction that there exists an IC mediator for which $v_i > E[u_i(g_\omega, s'_j(\omega))]$. This implies that, under \mathcal{M} , the expected utility of player $j = 3 - i$ is strictly less than $1 - E[u_i(g_\omega, s'_j(\omega))]$. However, player j then has a profitable deviation to s'_j . When playing s'_j , on each partition element $P_j(\omega)$ player j will get utility *at least*

$$\Pr[s'_j(\omega) = g_\omega | P_j(\omega)] / 2 = E[u_j(s'_j(\omega), g_\omega) | P_j(\omega)] = 1 - E[u_i(g_\omega, s'_j(\omega)) | P_j(\omega)],$$

and exactly this utility whenever i always chooses the correct good g_ω . Since this is a profitable deviation, the mediator \mathcal{M} is not IC, a contradiction. ■

For the rest of Section 3.1, denote a pair v that satisfies the conditions in Theorem 1 as a *feasible v* .

3.1.2 AN OPTIMAL MEDIATOR

We now describe an optimal mediator, at the same time showing that the conditions on v in Theorem 1 are also sufficient. We begin with a simple observation:

Observation 1 If \mathcal{M} is IR, IC, and $W(\mathcal{M}) = 1$, then \mathcal{M} is optimal.

This observation implies that the mediator in the example of the introduction, for the case without an Amazon, is optimal (when the base values are retailers' utilities under their dominant strategies). One drawback of that mediator, however, is that it depends on the retailers' strategies in the unmediated game—in particular, it utilizes the fact that, in equilibrium, neither retailer offers the correct good in the bottom-right cell of the matrix.

Our aim here is to design a mediator that does not depend on such strategies, but only on players' respective information and base values.

The idea underlying the mediator we construct is the following. Suppose $v_i \geq v_j$. Then if $v_i \leq 1/2$, then the mediator will facilitate full data-sharing, and so each player will offer the correct good in every state. If $v_i > 1/2$ then player i will obtain all the information and player j will obtain only partial information: she will get recommendation g_ω with probability less than 1, and otherwise will get recommendation $s'_j(\omega)$. We now formally describe the mediator, and then show that it is indeed IR, IC, and optimal.

Mediator 1 Mediator for $\mathcal{P} = \{1, 2\}$ with jointly complete information

1: **procedure** $\mathcal{M}_{noA}^v(V, W)$ $\triangleright V = P_1(\omega)$ and $W = P_2(\omega)$ for realized $\omega \in \Omega$
2: $\omega \leftarrow V \cap W$
3: $i \leftarrow \arg \max_{k \in \{1, 2\}} v_k$
4: $j \leftarrow 3 - i$
5: **if** $v_i \leq 1/2$ **then return** (g_ω, g_ω) \triangleright Full data-sharing
6: **else**
7: $g_j \leftarrow s'_j(\omega)$
8: Choose $\gamma \in [0, 1]$ uniformly at random.
9: **if** $\gamma < \frac{2-2v_i-\alpha_j}{1-\alpha_j}$ **then**
10: **return** (g_ω, g_ω)
11: **else if** $i = 1$ **then return** (g_ω, g_j)
12: **else return** (g_j, g_ω)

Theorem 2 For any feasible v , the mediator \mathcal{M}_{noA}^v above is IR, IC, and optimal.

Proof: First, consider the case in which $v_i \leq 1/2$. In this case, strategy \bar{s} in $\Gamma^{\mathcal{M}_{noA}^v}$ leads to expected utility $1/2$ to each player, and so the mediator is IR. Furthermore, the strategy \bar{s} is a BNE, since each player always chooses g_ω , which is dominant, and so the mediator is also IC. Finally, since in this case $W(\Gamma^{\mathcal{M}_{noA}^v}) = 1$, Observation 1 implies that the mediator is optimal.

Next, suppose $v_i > 1/2$. Observe that, conditional on line 10 being activated, the expected utility of player j is $1/2$, whereas conditional on line 11 or 12, the expected utility of player j is $\alpha_j/2$ (since with probability α_j she chooses the correct good, and player i always chooses the correct good). Thus, overall the expected utility of player j is

$$\frac{2 - 2v_i - \alpha_j}{1 - \alpha_j} \cdot \frac{1}{2} + \left(1 - \frac{2 - 2v_i - \alpha_j}{1 - \alpha_j}\right) \cdot \frac{\alpha_j}{2} = 1 - v_i.$$

Since player i always chooses the correct good, it must be the case that $W(\Gamma^{\mathcal{M}_{noA}^v}) = 1$, and so the utility of player i is v_i . Finally, since $v_1 + v_2 \leq 1$, this implies that the utility of player j is at least v_j . Thus, the mediator is IR.

Next, observe that \bar{s} in $\Gamma^{\mathcal{M}_{noA}^v}$ is dominant for player i , since she always chooses g_ω . It is also optimal for player j , since she either also chooses g_ω , or chooses the best response conditional on her information $P_j(\omega)$. Thus, the strategy \bar{s} is a BNE, and so the mediator is IC.

Finally, since $W(\Gamma^{\mathcal{M}_{noA}^v}) = 1$, Observation 1 implies that the mediator is optimal. ■

3.2 Coopetition Against an Amazon

In this section we analyze the game with an Amazon. As in Section 3.1, we first consider the following question: Under what base values v does there exist an IR, IC mediator? We then turn to our main result for the section, the construction of an optimal mediator.

3.2.1 WHEN DOES AN IR, IC MEDIATOR EXIST?

Once again, we begin with a simple observation. Since for any mediator \mathcal{M} and any type ω the total welfare of players 1 and 2 is at most $2/3$ (since Amazon always knows the type, offers the correct good, and gets at least $1/3$ of the surplus), it must be the case that $W(\mathcal{M}) \leq 2/3$. Thus, if $v_1 + v_2 > 2/3$, then there does not exist any IR mediator.

However, and again as above, even if $v_1 + v_2 \leq 2/3$ there may still not exist a mediator that is both IR and IC. In the example from the introduction, for instance, in the simple case with an Amazon and where $\alpha > 2\beta \geq 2\gamma > 4\delta$, the row retailer can guarantee herself a payoff of $(\alpha + \gamma)/3$ by ignoring the mediator and playing her dominant actions. This implies that the most the column retailer can obtain is $2/3 - (\alpha + \gamma)/3$. Any mediator that leads to a higher payoff to the column retailer must therefore fail to satisfy the IC constraints.

The result below consists of bounds on the base values under which an IR, IC mediator exists.

Theorem 3 *Suppose $v_i \geq v_j$. Then an IR, IC mediator exists only if $v_i \leq E \left[u_i(g_\omega, s'_j(\omega)) \right]$ and $v_j \leq 1 - 2v_i$.*

As with Theorem 1, we note that if the base values v are the equilibrium expected utilities of the game without a mediator, then the conditions in Theorem 3 are satisfied. Again, this follows from the observation that the mediator that recommends to each player her equilibrium action is both IR and IC.

Before proving the theorem we develop some notation and intuition. For any mediator \mathcal{M} and player $k \in \{1, 2\}$ let $\beta_k^{\mathcal{M}} = \Pr[(\mathcal{M}(\omega)_k = g_\omega) \cap (\mathcal{M}(\omega)_{3-k} \neq g_\omega)]$ and $\beta^{\mathcal{M}} = \Pr[\mathcal{M}(\omega)_k = \mathcal{M}(\omega)_{3-k} = g_\omega]$. With this notation, the expected utility of player k under \bar{s} in $\Gamma^{\mathcal{M}}$ is $\beta_k^{\mathcal{M}}/2 + \beta^{\mathcal{M}}/3$: For any ω , if player k offers g_ω and the other player, $3 - k$, does not, then the consumer chooses between players k and 0, and so each gets $1/2$. If both offer g_ω , the consumer chooses between players k , $3 - k$, and 0, and so each gets $1/3$.

An IR, IC mediator \mathcal{M} is one for which \bar{s} is a BNE, and in which $\beta_k^{\mathcal{M}}/2 + \beta^{\mathcal{M}}/3 \geq v_k$ for both players. In addition to the incentive constraints, we also have the constraints $\beta_i^{\mathcal{M}} + \beta_j^{\mathcal{M}} + \beta^{\mathcal{M}} \leq 1$, and all three variables non-negative.

Finally, the total welfare of players 1 and 2 in \mathcal{M} is $W(\mathcal{M}) = (\beta_1^{\mathcal{M}} + \beta_2^{\mathcal{M}})/2 + 2\beta^{\mathcal{M}}/3$, and so an optimal mediator is one that maximizes this sum subject to the IC and IR constraints above. Denote the linear program above, disregarding the IC constraint, as $\text{LP}(v_1, v_2)$:

$$\begin{aligned} & \text{maximize} && \frac{\beta_1 + \beta_2}{2} + \frac{2\beta}{3} \\ & \text{subject to} && \frac{\beta_i}{2} + \frac{\beta}{3} \geq v_i, \quad i = 1, 2 \\ & && \beta_i \geq 0, \quad i = 1, 2 \\ & && \beta \geq 0 \\ & && \beta_1 + \beta_2 + \beta \leq 1 \end{aligned}$$

To prove Theorem 3, we need to identify conditions on v under which the LP is feasible. To do this, it will be helpful to begin with some lemmas that characterize the LP's optimal solution:

Lemma 2 *Suppose $v_i \geq v_j$ and that $\text{LP}(v_1, v_2)$ is feasible. Then any optimal solution $(\beta_1, \beta_2, \beta)$ to $\text{LP}(v_1, v_2)$ satisfies:*

- (a) $\beta_1 + \beta_2 + \beta = 1$.
- (b) $\beta_j = 0$.
- (c) *If $v_i > 1/3$ then i 's IR constraint binds: $\beta_i/2 + \beta/3 = v_i$.*

Proof: Suppose (a) does not hold, and that $\beta_1 + \beta_2 + \beta < 1$. This means we can increase β to $\beta' > \beta$ without violating any of the constraints. The value of the LP is now $(\beta_1 + \beta_2)/2 + 2\beta'/3 > (\beta_1 + \beta_2)/2 + 2\beta/3$, contradicting the optimality $(\beta_1, \beta_2, \beta)$.

Now suppose (a) holds, but (b) does not. If $\beta_i < \beta_j$ then the solution $(\beta_2, \beta_1, \beta)$ is also feasible and optimal, so assume without loss of generality that $\beta_i \geq \beta_j$. Consider now the solution $\beta'_i = \beta_i - \beta_j$, $\beta'_j = 0$, and $\beta' = \beta + 2\beta_j$. Since

$$\beta'_i/2 + \beta'/3 = (\beta_i - \beta_j)/2 + (\beta + 2\beta_j)/3 > \beta_i/2 + \beta/3 \geq v_i$$

and

$$\beta'_j/2 + \beta'/3 = (\beta + 2\beta_j)/3 > \beta_j/2 + \beta/3 \geq v_j,$$

this new solution is feasible. Moreover,

$$(\beta'_1 + \beta'_2)/2 + 2\beta'/3 > (\beta_1 + \beta_2)/2 + 2\beta/3,$$

contradicting the optimality of $(\beta_1, \beta_2, \beta)$.

