

Information markets and nonmarkets[☆]

8

Dirk Bergemann^{a,*} and Marco Ottaviani^b

^a*Department of Economics, Yale University, New Haven, CT, United States*

^b*Department of Economics, Bocconi University, Milan, Italy*

*Corresponding author: e-mail address: dirk.bergemann@yale.edu

CONTENTS

1 Introduction	594
2 Buying and selling information	597
2.1 Value of information and experiment	597
2.2 Selling experiments	600
2.3 Selling realizations	602
2.4 Information products.....	604
2.5 Returns from information and data	605
3 Markets for information	606
3.1 Value of information in a normal quadratic environment	606
3.2 Selling information to competing firms	608
3.3 Information sharing among competing firms	610
3.4 Data intermediation	610
3.5 Data markets and data externalities	612
4 Instruments to trade and monetize information	614
4.1 Third-degree price discrimination.....	614
4.2 Ratchet effect and privacy	619
4.3 Ratings and recommender systems	621
4.4 Certification and expert markets	623
5 Forecasting and aggregation of predictions	624
5.1 Forecasting nonmarkets	624
5.1.1 Scoring rules.....	625
5.1.2 Forecasting contests	626
5.1.3 Reputational forecasting	626
5.1.4 Expert aggregation.....	627
5.2 Prediction markets	628

[☆] We thank Alessandro Bonatti, Ben Brooks, Anna Dreber, Tibor Heumann, Felix Holzmeister, Stephen Morris, Max Schaefer, Adam Siegel, and Christoph Siemroth for useful discussion. We thank two anonymous referees for detailed comments and suggestions. We thank James Atkins, Eugenio Piga, Michael Sullivan, and Shawn Thacker for excellent research assistance and the reviewers for constructive comments. Dirk Bergemann acknowledges financial support from NSF Grant SES 1948336, Sloan Foundation, and Omidyar Network.

5.2.1	<i>Equilibrium with heterogeneous beliefs and limited wealth</i>	629
5.2.2	<i>Equilibrium with heterogeneous beliefs and risk-averse traders</i>	631
5.2.3	<i>Aggregation of information and beliefs</i>	633
5.3	Automated market maker	636
5.4	Performance of prediction markets	637
6	Science as nonmarket	644
6.1	Organization of science	645
6.2	Measuring influence	647
6.3	Selection, publication, and funding	648
6.4	Economics of statistical inference	649
7	Conclusion	656
	References	657

Abstract

As large amounts of data become available and can be communicated more easily and processed more effectively, information has come to play a central role for economic activity and welfare in our age. This essay overviews contributions to the industrial organization of information markets and nonmarkets, while attempting to maintain a balance between foundational frameworks and more recent developments. We start by reviewing mechanism-design approaches to modeling the trade of information. We then cover ratings, predictions, and recommender systems. We turn to forecasting contests, prediction markets, and other institutions designed for collecting and aggregating information from decentralized participants. Finally, we discuss science as a prototypical information nonmarket with participants who interact in a non-anonymous way to produce and disseminate information. We aim to familiarize the reader with the central notions and insights in this burgeoning literature and also point to some critical open questions that future research will have to address.

Keywords

Information, Data, Data intermediaries, Information markets, Information nonmarkets, Science

1 Introduction

As large amounts of data become available and can be communicated more easily and processed more effectively, information has come to play a central role for economic activity and welfare in our age. In turn, markets for information where data is exchanged have become more prominent and significant in terms of trading volume. Trading information, however, is not merely about buying and selling access to a database. The ability to collect, mine, and analyze large datasets creates opportunities for exchanging information in the form of predictions, ratings, recommendations, as well as for customizing products and services. At the same time, the mechanisms

for trading information pose new challenges related to privacy, market power of information intermediaries, and the potential for distortions in almost any sector of the economy affected by the information and communication revolution.¹

The design of institutions for trading information requires attention to the peculiar features of information. An information good can be defined as anything that can be encoded, or digitized, as a stream of bits; see Shapiro and Varian (1999). In addition to data and, more generally, knowledge that can help decision making, prototypical examples of information goods are books, magazines, movies, music, and web pages. As recognized at least since Arrow (1962), information goods are characterized by a number of peculiar features:

1. **Uncertainty.** Information goods are typically hard to evaluate. For example, it is often difficult for a user to ascertain the quality of an innovative idea before applying it—and sometimes even long after that.
2. **Indivisibility.** Information goods are easy to duplicate. On the supply side, production of information typically entails a high fixed cost—typically non-reversible and thus sunk. After the first unit, however, it is cheap to duplicate information at near-zero marginal cost and with virtually infinite capacity. Said differently, information goods are non-rival: use by one consumer does not prevent use by other consumers.
3. **Inappropriability.** Information goods are difficult to exclude. Once information becomes available to one user, it is costly to prevent other users—or competing sellers—from having access to it.