Now suppose (a) and (b) hold, but that (c) does not. Then there is some $\varepsilon > 0$ such that $\beta'_i/2 + \beta'/3 > v_i$, where $\beta'_i = \beta_i - \varepsilon$ and $\beta' = \beta + \varepsilon$, and so $(\beta'_1, \beta'_2, \beta')$ is also a feasible solution. Furthermore,

$$\beta'_i/2 + 2\beta'/3 > \beta_i/2 + 2\beta/3,$$

contradicting the optimality of $(\beta_1, \beta_2, \beta)$. ■

One implication of Lemma 2 is that at an optimal solution, β is maximized. This follows from the observation that the objective function is $(\beta_1 + \beta_2 + \beta)/2 + \beta/6 = 1/2 + \beta/6$ when $\beta_1 + \beta_2 + \beta = 1$. In addition, Lemma 2 pins down the value of the optimal solution to $\text{LP}(v_1, v_2)$:

Lemma 3 *Suppose $v_i \geq v_j$ and that $\text{LP}(v_1, v_2)$ is feasible. Then the value $\text{opt}(v_1, v_2)$ of an optimal solution to $\text{LP}(v_1, v_2)$ is one of the following:*

- *If $v_i \leq 1/3$ then $\text{opt}(v_1, v_2) = 2/3$.*
- *If $v_i > 1/3$ then $\text{opt}(v_1, v_2) = 1 - v_i$.*

Proof: Lemma 2 implies that at the optimal solution we have $\beta_i + \beta = 1$, and so the value at the optimum is $\beta_i/2 + 2\beta/3$. In the first case, the optimal solution occurs at the point $\beta_1 = \beta_2 = 0$, $\beta = 1$: it maximizes the objective, while both IR constraints remain slack since $v_i \leq 1/3 = \beta/3$.

In the second case, Lemma 2 implies that $\beta_i/2 + \beta/3 = v_i$ and that $\beta_i = 1 - \beta$. These imply that $v_i = (1 - \beta)/2 + \beta/3$, and so that $\beta = 3 - 6v_i$ and $\beta_i = 1 - 3 + 6v_i$. Plugging these into the objective yields

$$\text{opt}(v_1, v_2) = \frac{\beta_i}{2} + \frac{2\beta}{3} = \frac{1 - 3 + 6v_i}{2} + \frac{2(3 - 6v_i)}{3} = 1 - v_i. \quad \blacksquare$$

Lemmas 2 and 3 characterize the maximal values attainable by an IR mediator, and are the analogues of the condition that $v_1 + v_2 \leq 1$ when there is *no* Amazon. However, before proving Theorem 3, we need two additional lemmas. Fix a mediator \mathcal{M} , and consider the hypothetical situation in which player i obtains additional information, and in each state ω she fully learns that the state is ω and so plays the dominant strategy $\hat{s}_i(\omega) = g_\omega$. In this new game, denote the best response of player j as $s'_j(\mathcal{M})$, the strategy that, for every $P_j(\omega)$ and $\mathcal{M}(\omega)_i$, chooses some g that is most-likely correct conditional on $P_j(\omega)$ and $\mathcal{M}(\omega)_i$. Also observe that here, player j is never the only player to choose g_ω , since player i always plays g_ω .

Player j is worse off in this hypothetical game than when all follow \mathcal{M} 's recommendations:

Lemma 4 $E[u_j(s'_j(\mathcal{M}), \hat{s}_i)] \leq E[u_j(\bar{s}_j, \bar{s}_i)]$ for any mediator \mathcal{M} .

Proof: Fix a particular state ω . For every action g of player j , who conditions on $P_j(\omega)$ and $\mathcal{M}(\omega)_j$, the utility $E[u_j(g, \hat{s}_i)|P_j(\omega), \mathcal{M}(\omega)_j] \leq E[u_j(g, \bar{s}_i)|P_j(\omega), \mathcal{M}(\omega)_j]$, since i makes correct choices more often in the former. This implies that $E[u_j(s'_j(\mathcal{M}), \hat{s}_i)|P_j(\omega), \mathcal{M}(\omega)_j] \leq E[u_j(\bar{s}_j, \bar{s}_i)|P_j(\omega), \mathcal{M}(\omega)_j]$, since in both sides of the inequality j chooses an optimal action conditional on her information, but in the RHS the utility from every choice g is higher than its corresponding utility in the LHS. This implies that $E[u_j(s'_j(\mathcal{M}), \hat{s}_i)] \leq E[u_j(\bar{s}_j, \bar{s}_i)]$. \blacksquare

Although in the hypothetical game player j is worse off, the sum of players' utilities is higher:

Lemma 5 *The sum of players' utilities is higher under $(s'_j(\mathcal{M}), \hat{s}_i)$ than under (\bar{s}_j, \bar{s}_i) in \mathcal{M} .*

Proof: Again, we show that this holds conditional on every $P_j(\omega)$ and $\mathcal{M}(\omega)_j$. Fix ω , and observe that

$$\Pr[\hat{s}_i(\omega) = \bar{s}_j(\omega) = g_\omega | P_j(\omega), \mathcal{M}(\omega)_j] \geq \Pr[\bar{s}_i(\omega) = \bar{s}_j(\omega) = g_\omega | P_j(\omega), \mathcal{M}(\omega)_j],$$

since \hat{s}_i always correctly chooses g_ω . Furthermore,

$$\Pr[\hat{s}_i(\omega) = s'_j(\mathcal{M}) = g_\omega | P_j(\omega), \mathcal{M}(\omega)_j] \geq \Pr[\hat{s}_i(\omega) = \bar{s}_j(\omega) = g_\omega | P_j(\omega), \mathcal{M}(\omega)_j],$$

since s'_j chooses the good that maximizes the probability of choosing g_ω . Thus,

$$\Pr [\hat{s}_i(\omega) = s'_j(\mathcal{M}) = g_\omega | P_j(\omega), \mathcal{M}(\omega)_j] \geq \Pr [\bar{s}_i(\omega) = \bar{s}_j(\omega) = g_\omega | P_j(\omega), \mathcal{M}(\omega)_j].$$

Now recall that $\text{LP}(v_1, v_2)$ is maximized when β , the probability that both players correctly choose g_ω is maximized. Thus, the value of $\text{LP}(v_1, v_2)$, the sum of players' utilities, is higher under (s'_j, \hat{s}_i) than under (\bar{s}_1, \bar{s}_2) . ■

We can now prove Theorem 3.

Proof: For the first condition, the upper bound on v_i , fix some IR, IC mediator \mathcal{M} . We will argue that the utility to player i under \mathcal{M} is at most $E [u_i(g_\omega, s'_j(\mathcal{M}))]$. Lemma 4 implies that player j 's utility decreases if player i additionally obtains all the information (even when j responds optimally), and Lemma 5 implies that in this same situation the sum of utilities increases. Together, these lemmas imply that i 's utility is higher when she obtains all the information and j only learns $P_j(\omega)$ and $\mathcal{M}(\omega)_j$, than when i also only learns $P_i(\omega)$ and $\mathcal{M}(\omega)_i$. Furthermore, when i has all the information, she is best off when j has the least information—that is, when $\mathcal{M}(\omega)_j = P_j(\omega)$. This is because any additional information that causes j to improve upon her payoff will do so at the expense of player i 's payoff (by transferring weight from $\beta_i^{\mathcal{M}}$ to $\beta^{\mathcal{M}}$). Thus, in any IR, IC mediator, the highest possible utility to player i is $E [u_i(g_\omega, s'_j(\omega))]$, since $\mathcal{M}(\omega)_j = P_j(\omega)$ and so $s'_j(\omega) = s'_j(\mathcal{M})$.

The second condition, $v_j \leq 1 - 2v_i$, follows from Lemma 3. If $v_i \leq 1/3$, then the maximal welfare of any IR mediator is $2/3$, and so $v_j \leq 2/3 - v_i \leq 1 - 2v_i$. If $v_i > 1/3$, then the lemma implies that the maximal welfare of any IR mediator is at most $1 - v_i$. Since player i gets at least v_i , player j can get at most $1 - 2v_i$. ■

For the rest of Section 3.2, denote a pair v that satisfies the conditions in Theorem 3 as a *feasible* v .

3.2.2 AN OPTIMAL MEDIATOR

We now describe an optimal mediator, at the same time showing that the conditions in Theorem 3 are not only necessary but also sufficient.

We begin with a definition.

Definition 2 *A mediator \mathcal{M} is fully revealing to player i if the mediator's recommendation to player i always coincides with the optimal good: $\bar{s}_i(\omega) = g_\omega$ for every ω . A mediator facilitates full data-sharing if it is fully revealing to both players 1 and 2.*

Fix the set of players to be $\mathcal{P} = \{0, 1, 2\}$. The idea underlying the mediator is the same as that of \mathcal{M}_{noA}^v . Suppose $v_i \geq v_j$. Then if $v_i \leq 1/3$, then the mediator will facilitate full data-sharing, and so each player will offer the correct good in every state. If $v_i > 1/3$ then player i will obtain all the information and player j will obtain only partial information: she will get recommendation g_ω with probability less than 1, and otherwise will get recommendation $s'_j(\omega)$. The challenge is to maximize the probability of recommendation g_ω subject to the IR constraints, and to do this in such a way that following the recommendations is a BNE.

We now formally describe the mediator, and then show that it is indeed IR, IC, and optimal. The main differences from \mathcal{M}_{noA}^v in Mediator 1 are lines 5 and 9.

Mediator 2 Mediator for $\mathcal{P} = \{0, 1, 2\}$ with jointly complete information

1:	procedure $\mathcal{M}_A^v(V, W)$	$\triangleright V = P_1(\omega)$ and $W = P_2(\omega)$ for realized $\omega \in \Omega$
2:	$\omega \leftarrow V \cap W$	
3:	$i \leftarrow \arg \max_{k \in \{1, 2\}} v_k$	
4:	$j \leftarrow 3 - i$	
5:	if $v_i \leq 1/3$ then return (g_ω, g_ω)	\triangleright Full data-sharing
6:	else	
7:	$g_j \leftarrow s'_j(\omega)$	
8:	Choose $\gamma \in [0, 1]$ uniformly at random.	
9:	if $\gamma < \frac{3-6v_i-\alpha_j}{1-\alpha_j}$ then	
10:	return (g_ω, g_ω)	
11:	else if $i = 1$ then return (g_ω, g_j)	
12:	else return (g_j, g_ω)	

Theorem 4 For any feasible v , the mediator \mathcal{M}_A^v above is IR, IC, and optimal.

An interesting feature of \mathcal{M}_A^v is that it either facilitates full data-sharing, or is fully revealing to one of the players. These features are inherent to *any* optimal IR mediator:

Theorem 5 For any feasible v and any optimal IR mediator \mathcal{M} , one of the following holds:

- \mathcal{M} is fully revealing to both players 1 and 2.
- \mathcal{M} is fully revealing to one of the players, and the IR constraint binds for that player (i.e., she is indifferent between \mathcal{M} and her base value).

Note that Mediator 1 for the case of no Amazon also satisfies the features described in Theorem 5. In that case, however, these features do not characterize all mediators, as there exist mediators that are optimal for the case of no Amazon that do not have them. One such mediator is the one described in the example from the introduction. With an Amazon, in contrast, these features characterize all optimal IR mediators.

We begin by showing that the mediator above is optimal. Our approach is to show that $W(\mathcal{M}_A^v)$ is equal to the value of the optimal solution to $\text{LP}(v_1, v_2)$, and that the strategy \bar{s} in \mathcal{M}_A^v is a BNE. Observe that $\text{LP}(v_1, v_2)$ does not contain the IC constraints, implying that they are not binding, and so only the IR constraints are relevant for optimality. This will no longer be the case without the assumption of jointly complete information, analyzed in Section 4 below.

We already laid most of the groundwork for the proofs of Theorems 4 and 5 when we proved Theorem 3, but do need one last lemma:

Lemma 6 Suppose v is feasible and $v_i > 1/3$. Then $\alpha_j = \Pr [s'_j(\omega) = g_\omega] \leq 3 - 6v_i$.

Proof: Since v is feasible, it holds that $E [u_i(g_\omega, s'_j(\omega))] \geq v_i$. Next, note that $E[u_i(g_\omega, s'_j(\omega))] = \alpha_j/3 + (1 - \alpha_j)/2$. Together, these facts imply that $\alpha_j \leq 3 - 6v_i$, as claimed. ■

We are now ready to prove Theorem 4.

Proof: First, consider the case in which $v_i \leq 1/3$. In this case, strategy \bar{s} in $\Gamma^{\mathcal{M}_A^v}$ leads to expected utility $1/3$ to each player, which satisfies the IR constraint. Furthermore, strategy \bar{s} is a BNE, since each player always chooses g_ω , which is dominant. Finally, $W(\mathcal{M}_A^v)$ is maximal by Lemma 3.

Next, suppose $v_i > 1/3$. Observe that \bar{s} in $\Gamma^{\mathcal{M}_A^v}$ is dominant for player i , since she always chooses g_ω . It is also optimal for player j , since she either also chooses g_ω , or chooses the best response conditional on her information $P_j(\omega)$. Thus, the mediator is IC.

We now show that \mathcal{M}_A^v is optimal. For simplicity, henceforth denote by $\mathcal{M} = \mathcal{M}_A^v$. First note that $\beta_j^{\mathcal{M}} = 0$. Next, by Lemma 6, it holds that $\alpha_j \leq 3 - 6v_i$. Since $v_i \in (1/3, 1/2]$, we have that $1 - \alpha_j \in [0, 1)$ and so $0 \leq 3 - 6v_i - \alpha_j < 1 - \alpha_j$, implying that line 9 of Mediator 2 is valid. From that line we have that

$$\beta^{\mathcal{M}} = \alpha_j + (1 - \alpha_j) \cdot \frac{3 - 6v_i - \alpha_j}{1 - \alpha_j} = 3 - 6v_i,$$

since both players choose g_ω when $s'_j(\omega) = g_\omega$ and when $s'_j(\omega) \neq g_\omega$ but $\gamma < (3 - 6v_i - \alpha_j)/(1 - \alpha_j)$. This means that

$$W(\mathcal{M}) = \frac{\beta_i^{\mathcal{M}}}{2} + \frac{2\beta^{\mathcal{M}}}{3} = \frac{1 - 3 + 6v_i}{2} + \frac{2(3 - 6v_i)}{3} = 1 - v_i,$$

which, by Lemma 3, is the value of the optimal solution to $\text{LP}(v_1, v_2)$.

Finally, since i 's utility is v_i and so j 's utility is $1 - 2v_i$, the mediator is IR. ■

Finally, we can also prove Theorem 5.

Proof: Suppose first that v is such that $v_j \leq v_i \leq 1/3$. In this case, full data-sharing is IC, and so any optimal mediator \mathcal{M} must have $W(\mathcal{M}) \geq 2/3$. However, the maximum value of $\text{LP}(v_1, v_2)$ is $2/3$, and the only way to obtain this is to have $\beta^{\mathcal{M}} = 1$. This is equivalent to full data-sharing, and so the only mediator that is optimal is the full data-sharing one.

Now suppose that $v_i > 1/3$. By the proof of Theorem 4, an optimal solution to $\text{LP}(v_1, v_2)$ is attainable by some mediator, in particular by \mathcal{M}_A^v . Thus, any mediator \mathcal{M} that is optimal must also yield $W(\mathcal{M})$ that is equal to the optimal value of $\text{LP}(v_1, v_2)$. By Lemma 2, any optimal solution to $\text{LP}(v_1, v_2)$ must have $\beta_j = 0$ and $\beta_i/2 + \beta/3 = v_i$. Thus, the mediator \mathcal{M} must satisfy these as well: the first equality implying that \mathcal{M} is fully revealing to player i , and the second that player i 's IR constraint binds. ■

4. No Jointly Complete Information

We now drop the assumption that players 1 and 2 have jointly complete information. Formally, for a given type ω let $S(\omega) = P_1(\omega) \cap P_2(\omega)$, and call each such $S(\omega)$ a *segment*. The assumption of jointly complete information states that all segments are singletons, and in this section we consider the more general setting in which segments may contain more than one type of consumer.

This setting presents new challenges, as there is now a new conflict between optimal strategies and welfare. To see this, consider some segment $S = \{\omega^1, \omega^2\}$ for which $\Pr[\omega^1] =$

$4 \cdot \Pr[\omega^2]$. Then, conditional on segment S , both players have a dominant strategy, namely, to offer $g = g_{\omega^1}$. To see this, observe that a player’s utility from offering g_{ω^1} is at least $\Pr[\omega^1]/2$ regardless of the other’s action, the utility from offering g_{ω^2} is at most $\Pr[\omega^2]$, and by assumption $\Pr[\omega^1]/2 = 2 \Pr[\omega^2] > \Pr[\omega^2]$. This leads to total welfare $\Pr[\omega^1]$. However, if players were to separate rather than pool—one offering g_{ω^1} and the other g_{ω^2} —then the total welfare would be $\Pr[\omega^1] + \Pr[\omega^2]$, which is higher. Note that this conflict between optimal strategies and welfare does not occur when players have jointly complete information, since then the pooling action, which is dominant, is also welfare-maximizing. Finally, although we illustrated this conflict for the setting without an Amazon, it is of course also present with an Amazon.

This conflict between optimal strategies and welfare maximization makes the design of mediators significantly more challenging. In principle, one could formulate the problem as a complex linear program, whose solution could be used to find the optimal mediator. Here, however, we take a different approach. We will show that *simple* mediators—in particular, variants of the ones from Section 3 for the setting with jointly complete information—are *approximately* optimal. And while this approach entails some loss in terms of optimality, we will show that these simple mediators are not only approximately optimal relative to the optimal IR, IC mediator, but rather that they are approximately optimal relative to a higher benchmark—namely, the welfare that can be achieved by *any* mediator, even one that is not IR or IC. Then, to complement our results, we will show that the approximation factors our mediators achieve are tight or nearly-tight, and that no other IR and IC mediator can in general achieve a better approximation factor relative to that same benchmark.

We begin with some notation. For any given segment $S(\omega)$, denote by g_ω^1 the good g that has the highest probability $\Pr[g = g_\omega | S(\omega)]$ in $S(\omega)$, and by g_ω^2 the good with the second highest probability, under the prior π . If there are several such goods, fix two such goods arbitrarily. Next, let $\phi_\omega^1 = \Pr[g_\omega^1 = g_\omega | S(\omega)]$ be the probability of g_ω^1 conditional on $S(\omega)$, and $\phi_\omega^2 = \Pr[g_\omega^2 = g_\omega | S(\omega)]$ be the respective probability of g_ω^2 . We will also use the notation that, for a given segment $S \in \{S(\omega) : \omega \in \Omega\}$, the goods g_S^1 and g_S^2 are the goods with the highest and second-highest probabilities in S , with ϕ_S^1 and ϕ_S^2 their respective conditional probabilities. Finally, let ϕ^1 be the total (unconditional) weight of goods g_S^1 over all segments S , and ϕ^2 be the total weight of goods g_S^2 : Formally,

$$\phi^k = \sum_{S \in \{S(\omega) : \omega \in \Omega\}} \phi_S^k \cdot \Pr[S] \quad \text{for each } k \in \{1, 2\}.$$

Now, as our main interest lies in coopetition against an Amazon, we will focus here on the setting with an Amazon. We note that, while we drop the assumption that players 1 and 2 have jointly complete information, we still assume that the Amazon player 0 has complete information. That is, even though in any given segment S players 1 and 2 do not know which type $\omega \in S$ is realized, player 0 does know, and offers each such type the correct good g_ω .

We proceed as follows. We begin in Section 4.1 with the design of an optimal mediator for a particular setting of parameters, namely, for the case in which $\phi_S^1 \leq \frac{3}{2} \cdot \phi_S^2$ for every segment S . We then turn to our main analysis in Section 4.2, for the case in which $\phi_S^1 > \frac{3}{2} \cdot \phi_S^2$ for every segment S . There we design the approximately optimal simple mediator mentioned

above, and analyze the tightness of the approximation factor. Finally, in Section 4.3 we show how to combine the two mediators in order to obtain an approximately optimal mediator for the general setting, without restrictions on the segments.

4.1 Optimal Mediator for $\phi_S^1 \leq \frac{3}{2} \cdot \phi_S^2$

We begin with the simpler case in which $\phi_S^1 \leq \frac{3}{2} \cdot \phi_S^2$ for every segment S . This case is simpler because here, conditional on any given segment, there is no conflict between optimal strategies and welfare maximization. In particular, both involve separating, meaning that the players make different offers—one offers g_S^1 and the other offers g_S^2 .

Claim 1 *Conditional on any segment S , if $\phi_S^1 \leq \frac{3}{2} \cdot \phi_S^2$ then the separating strategy profile is both welfare maximizing and an equilibrium.*

Proof: We will compare the only two reasonable player choices for welfare maximization conditional on segment S : (i) pooling on g_S^1 , or (ii) separating, with one player offering g_S^1 and the other g_S^2 .

- (i) Pooling: Here each of players 1 and 2 obtains utility $\phi_S^1/3$ —when type $\omega \in S$ with $g_\omega = g_S^1$ is realized (which happens with probability ϕ_S^1), the players split the surplus of 1 with player 0, and so each gets $1/3$. Welfare of players 1 and 2 is thus $\frac{2}{3} \cdot \phi_S^1$.
- (ii) Separating: Here the player who offers g_S^1 obtains utility $\phi_S^1/2$ —when type $\omega \in S$ with $g_\omega = g_S^1$ is realized, that player splits the surplus with player 0. Similarly, the player who offers g_S^2 obtains utility $\phi_S^2/2$. Welfare of players 1 and 2 is thus $(\phi_S^1 + \phi_S^2)/2$.

Observe that, when $\phi_S^1 \leq \frac{3}{2} \cdot \phi_S^2$, we have that $(\phi_S^1 + \phi_S^2)/2 > \frac{2}{3} \cdot \phi_S^1$, and so the welfare maximizing choices are the separating ones.

In addition to being welfare-maximizing, the separating actions also form an equilibrium (conditional on S). Fix a segment S , and suppose player 1 offers g_S^1 and player 2 offers g_S^2 . Then player 1 clearly does not wish to deviate, as she currently obtains utility $\phi_S^1/2$, a deviation to g_S^2 will lead to utility $\phi_S^2/3$, and a deviation to any other good will lead to utility at most $\phi_S^1/2$ (since g_S^1 has maximal probability in S). Player 2 also does not wish to deviate, as she currently obtains utility $\phi_S^2/2$, a deviation to g_S^1 will lead to utility $\phi_S^1/3$, and by assumption $\phi_S^2/2 \geq \phi_S^1/3$. ■

Given the lack of conflict between welfare maximization and equilibrium shown in Claim 1, we can now construct optimal mediators for this parameter setting. The mediator is straightforward: it facilitates full data-sharing between the players, and recommends the actions that lead to a separating equilibrium.

Mediator 3 Mediator for $\mathcal{P} = \{0, 1, 2\}$ without jointly complete information

1: **procedure** $\mathcal{M}_1^v(S)$ $\triangleright S = P_1(\omega) \cap P_2(\omega)$ for realized $\omega \in \Omega$
2: **if** $\phi^1 = \phi^2$ **then return** (g_S^1, g_S^2)
3: **else**
4: $i \leftarrow \arg \max_{k \in \{1, 2\}} v_k$
5: $j \leftarrow 3 - i$
6: Choose $\gamma \in [0, 1]$ uniformly at random.
7: **if** $\gamma < \frac{2v_i - \phi^2}{\phi^1 - \phi^2}$ **then**
8: **if** $i = 1$ **then return** (g_S^1, g_S^2)
9: **else return** (g_S^2, g_S^1)
10: **else if** $i = 1$ **then return** (g_S^2, g_S^1)
11: **else return** (g_S^1, g_S^2)

Theorem 6 *Suppose $v_1 + v_2 \leq (\phi^1 + \phi^2)/2$ and $\max\{v_1, v_2\} \leq \phi^1/2$. Then the mediator \mathcal{M}_1^v above is IR, IC, and optimal.*

Observe that the conditions on v in Theorem 6 are also necessary for the existence of any IR mediator, since the maximal welfare attainable by any mediator in this setting is $(\phi^1 + \phi^2)/2$, and the maximal utility any single player can attain is $\phi^1/2$.