These features of information goods create difficulties for market trading that must be overcome by institutional arrangements. By feature (1), sellers might pretend to be offering something more valuable than they really have. And buyers may only be confident in the value of the information they purchase after seeing it—yet at that stage, buyers might be in a position to expropriate the seller. Cost conditions given by (2) could lead to a natural monopoly if the information seller could perfectly control duplication—even though exclusion would be socially inefficient *ex post*. However, because of feature (3), buyers who cannot be excluded can duplicate the good themselves and share it with other consumers, thus competing with the seller.

Because of its public good nature linked to features (2) and (3), information “wants to be free.” Once information is produced, it is socially efficient to make it freely available. But if information is made free once it is produced, we have a classic hold-up problem: expecting no reward, the investment in information will not be made. So, the market can easily break down without judicious help from careful institutional design. The alternative of socializing the production of information makes

¹ The recent Furman report identifies “the central importance of data as a driver of concentration and a barrier to competition in digital markets” (Digital Competition Expert Panel (2019))—a theme echoed in the reports by Cr mer et al. (2019) and by the Stigler Committee on Digital Platforms (2019). Tirole (2020) reviews the challenges to economic policy posed by the rapid growth of the digital platforms.

the evaluation problem (1) more pressing. Who will determine whether the information (or the innovation that embodies it) is valuable? A high level of state capacity is needed if we are to entrust central government boards to pick winners.²

And while information is traded in market environments, trading may only be effective if it relies on sophisticated contracts or reputation; see, for example, Anton and Yao (1994), Varian (2000), and Henry and Ponce (2011). Given the difficulties of information markets, organizational forms alternative to the market have also developed for the production and dissemination of information. In a prominent example of information nonmarkets, the science reward system leverages the willingness of scientists and their communities to contribute to public knowledge.

This chapter overviews contributions to the industrial organization of information markets and nonmarkets, while attempting to maintain a balance between foundational frameworks and more recent developments. We aim to make the reader familiar with the central notions and insights in this literature and also point to some critical open questions that future research will have to address.

Marschak (1968) envisioned a comprehensive perspective on the economics of information: (i) data are gathered, (ii) data are communicated, and (iii) on the basis of the messages, decisions are made. In information theory there is a sharp separation between Shannon's (1948) theory of communication (focusing on efficient coding and transmission) and statistical decision theory (centered on the value of information for decision making). The work we review here squarely builds on the classic work by Blackwell (1951, 1953) on the comparison of experiments; for ties to the theory of communication we refer to the recent overview by Maćkowiak et al. (forthcoming).

There is a recent, yet distinct literature on Bayesian persuasion and information design that analyzes how a sender/designer can achieve an objective only by committing to a policy of information revelation; we refer the interested reader to the recent survey by Bergemann and Morris (2019). The notion of Bayesian persuasion adds commitment to the classic question of strategic information communication as posed by Crawford and Sobel (1982); we shall largely bypass this literature as well.

Section 2 proceeds by considering a variety of mechanism-design inspired approaches to buying and selling information. We adopt the perspective that information is an input into a (strategic) decision problem and study the optimal sale of supplemental information to heterogeneous, privately informed agents. In doing so, we distinguish between *ex ante* and *ex post* sale of information. Section 3 offers an introduction to the organization of markets for information. Section 4 discusses instruments that facilitate or hinder the trade of information. Here we discuss third-degree price discrimination and the ratchet effect associated with using information for price discrimination purposes. Then we turn to the role of ratings, predictions, and recommender systems as markets for indirect information.

² Yet, to complicate the matter further, there are important exceptions when more public information may be detrimental to social welfare, for example it may preclude the possibility of mutual insurance, see Hirshleifer (1971).