Proof: Suppose first that $\phi^1 = \phi^2$, and so the conditional in line 2 of Mediator 3 is true. In this case, each player obtains utility $\phi^1/2$ under \mathcal{M}_1^v , and the mediator is thus IR. In addition, since by Claim 1 separation is an equilibrium on every segment S , the mediator is also IC. Finally, the total welfare to the players is $(\phi^1 + \phi^2)/2$, which is maximal.

Next, suppose $\phi^1 > \phi^2$. Observe that, in this case, the mediator mixes between utility $\phi^1/2$ to player i and $\phi^2/2$ to player j , and utility $\phi^2/2$ to player i and $\phi^1/2$ to player j , depending on the value of γ . Regardless of the realizations of γ , the total welfare to the players is always $(\phi^1 + \phi^2)/2$. This is the optimal welfare attainable.

Next, the expected utility of player i in \mathcal{M}_1^v is

$$\frac{2v_i - \phi^2}{\phi^1 - \phi^2} \cdot \frac{\phi^1}{2} + \left(1 - \frac{2v_i - \phi^2}{\phi^1 - \phi^2}\right) \frac{\phi^2}{2} = v_i.$$

Furthermore, since $v_1 + v_2 \leq (\phi^1 + \phi^2)/2$ and total welfare is $(\phi^1 + \phi^2)/2$, the utility to player j is at least v_j . Thus, the mediator is IR.

Finally, since by Claim 1 separation is an equilibrium on every segment S , the mediator is IC. ■

4.2 Approximately Optimal Mediator for $\phi_S^1 > \frac{3}{2} \cdot \phi_S^2$

Suppose now that $\phi_S^1 > \frac{3}{2} \cdot \phi_S^2$ for every segment S . This case is more complicated since, conditional on any segment, there may be a conflict between optimal strategies and welfare maximization:

Claim 2 *Conditional on any segment S , if $\phi_S^1 > \frac{3}{2} \cdot \phi_S^2$ then pooling on g_S^1 is the dominant strategy. If $\phi_S^1 \leq 3 \cdot \phi_S^2$ then separating is welfare-maximizing, whereas if $\phi_S^1 > 3 \cdot \phi_S^2$ then pooling is welfare-maximizing.*

Proof: Again, we will compare the only two reasonable player choices conditional on segment S : (i) pooling on g_S^1 , or (ii) separating, with one player offering g_S^1 and the other g_S^2 . As in the proof of Claim 1, pooling leads to individual utilities $\phi_S^1/3$ and so welfare $\frac{2}{3} \cdot \phi_S^1$, whereas separating leads to individual utilities $\phi_S^1/2$ and $\phi_S^2/2$ and so welfare $(\phi_S^1 + \phi_S^2)/2$. Then $\frac{2}{3} \cdot \phi_S^1 > (\phi_S^1 + \phi_S^2)/2$ if and only if $\phi_S^1 > 3 \cdot \phi_S^2$, as claimed.

In terms of optimality, however, observe that choosing g_S^1 leads to utility at least $\phi_S^1/3$, whereas choosing any other good leads to utility at most $\phi_S^2/2$. The former is greater than the latter whenever $\phi_S^1 > \frac{3}{2} \cdot \phi_S^2$, and so constitutes a dominant action (conditional on S). ■

Our goal is to design a mediator that is approximately optimal relative to a particular benchmark. We denote the benchmark OPT , which is the maximal welfare attainable by any mediator, even disregarding the IR and IC constraints. We will show that the mediator below, which is a variant of the mediator \mathcal{M}_A^v for the case of jointly complete information, is IR, IC, and obtains welfare at least $\frac{3}{4} \cdot \text{OPT}$. We will then show that this approximation factor is tight.

Before describing the mediator, recall the strategy s'_j from Section 3 that, for every $P_j(\omega)$, chooses one of the goods that is most likely to be correct. Here we will use a different formulation of s'_j . In particular, let $s'_j(\omega) \in \arg \max_{g \in G} E[u_j(g, g_\omega^1) | P_j(\omega)]$. That is, s'_j is an optimal action for player j conditional on knowing only $P_j(\omega)$, in the hypothetical situation where i always plays the most-likely correct good g_ω^1 in every segment. Also, let α_j and α_i be the expected utilities of players j and i in this hypothetical situation: $\alpha_j = E[u_j(s'_j(\omega), g_\omega^1)]$ and $\alpha_i = E[u_i(g_\omega^1, s'_j(\omega))]$. Note that, under jointly complete information, the formulation of s'_j is identical to the one given in Section 3. Without jointly complete information, however, they may differ.

Mediator 4 Mediator for $\mathcal{P} = \{0, 1, 2\}$ without jointly complete information

- 1: **procedure** $\mathcal{M}_2^v(V, W)$ ▷ $V = P_1(\omega)$ and $W = P_2(\omega)$ for realized $\omega \in \Omega$
 - 2: $S \leftarrow V \cap W$
 - 3: $i \leftarrow \arg \max_{k \in \{1, 2\}} v_k$
 - 4: $j \leftarrow 3 - i$
 - 5: **if** $v_i \leq \phi^1/3$ **then return** (g_S^1, g_S^1) ▷ Full data-sharing
 - 6: **else**
 - 7: $g_j \leftarrow s'_j(\omega)$
 - 8: Choose $\gamma \in [0, 1]$ uniformly at random.
 - 9: **if** $\gamma < \frac{3\alpha_i - 3v_i}{3\alpha_i - \phi^1}$ **then**
 - 10: **return** (g_S^1, g_S^1)
 - 11: **else if** $i = 1$ **then return** (g_S^1, g_j)
 - 12: **else return** (g_j, g_S^1)
-

Theorem 7 Let $v_i \geq v_j$, and suppose $v_i \leq E[u_i(g_\omega^1, s'_j(\omega))]$ and $v_j \leq \phi^1/2 - v_i$. Then the mediator \mathcal{M}_2^v above is IR, IC, and has welfare $W(\mathcal{M}_2^v) \geq \frac{3}{4} \cdot \text{OPT}$.

We note that, unlike the conditions on v in Theorems 2, 4, and 6, the conditions here are not necessary but only sufficient.

Proof: Consider first the case in which $v_i \leq \phi^1/3$. Here, the utility to each player is $\phi^1/3 \geq v_i \geq v_j$, and so the mediator is IR. Furthermore, players always pool on g_ω^1 which is an equilibrium by Claim 2. Thus, in this case the mediator is also IC. Finally, total welfare here is $\frac{2}{3} \cdot \phi^1$. Below we will show that this is within a 3/4 factor of OPT.

Before turning to the approximation factor, consider the case in which $v_i > \phi^1/3$. Observe that the total welfare in this case is at least $\phi^1/2$. If on a given segment $S(\omega)$ the players pool on g_ω^1 , then welfare here is $\frac{2}{3} \cdot \phi_\omega^1 > \frac{1}{2} \cdot \phi_\omega^1$. If the players separate, then the welfare here is at least $\phi_\omega^1/2$, since player i always chooses g_ω^1 . Thus, total welfare is at least $\phi^1/2$.

Next, the expected utility of player i in this case is

$$\frac{3\alpha_i - 3v_i}{3\alpha_i - \phi^1} \cdot \frac{\phi^1}{3} + \left(1 - \frac{3\alpha_i - 3v_i}{3\alpha_i - \phi^1}\right) \alpha_i = v_i.$$

Furthermore, since the total welfare is at least $\phi^1/2$, player j obtains welfare at least $\phi^1/2 - v_i \geq v_j$, and so the mediator is IR here as well. In addition, since by Claim 2 the action g_ω^1 is dominant in every segment $S(\omega)$, and in the mediator player i always chooses g_ω^1 and player j either pools or best responds given her information, the mediator is IC.

Finally, we turn to the approximation factor. By Claim 2, the maximal welfare possible on a segment S in which $\phi_S^1 \leq 3 \cdot \phi_S^2$ is $(\phi_S^1 + \phi_S^2)/2$, and our mediator guarantees welfare at least $\frac{\phi_S^1}{2}$. Hence, here we have approximation ratio at least

$$\frac{\phi_S^1/2}{(\phi_S^1 + \phi_S^2)/2} \geq \frac{\phi_S^1/2}{(\phi_S^1 + \phi_S^1/3)/2} = \frac{3}{4}.$$

In addition, by Claim 2, the maximal welfare possible on a segment S in which $\phi_S^1 > 3 \cdot \phi_S^2$ is $2\phi_S^1/3$, and our mediator guarantees welfare at least $\frac{\phi_S^1}{2}$. Hence, here we have approximation ratio at least

$$\frac{\phi_S^1/2}{2\phi_S^1/3} = \frac{3}{4}.$$

Since a ratio of 3/4 is achieved for each segment, it is also achieved overall, and $W(\mathcal{M}_2^g) \geq \frac{3}{4} \cdot \text{OPT}$. ■

We now show that the approximation factor of 3/4 is tight, by giving an example wherein there does not exist an IR mediator, not even one that is not IC, that achieves a factor higher than 3/4.

Example 1 Fix some $\varepsilon > 0$, and suppose there are $\lceil 1/\varepsilon \rceil + 1$ segments. Within each segment S there are two correct goods, g_S^1 and g_S^2 , that satisfy $\phi_S^1 = \phi_S^2/\varepsilon$. Furthermore, all the g_S^1 goods are distinct, whereas for every S the good $g_S^2 = g^2$ is the same. Player 1 has all the information, and so $P_1(\omega) = S(\omega)$ for every ω , whereas player 2 has none, and so $P_2(\omega) = \Omega$ for every ω . Finally, $v_1 = \phi^1/2$ and $v_2 = 0$.

In this example, maximal welfare is achieved when all players pool on g_S^1 for every S , and so $\text{OPT} = \frac{2}{3} \cdot \phi^1$. However, any IR mediator must give player 1 utility $\phi^1/2$, and the only way to achieve this is to give player 2 no additional information, so that player 1 always chooses g_S^1 and player 2 always chooses g^2 (this is, in fact, what the mediator \mathcal{M}_2^v does). Total welfare here is then $(\phi^1 + \phi^2)/2 = (1 + \varepsilon)\phi^1/2$, and so the approximation ratio is

$$\frac{(1 + \varepsilon)\phi^1/2}{2\phi^1/3} = (1 + \varepsilon) \cdot \frac{3}{4} \xrightarrow{\varepsilon \rightarrow 0} \frac{3}{4}.$$

In the previous example, the main bottleneck to achieving an approximation factor higher than $3/4$ was the IR constraint. In this next example we show that even without the IR constraints, there does not exist an IC mediator that achieves an approximation factor higher than $4/5$.

Example 2 Fix some $\varepsilon > 0$, and suppose that in every segment S there are two correct goods, g_S^1 and g_S^2 , that satisfy $\phi_S^1 = (\frac{3}{2} + \varepsilon)\phi_S^2$. Furthermore, the goods in every segment are distinct from goods in other segments. By Claim 2, maximal welfare here is achieved when players separate on each segment, and so $\text{OPT} = (\phi^1 + \phi^2)/2$.

However, again by Claim 2, the dominant action for players is to pool on g_S^1 . Note that in any IC mediator, no player will ever choose a good $g = g_S^2$, regardless of her information, since she can always improve her payoff by deviating to the strictly better g_S^1 (here we use the fact that all goods are distinct). Thus, the most any IC mediator can attain is when all players pool on the top good, leading to welfare $\frac{2}{3} \cdot \phi^1$ and so approximation factor

$$\frac{\frac{2}{3} \cdot \phi^1}{\frac{\phi^1 + \phi^2}{2}} = \frac{\frac{2}{3} \cdot \phi^1}{\frac{\phi^1}{2} + \frac{\phi^1/(3/2 + \varepsilon)}{2}} = \frac{2 + \frac{4\varepsilon}{3}}{\frac{5}{2} + \varepsilon} \xrightarrow{\varepsilon \rightarrow 0} \frac{4}{5}.$$

4.3 Approximately Optimal Mediator for General Setting

Each of the mediators in the previous two sections is IR, IC, and approximately-optimal under a particular setting of parameters, but not in the other. Mediator \mathcal{M}_1^v of Mediator 3 is not IC when $\phi_S^1 > \frac{3}{2} \cdot \phi_S^2$ for some segment S , since players wish to pool but the mediator facilitates separating. \mathcal{M}_2^v of Mediator 4 is not IC when $\phi_S^1 < \frac{3}{2} \cdot \phi_S^2$ for some segment S , since players wish to separate but the mediator facilitates pooling. In this section we show that the two mediators can be combined in order to yield an IR, IC, and approximately optimal mediator for any parameter setting. The main idea is straightforward: to run \mathcal{M}_1^v of Mediator 3 on segments S for which $\phi_S^1 \leq \frac{3}{2} \cdot \phi_S^2$, and to run \mathcal{M}_2^v of Mediator 4 on segments S for which $\phi_S^1 > \frac{3}{2} \cdot \phi_S^2$.

Before formalizing this mediator, we need some notation. Let $S_1 = \{S : \phi_S^1 \leq \frac{3}{2} \cdot \phi_S^2\}$ and $S_2 = \{S : \phi_S^1 > \frac{3}{2} \cdot \phi_S^2\}$. Also, for each $\ell \in \{1, 2\}$ let ϕ_ℓ^1 be the total weight of the most-likely correct goods conditional on S_ℓ , and ϕ_ℓ^2 be the total weight of the second most-likely correct goods conditional S_ℓ . Formally, for each $k \in \{1, 2\}$,

$$\phi_\ell^k = \sum_{S \in S_\ell} \phi_S^k \cdot \Pr[S|S_\ell].$$

Finally, modify the definitions of s'_j , α_j , and α_i from Section 4.2 as follows. First, for any ω with $S(\omega) \in S_2$ let $s'_j(\omega) \in \arg \max_{g \in G} E[u_j(g, g_\omega^1) | P_j(\omega), S_2]$, so it is an optimal action for

player j conditional on the information that $\omega \in P_j(\omega)$ and $S(\omega) \in S_2$, in the hypothetical situation where i always plays the most-likely correct good g_ω^1 in every segment $S(\omega)$. Also, let α_j and α_i be the expected utilities of players j and i in this hypothetical situation: $\alpha_j = E \left[u_j(s'_j(\omega), g_\omega^1) | S_2 \right]$ and $\alpha_i = E \left[u_i(g_\omega^1, s'_j(\omega)) | S_2 \right]$.

Mediator 5 Mediator for $\mathcal{P} = \{0, 1, 2\}$ without jointly complete information

- 1: **procedure** $\mathcal{M}_3^v(V, W)$ $\triangleright V = P_1(\omega)$ and $W = P_2(\omega)$ for realized $\omega \in \Omega$
 - 2: $S \leftarrow V \cap W$
 - 3: **if** $\phi_S^1 \leq \frac{3}{2} \cdot \phi_S^2$ **then run** $\mathcal{M}_1^v(S)$ with (ϕ_1^1, ϕ_1^2) replacing (ϕ^1, ϕ^2)
 - 4: **else run** $\mathcal{M}_2^v(V, W)$ with (ϕ_2^1, ϕ_2^2) replacing (ϕ^1, ϕ^2)
-

Theorem 8 *Let $v_i \geq v_j$, and suppose that the following hold: $v_1 + v_2 \leq (\phi_1^1 + \phi_1^2)/2$, $v_i \leq \max \left\{ E \left[u_i(g_\omega^1, s'_j(\omega)) | S_2 \right], \frac{\phi_1^1}{2} \right\}$, and $v_j \leq \phi_2^1/2 - v_i$. Then the mediator \mathcal{M}_3^v above is IR, IC, and has welfare $W(\mathcal{M}_3^v) \geq \frac{3}{4} \cdot \text{OPT}$.*

Proof: Conditional on S_1 , the conditions on v of Theorem 6 are satisfied, and so \mathcal{M}_1^v is IR and IC. Conditional on S_2 , the conditions on v of Theorem 7 are satisfied, and so \mathcal{M}_2^v is IR and IC. That is, conditional on either S_1 or S_2 , the expected utilities of either player k is at least v_k . Thus, k 's expected utility in \mathcal{M}_3^v , which is a mixture of her utilities conditional on S_1 and on S_2 , must also be greater than v_k . This implies that \mathcal{M}_3^v is IR.

Furthermore, since \mathcal{M}_1^v and \mathcal{M}_2^v are IC, for either player k the strategy \bar{s}_k is a best-response to \bar{s}_{3-k} conditional on either S_1 or S_2 . Thus, \mathcal{M}_3^v is also IC.

Finally, we consider the welfare guaranteed by \mathcal{M}_2^v . Conditional on S_1 , the welfare attained is $(\phi_1^1 + \phi_1^2)/2$, by Theorem 6. This is the optimal attainable by any mediator when $\phi_S^1 \leq \frac{3}{2} \cdot \phi_S^2$ for every segment S , by Claim 1. Hence, conditional on S_1 , the welfare achieved by \mathcal{M}_3^v is equal to OPT (conditional on S_1).

Conditional on S_2 , the welfare attained on any segment S is at least $\frac{3}{4}$ times the welfare attainable on this segment by any mediator (by the proof of Theorem 7). Hence, the welfare achieved by \mathcal{M}_3^v is at least $\frac{3}{4} \cdot \text{OPT}$ (conditional on S_2).

Since conditional on either S_1 or S_2 yields welfare at least $\frac{3}{4}$ times the optimal, the total welfare of \mathcal{M}_3^v is $W(\mathcal{M}_3^v) \geq \frac{3}{4} \cdot \text{OPT}$. ■

5. Do Players Share More Data With or Without an Amazon?

In this section we consider the intriguing question of whether players optimally share more data in the presence of an Amazon or in its absence. While it seems intuitive that players would share more data when facing stronger outside competition, we show here that this is not necessarily the case. We also show that the reason the intuition fails is related to the possible conflict between equilibrium and welfare maximization identified in Claim 2.

Because we are interested in comparing data-sharing across two different environments—without and with an Amazon—we focus on our main interpretation of the base values v , namely, that they correspond to the equilibrium expected utilities absent a mediator. Of course, because the two environments differ, so will these equilibrium expected utilities. We will thus be interested in comparing the benefit of data-sharing relative to firms' no-sharing

equilibrium utilities v^{noA} absent an Amazon with the benefit of data-sharing relative to firms' no-sharing equilibrium utilities v^A in the presence of an Amazon.

We begin our analysis in Section 5.1, where we show that, when there is jointly complete information, the intuition that firms share more data in the presence of an Amazon than in its absence does hold. In particular, we show that if, when there is no Amazon, data-sharing strictly increases welfare relative to the welfare-maximizing equilibrium sans data-sharing, then optimal data-sharing strictly increases welfare in the environment with an Amazon relative to any equilibrium sans data-sharing.

In Section 5.2 we then drop the assumption of jointly complete information, and show that this result no longer holds. In particular, we describe two settings:

1. A setting where full data-sharing is optimal in the presence of an Amazon, but no data-sharing is optimal in the absence of an Amazon.
2. A setting where full data-sharing is optimal in the *absence* of an Amazon, but no data-sharing is optimal in the *presence* of an Amazon. In this setting, data-sharing is strictly beneficial to firms in the absence of an Amazon but not in its presence.

In Section 5.2 we also identify and discuss the driving force behind this difference, namely, the possible conflict between equilibrium and welfare maximization.

5.1 Conditions for Sharing with Jointly Complete Information

In this section we show that, under jointly complete information, when players can strictly benefit from data-sharing in the absence of an Amazon, they can also strictly benefit from data-sharing in the presence of an Amazon.

Theorem 9 *Fix an unmediated game $\Gamma = (\Omega, \mathcal{P}, G^2, (P_i(\cdot))_{i \in \mathcal{P}}, (u_i)_{i \in \mathcal{P}}, \pi)$ in which there is jointly complete information and no Amazon, and let (v_1, v_2) be the expected utilities in a BNE of Γ for which $v_1 + v_2$ is maximal. Also, let $\Gamma' = (\Omega, \mathcal{P}, G^3, (P_i(\cdot))_{i \in \mathcal{P}}, (u_i)_{i \in \mathcal{P}}, \pi)$ be the same unmediated game, except with an Amazon, and let (v'_1, v'_2) be the expected utilities in some BNE of Γ' . Then if there exists an IR, IC mediator \mathcal{M} in the game without an Amazon such that $W(\mathcal{M}) > v_1 + v_2$, there also exists an IR, IC mediator \mathcal{M}' in the game with an Amazon such that $W(\mathcal{M}') > v'_1 + v'_2$.*

The converse of Theorem 9 is not true—there are examples where players strictly benefit from data-sharing in the presence of an Amazon but not in its absence. A simple example is the game described in Figure 1 in the Introduction, with parameters $\alpha > \beta = \gamma > \delta = 0$. In the absence of an Amazon, the unique equilibrium there yields maximal welfare of 1, since at least one consumer always offers the correct good. Thus, data-sharing cannot strictly improve welfare. In the presence of an Amazon, however, data-sharing can be strictly beneficial, as described in the Introduction.

Proof: Suppose towards a contradiction that there exists an IR, IC mediator \mathcal{M} in the game without an Amazon such that $W(\mathcal{M}) > v_1 + v_2$, but that for some (v'_1, v'_2) there does not exist an IR, IC mediator \mathcal{M}' in the game with an Amazon such that $W(\mathcal{M}') > v'_1 + v'_2$. Let E be the equilibrium in the game with an Amazon that yields payoffs (v'_1, v'_2) , and let

\mathcal{M}_0 be the mediator that does nothing other than recommend players their actions from the equilibrium E . This mediator is IR, IC, and yields welfare $W(\mathcal{M}_0) = v'_1 + v'_2$. Since, by assumption, no IR, IC mediator yields higher welfare, \mathcal{M}_0 is optimal relative to the feasible pair (v'_1, v'_2) . Thus, by Theorem 5, \mathcal{M}_0 is fully revealing to at least one of the two players. In particular, this implies that the BNE E in the game with an Amazon is one in which at least one of the players always offers the correct good g_ω to every consumer ω .

We now argue that this implies that the same E is an equilibrium also in the game without an Amazon. The player i that always offers the correct good is clearly playing a dominant strategy, both in the game with an Amazon and in the game without, as any deviation can only lower her utility. The other player j is playing a best response in the game with an Amazon. Because both player i and the Amazon always offer the correct good, j 's strategy must maximize the probability of offering the correct good conditional on knowing (only) $P_j(\omega)$. However, this strategy is also an equilibrium in the game without an Amazon, since i always offers the correct good.

The equilibrium E is such that at least one of the players always offers the correct good, and so in the game without an Amazon it yields a total welfare of 1. However, this contradicts the assumption that there exists a mediator \mathcal{M} that strictly improves welfare, since no mediator can yield welfare $W(\mathcal{M}) > 1$. ■

5.2 Data-Sharing with No Jointly Complete Information

The analysis of Section 5.1 shows that when players strictly benefit from data-sharing in the absence of an Amazon, they necessarily also benefit from data-sharing in the presence of an Amazon. This is consistent with the intuition that players can gain more by sharing data when they also face outside competition. However, this result relies on the assumption of jointly complete information. In this section we drop the assumption, and show that in that case this intuition is not complete. In particular, while there are settings in which players benefit more from sharing data in the presence of an Amazon, there are also settings in which players benefit from data-sharing in the absence of an Amazon but *not* in the presence of an Amazon.