Section 5 turns to institutional arrangements for collecting information and aggregating expert judgments. Forecasting contests and prediction markets are specific institutions designed for collecting and aggregating information from decentralized participants. At least since Hayek (1945), economists have long praised the ability of markets to aggregate the dribs and drabs of information that are available to different individuals. However, close theoretical and empirical scrutiny reveals a much more nuanced picture. Section 6 analyzes science as a form of production of information and knowledge that is organized as a nonmarket where evaluation and non-pecuniary incentives play a significant role. Section 7 concludes. Even though the literature on the economics of information has grown rapidly as the internet and digitalization have made information central to many markets over the last decades, plenty of important questions remain unanswered.

2 Buying and selling information

We begin with the classic decision-theoretic model of Blackwell (1951, 1953). Within this framework, we examine the incentives to buy and sell information in a canonical and comprehensive model—one with a single seller and a single buyer of information. As a by-product of the analysis, we will be able to speak about the optimal information structure. Then, we discuss a variety of mechanism design inspired approaches to selling information when data buyers are privately informed about their beliefs or preferences. We focus first on the direct sale of information in which contracting takes place at the *ex ante* stage. Here the buyer purchases an information structure (i.e., a Blackwell experiment), as opposed to paying for specific realizations of the seller’s informative signals. We then turn to different contracting assumptions that extend the analysis to the sale of individual signal realizations and to the indirect sale of information.

2.1 Value of information and experiment

The setting can be described as follows. A decision maker, the buyer of information (or data), faces a decision problem under uncertainty. The state of nature ω is drawn from a finite set $\Omega = \{\omega_1, \dots, \omega_i, \dots, \omega_I\}$. The buyer chooses an action a from a finite set $A = \{a_1, \dots, a_j, \dots, a_J\}$. The *ex post* utility is denoted by

$$u(\omega_i, a_j) \triangleq u_{ij} \in \mathbb{R}_+. \quad (1)$$

The *ex post* payoffs can thus be represented by an $I \times J$ matrix:

$$\begin{array}{c|ccc}
 u & a_1 & \cdots & a_J \\
 \hline
 \omega_1 & u_{11} & \cdots & u_{1J} \\
 \vdots & \vdots & & \vdots \\
 \omega_I & u_{I1} & \cdots & u_{IJ}
 \end{array} . \quad (2)$$

The decision maker has an *interim belief* θ about the state, with $\theta = (\theta_1, \dots, \theta_I) \in \Delta\Omega$, where θ_i denotes the interim probability that the decision maker assigns to state ω_i . This decision-theoretic framework for the Blackwell comparison of experiments is also covered by Blackwell and Girshick (1954) and Laffont (1989).

We augment this framework to account for a seller of information who can sell additional or supplemental information to the decision maker, the buyer. In line with the interpretation of selling supplemental information, the *initial belief* $\theta \in \Delta\Omega$ of the buyer can be interpreted as being generated from the common prior and privately observed signals. Thus, suppose there is a common prior $\mu \in \Delta\Omega$ shared by the buyer and seller. The buyer privately observes a signal $r \in R$ according to a commonly known experiment $\lambda : \Omega \rightarrow \Delta R$. The buyer then forms his interim belief via Bayes' rule

$$\theta(\omega | r) \triangleq \frac{\lambda(r | \omega) \mu(\omega)}{\sum_{\omega' \in \Omega} \lambda(r | \omega') \mu(\omega')}.$$

The interim beliefs $\theta(\omega | r)$, simply denoted by θ , are thus the private information of the buyer. From the seller's perspective, the common prior $\mu \in \Delta\Omega$ and the distribution of signals $\lambda : \Omega \rightarrow \Delta R$ induce a distribution $F \in \Delta(\Delta\Omega)$ of interim beliefs.

The buyer seeks to augment his initial private information by obtaining additional information from the seller in order to improve the quality of his decision making. A statistical *experiment* (equivalently, an information structure) $E = (S, \pi)$ consists of a set S of signals s and a likelihood function

$$\pi : \Omega \rightarrow \Delta S. \tag{3}$$

For a given experiment $E = (S, \pi)$, let S denote the (finite) set of signals that are in the support of the experiment and π_{ik} the conditional probability of signal $s_k \in S$ in state ω_i . Letting $K \triangleq |S|$, we have

$$\pi_{ik} \triangleq \Pr[s_k | \omega_i],$$

where $\pi_{ik} \geq 0$ and $\sum_{k=1}^K \pi_{ik} = 1$ for all i . We then obtain the stochastic matrix

$$\begin{array}{c|cccc} E & s_1 & \cdots & s_k & \cdots & s_K \\ \hline \omega_1 & \pi_{11} & \cdots & \pi_{1k} & \cdots & \pi_{1K} \\ \vdots & \vdots & & & & \vdots \\ \omega_i & \pi_{i1} & & \pi_{ik} & & \pi_{iK} \\ \vdots & \vdots & & & & \vdots \\ \omega_I & \pi_{I1} & \cdots & \pi_{Ik} & \cdots & \pi_{IK} \end{array}. \tag{4}$$