More specifically, in Section 5.2.1 below we describe a setting where:

- In the absence of an Amazon, the optimal welfare is achieved when players share no data and not when they fully share data;
- In the presence of an Amazon, the optimal welfare is achieved when players fully share data and not when they share no data.

In contrast, in Section 5.2.1 we describe a setting where:

- In the absence of an Amazon, the optimal welfare is achieved when players fully share data and not when they share no data;
- In the presence of an Amazon, the optimal welfare is achieved when players share no data and not when they fully share data.

The main driving force behind the distinction between the two settings is whether or not there is a conflict between equilibrium and welfare maximization, as formalized in Claim 2.

The situation described in Section 5.2.1 is one where $\phi_S^1 \geq 3 \cdot \phi_S^2$ for every segment S . In this case, under full data-sharing, it is a dominant strategy for players to pool on g_S^1 (both with and without an Amazon). In the presence of an Amazon, such pooling is also welfare maximizing, by Claim 2. In the absence of an Amazon, however, separating is always welfare maximizing. Thus, full data-sharing leads to the optimal welfare in the former case, but not in the latter.

In contrast, the situation described in Section 5.2.2 is one where $\phi_S^1 \in (\frac{3}{2}\phi_S^2, 2\phi_S^2)$. In this case, under full data-sharing, in the absence of an Amazon it is an equilibrium for players to separate, whereas in the presence of an Amazon it is dominant for players to pool (again, by Claim 2). Furthermore, in both settings, separating is the welfare maximizing outcome. Thus, full data-sharing leads to the optimal welfare in the absence of an Amazon, but not in its presence.

5.2.1 MORE SHARING WITH AN AMAZON

When there is an Amazon, if for each segment S it holds that $\phi_S^1 \geq 3 \cdot \phi_S^2$, then maximal welfare is achieved when players pool on ϕ_S^1 (by Claim 2). This maximal welfare can be achieved by full data-sharing (assuming the IR constraint is satisfied): On each segment S , the mediator recommends g_S^1 to both players, and this is optimal for them (again by Claim 2). In contrast, when there is no Amazon, full data-sharing may *not* be optimal (regardless of the IR constraint), as here maximal welfare is attained by separation.

In this section we construct an example in which full data-sharing is strictly suboptimal when there is no Amazon. While the example from the introduction illustrated this, it is not quite what we desire, since in that example the mediator with an Amazon is *also* an optimal mediator without. What we would like is an example where full data-sharing is optimal only with an Amazon. In the analysis that follows, we first ignore the IR constraint. At the end of the section we then show that, with the proper modification, the argument holds also when taking the IR constraint into account.

In the following construction, for simplicity we refer to different types of consumers ω by the good g_ω they desire. There are two segments, $S^1 = \{g_1\}$ and $S^2 = \{g_1, g_2\}$, and players' partitions are $\Pi_1 = \{S^1, S^2\}$ and $\Pi_2 = \Omega$. That is, player 1 learns whether the realized ω lies in the segment S^1 (and so the correct good is g_1) or S^2 (and so the correct good is g_1 or g_2), whereas player 2 does not learn anything. The prior is such that both segments are equally likely, and that $\Pr [g_2|S^2] > 3 \Pr [g_1|S^2] > 0$. That is, in segment S^2 , the good g_2 is more than three times more likely to be correct than good g_1 .

In the presence of an Amazon, the welfare maximizing choices lead to pooling on g_1 in segment S^1 and pooling on g_2 in segment S^2 , by Claim 2. This can be achieved by full data-sharing, leading to an equilibrium in which players take these pooling actions (again, by Claim 2). In contrast, with no data-sharing, player 2 chooses the action g_1 (in both segments), leading to lower welfare due to separation on S^2 .

Now suppose there is no Amazon. Here the welfare maximizing choices have players separating on segment S^2 . Without any additional information, player 2 chooses action g_1 , leading to such optimal separation. With data-sharing, however, player 2 takes action g_2 on S^2 , pooling with player 1, leading to lower welfare. Thus, in this example, full data-sharing

is optimal in the presence of an Amazon but not in its absence. Furthermore, without an Amazon, it is optimal to not share any data.

Finally, recall that we thus far neglected the IR constraint. We now modify the example so that the conclusion holds when also taking the IR constraints into account. In particular, duplicate the construction with a different set of segments and consumers, but where players 1 and 2 have the opposite information. Formally, in addition to the two segments, $S^1 = \{g_1\}$ and $S^2 = \{g_1, g_2\}$, suppose there are also segments $\bar{S}^1 = \{\bar{g}_1\}$ and $\bar{S}^2 = \{\bar{g}_1, \bar{g}_2\}$, where the goods g^1 and g^2 are distinct from the goods \bar{g}^1 and \bar{g}^2 . Furthermore, the partitions are $\Pi_1 = \{S^1, S^2, \bar{S}^1 \vee \bar{S}^2\}$ and $\Pi_2 = \{S^1 \vee S^2, \bar{S}^1, \bar{S}^2\}$. That is, if the realized segment is S^1 or S^2 , then again player 1 learns the segment and player 2 learns only that the realized segment is either S^1 or S^2 . Symmetrically, if the realized segment is \bar{S}^1 or \bar{S}^2 , then player 2 learns the segment and player 1 learns only that the realized segment is either \bar{S}^1 or \bar{S}^2 . Finally, suppose that the realized segment is S^1 or S^2 with probability $1/2$, and \bar{S}^1 or \bar{S}^2 with probability $1/2$.

Because the example is symmetric, the utilities of the two players, absent data-sharing, are equal. Similarly, the utilities of the two players, under full data-sharing, are also equal. Thus, whenever full data-sharing leads to strictly higher welfare than no data-sharing (as in the presence of an Amazon above), then it also satisfies the IR constraint. And if no data-sharing is optimal in the original example (as in the absence of an Amazon above), then it is also optimal in the modified setting. Thus, the conclusion that full data-sharing may be optimal in the presence of an Amazon but not in its absence holds also when we take the IR constraint into account.

5.2.2 MORE SHARING WITHOUT AN AMAZON

When there is no Amazon, the optimal welfare is achieved when, for each segment S , players separate—one choosing g_S^1 and the other g_S^2 —as this leads to maximal utility $g_S^1 + g_S^2$ in each segment. Furthermore, if $\phi_S^1 < 2\phi_S^2$ for every segment S , then this maximal welfare can be achieved by full data-sharing: On each segment S , the mediator recommends g_S^1 to one of the players, and g_S^2 to the other. In contrast, in the presence of an Amazon, if $\phi_S^1 < 2\phi_S^2$ for every segment S then full data-sharing may *not* be optimal. In this section we construct an example in which full data-sharing is strictly suboptimal with an Amazon. As in Section 5.2.1 above, in the following construction, for simplicity we refer to different types of consumers ω by the good g_ω they desire. We also ignore the IR constraint. However, again as in Section 5.2.1 above, the conclusion holds also when we take the IR constraint into account if we consider a modified, symmetric example in which the segments and partitions are duplicated.

Suppose there are three segments, $S^1 = \{g_1\}$, $S^2 = \{g_3\}$, and $S^3 = \{g_1, g_2, g_3\}$, and players' partitions are $\Pi_1 = \{S^1, S^2 \vee S^3\}$ and $\Pi_2 = \{S^2, S^1 \vee S^3\}$. That is, player 1 learns whether the realized ω lies in the segment S^1 , and so $g_\omega = g_1$, or whether ω lies in one of S^2 or S^3 , and so the correct good for the realized type is one of g_1 , g_2 , or g_3 (with relative probabilities to be specified shortly). Player 2 is similar, but learns whether the realized segment is S^2 or one of S^1 or S^3 . The prior is such that the segments are equally likely. Denote by $\phi_3^k = \Pr[g_k|S_3]$ the probability that good g_k is correct, conditional on segment S^3 , and suppose that $\phi_3^1 \in \left(\frac{3 \cdot \phi_3^2}{2}, 2 \cdot \phi_3^2\right)$ and that $\phi_3^2 > \phi_3^3 > \phi_3^1/3$.

Suppose there is no Amazon. The optimal welfare here, $(2 + \phi_3^1 + \phi_3^2)/3$, is attained when at least one player chooses g_1 conditional on S^1 , at least one player chooses g_3 conditional on S^2 , and when one player chooses g_1 and another chooses g_2 conditional on S^3 . Full data-sharing leads players to learn the correct segment, and so to maximal welfare: on segments S^1 and S^2 players pool on the correct good, and on segment S^3 it is an equilibrium for players to separate.

In contrast, without data-sharing, while player 2 will end up choosing g_1 on S^3 , player 1 will end up choosing g_3 on S^3 . This is because player 1 cannot differentiate between S^2 and S^3 , and conditional on $S^2 \vee S^3$ the strategy of choosing good g_3 is dominant. This, however, leads to total welfare $(2 + \phi_3^1 + \phi_3^3)/3 < (2 + \phi_3^1 + \phi_3^2)/3$, and so full data-sharing is better than no sharing.

Now suppose there is an Amazon. When there is no data-sharing, then conditional on segments S^1 or S^2 , welfare will be $2/3$, since for both players it is a dominant strategy to offer the correct good. On segment S^3 the strategies g_3 and g_1 of players 1 and 2, respectively, will again be dominant (as in the case of no Amazon), leading to conditional welfare $(\phi_3^1 + \phi_3^3)/2$ here.

What happens will full data-sharing? Here, on segment S^3 , players will *pool* on action g_1 , by Claim 2 and the assumption that $\phi_3^1 > \frac{3}{2}\phi_3^2$. Thus, welfare conditional on S^3 will be

$$\frac{2}{3} \cdot \phi_3^1 < \frac{\phi_3^1 + \phi_3^3}{2},$$

where the inequality follows from the assumption that $\phi_3^3 > \phi_3^1/3$. Since, conditional on S^1 or S^2 , welfare is the same with full sharing and with no sharing, total (unconditional) welfare is higher with no sharing than with full sharing.

Furthermore, observe that under no data-sharing, welfare is maximal conditional on S_1 or S_2 . Under S_3 , higher welfare could (only) be attained if players were to separate on g_1 and g_2 . However, such separation can never occur in equilibrium: Any player that is supposed to offer g_2 can strictly benefit by deviating to offering g_1 , regardless of the other player's action (by Claim 2 and since $\phi_3^1 > \frac{3}{2}\phi_3^2$). Therefore, under no data-sharing players obtain maximal welfare subject to the IC constraint being satisfied.

Thus, in this example, full data-sharing is optimal in the absence of an Amazon, but not in its presence. In contrast, no data-sharing is optimal in the presence of an Amazon, but not in its absence. In particular, this implies that, in this example, data-sharing is strictly beneficial in the absence of an Amazon, but not in its presence.

6. Conclusion

In this paper we proposed a simple model to study cooperative data-sharing between firms. We designed several data-sharing schemes, in the form of mediators, that are optimal in their respective domains, with and without an Amazon.

In our model, we assumed that the mediator only facilitates data-sharing between firms. One natural extension is to allow the mediator to also use monetary transfers between them. In Appendix A.1 we analyze this extension. We show that transfers can loosen the IR constraint, implying that, under jointly complete information, full data-sharing is often optimal. However, when there is no jointly complete information, full data-sharing may

violate the IC constraint, and in this case our mediators (even without transfers) are close to optimal.

One limitation of our model is that it considers data-sharing between only two firms. What happens when there are $n > 2$ small firms, and a mediator that can facilitate sharing between all of them? We leave a thorough analysis of optimal mediators in this setting for future work, but provide some preliminary results in Appendix A.2. In particular, we analyze the case of jointly complete information. We show that, in cooperation without an Amazon, a generalization of the mediator from the introductory example to n players is optimal. In cooperation against an Amazon, however, the situation is more complicated. Nonetheless, we provide necessary and sufficient conditions for full data-sharing to be optimal in this case.

Another limitation of our model is that we assume that firms truthfully reveal their data to the mediator. As discussed in Section 2, this is a natural assumption in our context, and can be justified both when mediators are service providers that already store the firms' data, and when the firms interact repeatedly. As noted there, however, it is nonetheless an interesting question to determine the minimal powers a mediator requires in order to incentivize truthful reporting. Suppose players are not assumed to truthfully reveal their data, but have the ability to manipulate their reports. Suppose also that the mediator observes the actual goods offered by the players following the mediator's recommendations. Finally, for simplicity, suppose that there is jointly complete information, no Amazon, and that both $v_1 < 1/2$ and $v_2 < 1/2$. The relevant mediator is \mathcal{M}_{noA}^v described in Theorem 2, and its recommendations are the correct good to both players (line 5 of Mediator 1). In this case, if the mediator observes that one of the players did not offer the recommended good, then the mediator can infer that this firm did not truthfully report its data. Thus, the mediator can "catch" cheaters, implying that it is not necessary to assume full truthfulness. We leave the question of whether or not this is possible in other settings, and particularly under no jointly complete information, to future work.

A final limitation of our model is that players' payoffs are assumed to be symmetric. In a more general model, one player's payoff from offering the correct good could be higher in some states, and another player's payoff could be higher in other states. Again, we leave the analysis of this generalization to future work.

Coopetitive data-sharing is crucial for firms' survival in current online markets, and the model and results of this paper have only scratched the surface of what can and should be done. In particular, the paper leaves numerous directions open for future research. In addition to thorough analyses of optimal data-sharing with more than 2 players, weaker mediators, and asymmetric players, a particularly important direction is to analyze data-sharing in a variety of markets different from the one studied here.

Acknowledgements

Gradwohl gratefully acknowledges the support of National Science Foundation award number 1718670. The work of Tennenholtz was supported by funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant number 740435).

Appendix A. Extensions

A.1 Transfers

In this section we consider an extension to the model: the possibility of monetary transfers between players. With transfers, the mediator may require one player to pay an amount c , and another to receive that amount c , when opting in. Assume total utilities are quasi-linear in these transfers—that is, a player’s total utility is her utility from the game plus the change in her monetary position. How does the possibility of such transfers affect the design of mediators?

A first observation is that the optimality of mediator \mathcal{M}_{noA}^v , under jointly complete information and without an Amazon, is not affected, since that mediator leads to a globally optimal sum of player’s utilities. Any transfers will only change the distribution of utilities (and hence whether or not the mediator is IR), not their sum.

When players have jointly complete information and they compete against an Amazon, however, the optimal mediator can be different. In particular, when there are transfers, then a mediator that is fully revealing to both players is optimal whenever $v_1 + v_2 \leq 2/3$. If both $v_i \leq 1/3$ and $v_j \leq 1/3$, then the mediator that is fully revealing to both is optimal even without a transfer (this is what \mathcal{M}_A^v does). If $v_i \geq 1/3 \geq v_j$, let the mediator set a transfer $c = (v_i - 1/3)$, have player j transfer c to player i , and then always recommend g_ω to both players. After the transfer this leads to utility $1/3 + c$ to player i and $1/3 - c$ to player j . Observe that $1/3 + c = 1/3 + (v_i - 1/3) = v_i$ and $1/3 - c = 1/3 - (v_i - 1/3) = 2/3 - v_i \geq v_j$, and so the mediator is IR. The mediator is IC since each player plays the dominant action g_ω for every ω . And, finally, the mediator is optimal, since the sum of utilities is $2/3$, which is globally maximal.

When there is no jointly complete information the situation is slightly more complicated, and here we consider three settings of parameters. First, if $\phi_S^1 \leq \frac{3}{2} \cdot \phi_S^2$, then the mediator \mathcal{M}_1^v of Mediator 3 is IR, IC, and optimal, even without transfers. Second, if $\phi_S^1 > 3 \cdot \phi_S^2$, then by Claim 2 the pooling strategy is both dominant and welfare maximizing. In this case, a mediator like the mediator with jointly complete information and transfers above will be optimal whenever $v_1 + v_2 \leq \frac{2}{3} \cdot \phi^1$. If both $v_i \leq \phi^1/3$ and $v_j \leq \phi^1/3$, then the mediator that is fully revealing to both is optimal even without a transfer (as in \mathcal{M}_2^g). If $v_i \geq \phi^1/3 \geq v_j$, let the mediator set a transfer $c = (v_i - \phi^1/3)$, have player j transfer c to player i , and then always recommend g_ω to both players. This mediator is IR, IC, and optimal, by the same reasoning as above.

Third, if $\phi_S^1 \in (\frac{3}{2} \cdot \phi_S^2, 3 \cdot \phi_S^2)$, then full data-sharing does not always satisfy the IC constraint. This is demonstrated by the example in Section 5.2.1, in which the optimal IC mediator is one in which there is no data-sharing (and it is strictly better than full data-sharing). Since the example is symmetric, players’ payoffs are identical, and so transfers do not alleviate the situation. Nonetheless, recall that mediator \mathcal{M}_3^v (analyzed in Theorem 8) yields welfare at least $\frac{3}{4} \cdot \text{OPT}$, where OPT is the maximal welfare achieved by any mediator, including one that is not IR or IC. Furthermore, Example 2 shows that no mediator (not even one that does not satisfy IR) can in general achieve more than $\frac{4}{5} \cdot \text{OPT}$. This implies that even the best mediator *with transfers* can yield at most $\frac{4}{5} \cdot \text{OPT}$. Our mediator \mathcal{M}_3^v achieves close to this—namely, $\frac{3}{4} \cdot \text{OPT}$ —*without* transfers.

A.2 Many Players

Throughout this paper we analyzed the case of 2 players and possibly an Amazon. In this section we discuss an extension of the model and results to $n > 2$ players. We leave a thorough analysis of optimal mediators in this setting for future work, but here we provide some preliminary results. In particular, we consider the case in which players have jointly complete information: $\bigcap_{i \in \{1, \dots, n\}} P_i(\omega) = \{\omega\}$ for every $\omega \in \Omega$, where $P_i(\omega)$ is the partition element of player i when consumer type ω is realized.

The main take-away from this section is that in coopetition without an Amazon, a generalization of the mediator from the introductory example to n players is optimal. In coopetition against an Amazon, however, the situation is more complicated. We provide necessary and sufficient conditions for full data-sharing to be optimal in this case.

A.2.1 MANY PLAYERS WITHOUT AMAZON

Let $E = (s_1, \dots, s_n)$ be a BNE of the unmediated game, where $s_i(\omega)$ is the equilibrium strategy of player i when she obtains information $P_i(\omega)$. We will design an optimal mediator for the setting in which, for every i , the base value v_i is at most the utility of player i in E . Consider the following mediator, and observe that it uses the equilibrium E in its construction (in contrast with \mathcal{M}_{noA}^v , which only uses the base values v).

Mediator 6 Mediator for $\mathcal{P} = \{1, \dots, n\}$ given equilibrium $E = (s_1, \dots, s_n)$

- 1: **procedure** $\mathcal{M}_{noA}^E(V_1, \dots, V_n)$ $\triangleright V_i = P_i(\omega)$ for realized $\omega \in \Omega$ and every i
 - 2: $\omega \leftarrow \bigcap_{i \in \{1, \dots, n\}} V_i$
 - 3: $g_i \leftarrow$ a draw from distribution $s_i(\omega)$ for every i
 - 4: **if** $(g_i \neq g_\omega \ \forall i \in \mathcal{P})$ **then return** $(g_\omega, \dots, g_\omega)$
 - 5: **else return** (g_1, \dots, g_n)
-

Theorem 10 *Fix a BNE $E = (s_1, \dots, s_n)$ of the unmediated game with no Amazon, and for every i let v_i be at most the expected utility of player i in E . Then \mathcal{M}_{noA}^E is IR, IC, and optimal.*

In mediator \mathcal{M}_{noA}^E , every player obtains the same utility as under E , plus additional utility whenever no player correctly chooses g_ω . Thus, the mediator is IR. Furthermore, observe that, under this mediator, at least one player chooses g_ω for every ω . This implies that the total welfare is 1, which is maximal. The main challenge in the proof is to show that the mediator is IC: in the mediator each player either plays her equilibrium strategy $s_i(\omega)$ or the dominant action g_ω , and the proof involves showing that, conditional on not obtaining recommendation g_ω , the strategy $s_i(\omega)$ is still optimal.

Proof: Recall that \bar{s} is the strategy profile in which players always follow \mathcal{M}_{noA}^E 's recommendation. We first show that all players receive (weakly) higher expected utilities under \bar{s} than under E . Consider some $\omega \in \Omega$, and some realizations (g_1, \dots, g_n) of $(s_1(\omega), \dots, s_n(\omega))$. Consider two events: If $g_i \neq g_\omega$ for all i then without the mediator, all players obtain utility zero. The mediator here recommends g_ω to all, leading to positive utility to all players. Conditional on this event, then, the mediator is beneficial to all

players. The other event is the complement of this first event, in which case the mediator recommends players' BNE actions. Since the mediator's recommendation does not change players' actions relative to E , it also does not affect their utilities. Overall, then, all players prefer \bar{s} in \mathcal{M}_{noA}^E to E , and so the mediator is IR.

Next, we show that \bar{s} is a BNE of $\Gamma^{\mathcal{M}_{noA}^E}$. Fix ω and a player i , and suppose all other players j play \bar{s}_j . Denote this profile of other players as \bar{s}_{-i} . To simplify notation, denote by $\mathcal{M} = \mathcal{M}_{noA}^E$ and by \mathcal{M}_i the distribution over recommendations to player i in state ω . Also, fix some $g \in \text{supp}(s_i(\omega))$, and let $\alpha = \Pr[(s_i(\omega) \neq g) | P_i(\omega) \cap (\mathcal{M}(\omega)_i = g)]$. Finally, for simplicity, denote by $s_i = s_i(\omega)$, the strategy of player i in state ω . Then when the mediator recommends action g , player i 's expected utility from following that recommendation is

$$\begin{aligned} & E[u_i(g, \bar{s}_{-i}) | P_i(\omega) \cap (\mathcal{M}_i = g)] \\ &= \alpha E[u_i(g, \bar{s}_{-i}) | P_i(\omega) \cap (\mathcal{M}_i = g) \cap (s_i \neq g)] \\ &\quad + (1 - \alpha) E[u_i(g, \bar{s}_{-i}) | P_i(\omega) \cap (\mathcal{M}_i = g) \cap (s_i = g)] \end{aligned} \tag{1}$$

$$= \alpha/n + (1 - \alpha) E[u_i(g, \bar{s}_{-i}) | P_i(\omega) \cap (\mathcal{M}_i = g) \cap (s_i = g)] \tag{2}$$

$$= \alpha/n + (1 - \alpha) E[u_i(g, s_{-i}) | P_i(\omega) \cap (\mathcal{M}_i = g) \cap (s_i = g)] \tag{3}$$

$$\begin{aligned} &= \alpha/n + \frac{1 - \alpha}{\Pr[\mathcal{M}_i = g | P_i(\omega) \cap (s_i = g)]} \cdot \left(E[u_i(g, s_{-i}) | P_i(\omega) \cap (s_i = g)] \right. \\ &\quad \left. - E[u_i(g, s_{-i}) | P_i(\omega) \cap (\mathcal{M}_i \neq g) \cap (s_i = g)] \cdot \Pr[\mathcal{M}_i \neq g | P_i(\omega) \cap (s_i = g)] \right) \end{aligned} \tag{4}$$

$$= \alpha/n + \frac{(1 - \alpha) \cdot (E[u_i(g, s_{-i}) | P_i(\omega) \cap (s_i = g)])}{\Pr[\mathcal{M}_i = g | P_i(\omega) \cap (s_i = g)]}, \tag{5}$$

where (1) follows from the law of total expectation, (2) follows since $s_i \neq \mathcal{M}_i$ implies that the mediator recommended g_ω to all players, (3) follows since $s_i = \mathcal{M}_i$ implies that the mediator recommended the equilibrium action to all players, (4) follows from another application of the law of total expectation, and (5) follows since $\mathcal{M}_i \neq g = s_i$ implies that $g \neq g_\omega$.