We take the realization of the buyer's private signal $r \in R$ and that of the signal $s \in S$ from any experiment E as independent, conditional on the state ω . In other words, the buyer and the seller draw their information from independent sources.

We first describe the value of the buyer's initial information and then determine the incremental value of an experiment $E = (S, \pi)$. The value of the buyer's problem under the private information generated by λ *only* is given by choosing an action $a(\theta)$ that maximizes the expected utility given the interim belief θ , i.e.,

$$a(\theta) \in \arg \max_{a_j \in A} \left\{ \sum_{i=1}^I \theta_i u_{ij} \right\}.$$

The expected utility from interim belief θ is therefore given by

$$u(\theta) \triangleq \max_{a_j \in A} \left\{ \sum_{i=1}^I \theta_i u_{ij} \right\}.$$

By contrast, if the buyer has access to an experiment $E = (S, \pi)$, he observes the signal realization $s_k \in S$, updates his beliefs, and then chooses an appropriate action. The marginal distribution of signals s_k from the perspective of type θ is given by

$$\Pr[s_k | \theta] = \sum_{i=1}^I \theta_i \pi_{ik},$$

where π_{ik} is the conditional probability that signal k is generated given the realization of state ω_i . Consequently, for any signal s_k that occurs with strictly positive probability $\Pr[s_k | \theta]$, an action that maximizes the expected utility of type θ is given by

$$a(s_k | \theta) \in \arg \max_{a_j \in A} \left\{ \sum_{i=1}^I \left(\frac{\theta_i \pi_{ik}}{\sum_{i'=1}^I \theta_{i'} \pi_{i'k}} \right) u_{ij} \right\}, \quad (5)$$

which leads to a conditional expected utility equal to

$$u(s_k | \theta) \triangleq \max_{a_j \in A} \left\{ \sum_{i=1}^I \left(\frac{\theta_i \pi_{ik}}{\sum_{i'=1}^I \theta_{i'} \pi_{i'k}} \right) u_{ij} \right\}. \quad (6)$$

Integrating over all signal realizations s_k and subtracting the expected utility obtained with the initial information, the (net) value of experiment E for type θ is

$$\begin{aligned} V(E, \theta) &\triangleq \mathbb{E}[u(s | \theta)] - u(\theta) \\ &= \sum_{k=1}^K \max_{a_j \in A} \left\{ \sum_{i=1}^I \theta_i \pi_{ik} u_{ij} \right\} - \max_{a_j \in A} \left\{ \sum_{i=1}^I \theta_i u_{ij} \right\}. \end{aligned} \quad (7)$$

The value of the initial information, given by $\max_{a_j \in A} \left\{ \sum_{i=1}^I \theta_i u_{ij} \right\}$, is generated by the action that has the highest probability-weighted utility. The value of an experi-

ment E is generated by choosing an action on the basis of the posterior belief induced by each signal s_k .³

The problem that we describe here was originally formulated by Raiffa and Schlaifer (1967), Section 5.5, as the expected net gain of an experiment.⁴ The conceptual core of this section is the notion of a statistical experiment and the resulting value of information. Section 2.2 then presents results of how experiments are optimally designed and priced. Section 2.3 suggests how an experiment may be segmented into realizations that can then be traded individually rather than as a bundle in the form of an experiment. Section 2.4 illustrates how statistical experiments generate information products that become marketable products. Finally, Section 2.5 presents some empirical work that aims to estimate and evaluate the value of information and data.

2.2 Selling experiments

In the context of buying and selling information, we can interpret the model given above by (1)–(3) as one where a data buyer faces a decision problem under uncertainty. A data seller owns a database containing information about a “state” variable that is relevant to the buyer’s decision. Initially, the data buyer has only partial information about the state. This information is private to the data buyer and unknown to the data seller. The precision of the buyer’s private information determines his willingness to pay for any supplemental information. Thus, from the perspective of the data seller, there are many possible types of the data buyer.

Bergemann et al. (2018) investigate the revenue-maximizing information policy, i.e., how much information the data seller should provide and how she should price access to the data. In order to screen the heterogeneous data buyer types, it is optimal for the seller to offer a menu of information products. In the present context, these products are statistical experiments—signals that reveal information about the payoff-relevant state. Only the information product itself is assumed to be contractible. By contrast, payments *cannot* be made contingent on either the buyer’s action or the realized state and signal. Consequently, the value of an experiment to a buyer is determined by the buyer’s private belief and can be computed independently of the price of the experiment. Bergemann et al. (2018) recast the resulting screening problem as a nonlinear pricing problem wherein the buyer’s type is given by his prior belief. In other words, the seller’s problem is to *design* and *price* different versions of experiments, that is, different information products from the same underlying database.

The optimal menu of experiments has the following two structural properties: (i) the fully informative experiment is part of an optimal menu; and (ii) every ex-

³ The value of information for the data buyer differs from a consumer’s value for multiple goods or bundles of characteristics. In particular, the first max operator in (7) corresponds to the optimality condition for the buyer’s action given the available information. The second max operator reflects the type-dependent nature of participation constraints.

⁴ Raiffa and Schlaifer (1967) attribute the specific results for normal priors and normal signals to Grundy et al. (1956) and the minimax solution to the same problem is attributed to Somerville (1954).

periment in any optimal menu is non-dispersed, i.e., $\pi_{ij} = 0$ for some $i \neq j$. Thus, the optimal menu includes the efficient information structure, and every item in the optimal menu eliminates the data buyer's uncertainty along some dimension.

The design of information can be rephrased in terms of hypothesis testing and the analysis can be interpreted as a pricing model for statistical tests. A useful frame is given by Bayesian hypothesis testing. To illustrate this point, consider the example of a lender who must decide whether to grant a loan to a prospective borrower at the prevailing market rate. The lender wants to test a null hypothesis H_0 (borrower is low-risk) against an alternative H_1 (high risk). The central issue for the data seller is that she does not know the data buyer's interim beliefs and, hence, the buyer's willingness to pay for this information. The seller can design any binary ("pass/fail") test that reports whether the riskiness measure is above or below a particular threshold. Each test is intended for a different buyer type θ , and yields a different combination of type I and type II *statistical errors*. We can represent this experiment as a statistical test of the null hypothesis $H_0 = \{\omega_2\}$,

E	s_1	s_2
ω_1	$1 - \beta$	β
ω_2	α	$1 - \alpha$

where α and β denote the probability of a type I and type II statistical error, respectively.

The main idea behind the revenue-maximizing mechanism for the information seller is akin to offering "damaged goods" to low-value buyers. However, when selling information goods—see Shapiro and Varian (1999)—product versioning allows for richer and more profitable distortions than when selling physical goods. This is due to a peculiar property of information products: because buyers value different dimensions (i.e., information about specific state realizations), the buyers with the lowest willingness to pay also have very specific preferences. In the context of credit markets, very aggressive lenders are interested in very negative information only, and are willing to grant a loan otherwise.

The seller can thus screen the buyer's private information by leveraging a key insight—that information is only valuable if it changes optimal actions. A complete characterization of the optimal selling mechanism with many states and many actions is currently not known. The data seller's problem bears some resemblance to the canonical bundling problem with many items. With more than two states, the types of a buyer are multidimensional and it is well-known—see, for example, Pavlov (2011)—that the optimality of single-price result of Myerson (1981) and Riley and Zeckhauser (1983) does not hold. Indeed, the optimal menu involves stochastic bundling quite generally, and the structure of the bundles offered can be quite rich.⁵ Stochastic bundles are analogous to the partially informative experiments in

⁵ Daskalakis et al. (2017) construct an example wherein the types follow a Beta distribution and the optimal menu contains a continuum of stochastic allocations.

our model. To further distinguish the current model from these classic multidimensional problems, it is noteworthy that stochastic bundling can arise in the current setting even when the types of the buyer are one-dimensional. Given the computational complexity of the optimal menu, Bergemann et al. (2021b) provide sufficient conditions when selling complete information is approximately optimal.