Similarly, i 's expected utility from not following the mediator's recommendation g and playing $g' \neq g$ instead is

$$\begin{aligned} & E[u_i(g', \bar{s}_{-i}) | P_i(\omega) \cap (\mathcal{M}_i = g)] \\ &= \alpha E[u_i(g', \bar{s}_{-i}) | P_i(\omega) \cap (\mathcal{M}_i = g) \cap (s_i \neq g)] \\ &\quad + (1 - \alpha) E[u_i(g', \bar{s}_{-i}) | P_i(\omega) \cap (\mathcal{M}_i = g) \cap (s_i = g)] \\ &= (1 - \alpha) E[u_i(g', \bar{s}_{-i}) | P_i(\omega) \cap (\mathcal{M}_i = g) \cap (s_i = g)] \\ &= (1 - \alpha) E[u_i(g', s_{-i}) | P_i(\omega) \cap (\mathcal{M}_i = g) \cap (s_i = g)] \\ &= \frac{1 - \alpha}{\Pr[\mathcal{M}_i = g | P_i(\omega) \cap (s_i = g)]} \cdot \left(E[u_i(g', s_{-i}) | P_i(\omega) \cap (s_i = g)] \right. \\ &\quad \left. - E[u_i(g', s_{-i}) | P_i(\omega) \cap (\mathcal{M}_i \neq g) \cap (s_i = g)] \cdot \Pr[\mathcal{M}_i \neq g | P_i(\omega) \cap (s_i = g)] \right) \\ &\leq \frac{(1 - \alpha) \cdot (E[u_i(g', s_{-i}) | P_i(\omega) \cap (s_i = g)])}{\Pr[\mathcal{M}_i = g | P_i(\omega) \cap (s_i = g)]}. \end{aligned}$$

Since E is an equilibrium, it holds that

$$E[u_i(g, s_{-i}(\omega)) | P_i(\omega) \cap s_i(\omega) = g] \geq E[u_i(g', s_{-i}(\omega)) | P_i(\omega) \cap s_i(\omega) = g],$$

and so

$$E[u_i(g, s_{-i}(\omega)) | P_i(\omega) \cap (\mathcal{M}(\omega)_i = g)] \geq E[u_i(g', s_{-i}(\omega)) | P_i(\omega) \cap (\mathcal{M}(\omega)_i = g)].$$

Additionally, note that if $\mathcal{M}(\omega)_i = g$ for some $g \notin \text{supp}(s_i(\omega))$ then player i is certain that $g = g_\omega$, and so following the mediator's recommendation is optimal. Thus, in both cases, following the mediator's recommendation is optimal, and so \bar{s} is a BNE.

Finally, we show that \mathcal{M}_{noA}^E is optimal with respect to E . In any state ω , the sum of players' utilities conditional on that state is at most 1. This maximal utility is achieved whenever at least one player plays g_ω . Under \mathcal{M}_{noA}^E , at least one player plays g_ω in *every* ω . Hence, the sum of players' (unconditional) expected utilities is $W(\mathcal{M}_{noA}^E) = 1$, the maximum possible in the game. ■

A.2.2 MANY PLAYERS AGAINST AMAZON

Suppose now that there are n players and an Amazon, that there is jointly complete information, and that v is an arbitrary profile of base values.

Claim 3 *The mediator \mathcal{M} that is fully revealing to all players $i \in \{1, \dots, n\}$ and always recommends g_ω is IR, IC, and optimal if and only if $v_i \leq 1/(n+1)$ for all i .*

Proof: The strategy \bar{s} in a fully revealing \mathcal{M} is a BNE, since each player plays g_ω , which is dominant. It is IR, since each player i 's expected utility in \mathcal{M} is $1/(n+1) \geq v_i$. Finally, it is optimal: Whenever k players correctly choose g_ω , the sum of players' utilities is $k/(k+1)$. Under \mathcal{M} , the sum of players' utilities is $n/(n+1)$, which is the most they can jointly obtain.

The reverse direction holds by the observation that if a player has base value $v_i > 1/(n+1)$ then her IR constraint is violated under a fully revealing mediator. ■

References

- Anderson, C. (2006). *The Long Tail: Why the Future of Business is Selling Less of More*. Hachette Books.
- Aumann, R. J. (1976). Agreeing to disagree. *The Annals of Statistics*, 4(6), 1236–1239.
- Ben-Porat, O., & Tennenholtz, M. (2019). Regression equilibrium. In *Proceedings of the ACM Conference on Economics and Computation*, pp. 173–191.
- Bengtsson, M., & Kock, S. (2000). “coopetition” in business networks – to cooperate and compete simultaneously. *Industrial Marketing Management*, 29(5), 411–426.
- Bergemann, D., & Bonatti, A. (2019). Markets for information: An introduction. *Annual Review of Economics*, 11, 85–107.

- Bergemann, D., & Morris, S. (2013). Robust predictions in games with incomplete information. *Econometrica*, *81*(4), 1251–1308.
- Bergemann, D., & Morris, S. (2016). Bayes correlated equilibrium and the comparison of information structures in games. *Theoretical Economics*, *11*(2), 487–522.
- Bergemann, D., & Morris, S. (2019). Information design: A unified perspective. *Journal of Economic Literature*, *57*(1), 44–95.
- Bimpikis, K., Crapis, D., & Tahbaz-Salehi, A. (2019). Information sale and competition. *Management Science*, *65*(6), 2646–2664.
- Brandenburger, A. M., & Nalebuff, B. J. (2011). *Co-opetition*. Crown Business.
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. (2020). Language models are few-shot learners. In *Advances in Neural Information Processing Systems*, Vol. 33, pp. 1877–1901.
- Chen, Y., Nissim, K., & Waggoner, B. (2015). Fair information sharing for treasure hunting. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Clarke, R. N. (1983). Collusion and the incentives for information sharing. *The Bell Journal of Economics*, *14*(2), 383–394.
- de Zegher, J. F., & Lo, I. (2020). Crowdsourcing market information from competitors. Available at SSRN 3537625.
- Dutta, P. S., Jennings, N. R., & Moreau, L. (2005). Cooperative information sharing to improve distributed learning in multi-agent systems. *Journal of Artificial Intelligence Research*, *24*, 407–463.
- European Commission (2020). A European strategy for data. <https://eur-lex.europa.eu/legal-content/EN/TXT/?qid=1593073685620&uri=CELEX:52020DC0066>. Accessed: 2021-05-13.
- Feng, Y., Gradwohl, R., Hartline, J., Johnsen, A., & Nekipelov, D. (2022). Bias-variance games. In *Proceedings of the ACM Conference on Economics and Computation*, pp. 328–329.
- Forges, F. (1993). Five legitimate definitions of correlated equilibrium in games with incomplete information. *Theory and Decision*, *35*(3), 277–310.
- Galperti, S., & Perego, J. (2018). A dual perspective on information design. Available at SSRN 3297406.
- Gehrig, T., & Stenbacka, R. (2007). Information sharing and lending market competition with switching costs and poaching. *European Economic Review*, *51*(1), 77–99.
- Gentzkow, M., & Kamenica, E. (2016). Competition in persuasion. *The Review of Economic Studies*, *84*(1), 300–322.
- Ginart, T., Zhang, E., Kwon, Y., & Zou, J. (2021). Competing AI: How does competition feedback affect machine learning?. In *Proceedings of the International Conference on Artificial Intelligence and Statistics*, pp. 1693–1701.
- Gu, Y., Madio, L., & Reggiani, C. (2019). Data brokers co-opetition. Available at SSRN 3308384.

- Ha, A. Y., & Tong, S. (2008). Contracting and information sharing under supply chain competition. *Management Science*, 54(4), 701–715.
- Ichihashi, S. (2020). Competing data intermediaries. *Manuscript, available at <https://shota2.github.io/research/data.pdf>*.
- Jappelli, T., & Pagano, M. (2002). Information sharing, lending and defaults: Cross-country evidence. *Journal of Banking & Finance*, 26(10), 2017–2045.
- Jentzsch, N., Sapi, G., & Suleymanova, I. (2013). Targeted pricing and customer data sharing among rivals. *International Journal of Industrial Organization*, 31(2), 131–144.
- Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., et al. (2019). Advances and open problems in federated learning. *arXiv preprint arXiv:1912.04977*.
- Kamenica, E., & Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, 101(6), 2590–2615.
- Kearns, M., Pai, M., Roth, A., & Ullman, J. (2014). Mechanism design in large games: Incentives and privacy. In *Proceedings of the Conference on Innovations in Theoretical Computer Science*, pp. 403–410.
- Kwon, Y., Ginart, A., & Zou, J. (2022). Competition over data: how does data purchase affect users?. *arXiv preprint arXiv:2201.10774*.
- Lester, B., Al-Rfou, R., & Constant, N. (2021). The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691*.
- Liu, Q., & Serfes, K. (2006). Customer information sharing among rival firms. *European Economic Review*, 50(6), 1571–1600.
- Loshen, D. (2014). Leveraging cloud-based analytics for crowd-sourced intelligence. <https://gigaom.com/report/leveraging-cloud-based-analytics-for-crowdsourced-intelligence/>. [Gigaom Research; accessed 6-June-2020].
- Lyu, L., Yu, J., Nandakumar, K., Li, Y., Ma, X., Jin, J., Yu, H., & Ng, K. S. (2020). Towards fair and privacy-preserving federated deep models. *IEEE Transactions on Parallel and Distributed Systems*, 31(11), 2524–2541.
- Mansour, Y., Slivkins, A., & Wu, Z. S. (2018). Competing bandits: Learning under competition. In *Proceedings of the Conference on Innovations in Theoretical Computer Science*.
- Martens, B., De Streel, A., Graef, I., Tombal, T., & Duch-Brown, N. (2020). Business-to-business data sharing: An economic and legal analysis. *European Commission, JRC Digital Economy Working Paper Series, 2020(05)*. <https://ec.europa.eu/jrc/sites/jrcsh/files/jrc121336.pdf>. Accessed: 2021-05-13.
- Mathevet, L., Perego, J., & Taneva, I. (2020). On information design in games. *Journal of Political Economy*, 128(4), 1370–1404.
- Monderer, D., & Tennenholtz, M. (2004). k-implementation. *Journal of Artificial Intelligence Research*, 21, 37–62.

- Monderer, D., & Tennenholtz, M. (2009). Strong mediated equilibrium. *Artificial Intelligence*, 173(1), 180–195.
- Myerson, R. B. (1991). *Game Theory: Analysis of Conflict*. Harvard University Press.
- Pagano, M., & Jappelli, T. (1993). Information sharing in credit markets. *The Journal of Finance*, 48(5), 1693–1718.
- Pathak, B., Garfinkel, R., Gopal, R. D., Venkatesan, R., & Yin, F. (2010). Empirical analysis of the impact of recommender systems on sales. *Journal of Management Information Systems*, 27(2), 159–188.
- Raith, M. (1996). A general model of information sharing in oligopoly. *Journal of Economic Theory*, 71(1), 260–288.
- Rochlin, I., & Sarne, D. (2015). Constraining information sharing to improve cooperative information gathering. *Journal of Artificial Intelligence Research*, 54, 437–469.
- Shamir, N., & Shin, H. (2016). Public forecast information sharing in a market with competing supply chains. *Management Science*, 62(10), 2994–3022.
- Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, 63(2), 129.
- Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2019). Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 10(2), 1–19.
- Zhan, Y., Zhang, J., Hong, Z., Wu, L., Li, P., & Guo, S. (2021). A survey of incentive mechanism design for federated learning. *IEEE Transactions on Emerging Topics in Computing*, 10(2), 1035–1044.