There is a substantial recent literature on the value of information in health economics where it appears as the expected value of sample information, see Ades et al. (2004). A new issue here is that the set of feasible statistical experiments may not be unconstrained, but rather constrained by the available tests and treatments. Drakopoulos and Randhawa (2020) and Ely et al. (2021) compare the value of a limited set of imperfect medical tests. In a model with two actions (social distancing or not) and two states, Drakopoulos and Randhawa (2020) characterize the optimal test within the available supply of tests. Ely et al. (2021) consider a related allocation problem among finitely many tests and seek to maximize the social welfare, abstracting away from incentive issues.

We will discuss in Section 2.4 how to bring the results of the above analyses to bear on the design of real-world information products. We will focus on offline and online brokers of big data—firms such as Acxiom, Nielsen, and Oracle—that sell information about individual consumers to business customers. Information used for marketing purposes is typically sold through data appends and marketing lists. Data appends reveal supplemental information about a firm’s existing or potential customers, allowing the firm to place them into more precise market segments. For instance, all major data brokers offer data management platforms (DMPs), customized software that enables websites to track their users and integrate their own data with third-party data. Most risk-mitigation data products offered by credit rating agencies are also of this kind.

2.3 Selling realizations

We have so far focused on the sale of *data appends* in the form of (ex ante) information structures. In contrast, the sale of *original lists* can be modeled as an informative experiment that reveals whether a potential consumer matches a pre-specified set of characteristics, in which case the buyer receives a contact and pays a price. This is true both when an *original list* is sold directly (e.g., in the case of information about ACT test takers) and when it is sold indirectly (as in the case of sponsored search or targeted display advertising). In these cases, the price paid by the buyer depends on the realization of the seller’s information.

Bergemann and Bonatti (2015) considers the trade of information bits (“cookies”) that are an input into a decision problem. In particular, a single firm (a buyer of information) has heterogeneous match values with a set of consumers. In order to realize the potential match value, the firm must choose a continuous investment level. The optimal investment level (e.g., advertising spending) depends on the consumer’s match value v . To capture the role of browser cookies they consider a special information structure, namely one in which individual consumers’ types are learned perfectly

or not at all. Through the purchase of information, the firm is then able to segment consumers into a *targeted* group that receives personalized levels of advertising, and a *residual* set that receives a uniform level of advertising.

Under conditions on the matching technology and on the distribution of match values, the data-buying policy takes the form of a single cutoff match value. That is, advertisers buy information about all users with match values above a cutoff (*positive targeting*) or below the cutoff (*negative targeting*).

Babaioff et al. (2012) study a related model of selling lists (i.e., pricing conditional on signal realizations) when buyers are heterogeneous and privately informed. In particular, the data buyer's value depends on two variables: one is known by the seller, while the other one is the buyer's type. The paper develops algorithms to characterize the optimal mechanism, and derives conditions under which the seller can extract the entire surplus by exploiting the correlation between their information and the buyer's type.

Eső and Szentes (2007a) as well as Li and Shi (2017) consider the case where signal realizations are not directly contractible, but the buyer's actions are. In these models, the seller of a good controls both its price and the information provided to the buyer, with the goal of screening their private, partial information. In the context of online markets, the seller is a provider of advertising space who can offer arbitrarily fine targeting criteria to advertisers. (Recall the earlier discussion of indirect sales of information through Facebook or Google advertising.)

Eső and Szentes (2007a) focus on the case where the seller releases information that is orthogonal to the buyer's type. (This is without loss if, for example, the buyer's type is a preference parameter, and the seller reveals information about the quality of the product.) The seller-optimal mechanism when a single buyer is present reveals all the information and offers a menu of European call options where a lower strike price costs more up front. In the case of competing buyers, a two-stage "handicap auction" is optimal. Intuitively, a positive strike price distorts the buyer's decisions, but the result suggests that it is more profitable to distort ex post decisions rather than the initial information. More recently, Li and Shi (2017) show that discriminatory disclosure of information—providing different buyer types with different signals—dominates full disclosure when the seller is not restricted to orthogonal disclosure.

More recently, Guo et al. (2020) establish that the optimal discriminatory disclosure policy consists of a pair of intervals. If the buyer's private information is correlated with the information controlled by the seller, then the optimal revenue generally cannot be attained by any selling mechanism without a discriminatory disclosure policy.

In many cases, an advertiser can use additional third-party data to refine the targeting criteria offered by a publisher. Eső and Szentes (2007b) consider a related model of selling advice. Reinterpreting their model, an advertiser buys information about a prospective consumer before deciding whether or not to advertise their product. As the transaction takes place contextually to the advertising campaign, the data buyer's action is contractible. In some special cases, the data seller discloses all the information to all buyer types. Distortions to the buyer's actions then come from a marginal

price of advice. In other words, the data seller grants access to her database (perhaps against a subscription fee) but charges a marginal price for the data only upon the buyer's investment. In practice, it is often the case that the advertiser is charged for data on a cost-per-mille (CPM) basis, in which case the price of data adds to the marginal cost of the advertising space.

Yang (2020) considers a richer model of data trade. A revenue-maximizing data broker sells market segmentations to a producer with private production cost who in turn sells a product to a unit mass of consumers with heterogeneous values. He completely characterizes the revenue-maximizing mechanisms for the data broker. In particular, every optimal mechanism induces quasi-perfect price discrimination—the data broker sells the producer a market segmentation described by a cost-dependent cutoff, such that all the consumers with values above the cutoff end up buying and paying their values while the rest of consumers do not buy. He thus considers a model with private information on the side of the buyers (and held by the platform) and private information by the firm about its costs of production. The data platform sells the firm segments of consumer markets that allows the firm to engage in third-degree price discrimination as characterized in the work of Bergemann et al. (2015).⁶

2.4 Information products

A recent report by the Federal Trade Commission (2014) classifies information products offered by data brokers regarding individual consumers along two key dimensions:

- Who identifies the prospective consumer? Is the data seller providing the data buyer with a new list of prospects? Or is the data seller appending information about an individual (or a group) that the buyer has already identified?
- Does the data seller provide information (direct sale) or access to a consumer (indirect sale)? In other words, does the data buyer have the means to independently contact the consumer? Or does the data seller provide an exclusive opportunity for the data buyer to reach a consumer?

In the terminology of the Federal Trade Commission (2014) report, *original lists* are the main object for sale by marketing and lead-generation companies, as well as by providers of financial data (e.g., Bloomberg). An original list is often simply a customer segment, i.e., a collection of potential consumers with certain characteristics. The *audience segments* sold by Nielsen, Acxiom, and Epsilon are the most common examples of such lists. Individual sites can also sell original lists. For example, Evite.com may sell lists of consumers attending a party in a given location, and

⁶ In practice, “selling” consumer data can take a wide variety of forms, which include not only traditional physical transactions but also integrated data-sharing agreements/activities. For instance, in a recent report by The New York Times, Facebook has formed ongoing partnerships with other firms, including Netflix, Spotify, Apple, and Microsoft, and granted these companies access to different aspects of consumer data “in ways that advanced its own interests,” see New York Times (2018).

AddThis may sell lists of consumers who have shared a given news article. *Data appendixes* reveal supplemental information about a firm's existing or potential customers. In the context of marketing, Nielsen Catalina Solutions and Oracle Datalogix connect an individual's offline and online purchases with the digital media they consume.

Bergemann and Bonatti (2019) provide an introduction to markets for information with a particular emphasis on data markets and data intermediation in e-commerce. In the following, we will frequently use examples from the economics of the internet as applications and illustrations.

2.5 Returns from information and data

Machine learning algorithms rely on the large-scale collection and algorithmic evaluation of data to perform prediction tasks. In the context of user-oriented applications, firms often aim at forecasting consumers' preferences and valuations in order to tailor content or to perform various forms of price discrimination. More broadly, data can be seen as a catalyst for learning and innovation.

Claussen et al. (2019) and Yoganarasimhan (2020) primarily focus on the returns from personalized data: Both studies show that additional data collected on an individual allow to improve the accuracy of personalized predictions at a diminishing rate. Chiou and Tucker (2017) and He et al. (2017) focus on the role of aggregate data for the average prediction accuracy across different users. Their findings are mixed: While Chiou and Tucker (2017) find no evidence for an improvement in average prediction accuracy, the findings of He et al. (2017) suggest substantial but diminishing returns to scale.

Other studies address potential synergies between personalized and collective data. Schaefer and Sapi (2020) document a complementarity between personal data and collective data using real-world search traffic data from Yahoo. The authors argue that the documented effect is consistent with personalized data continuously improving the technological capability of the firm. Lee and Wright (2021) analytically decompose returns from data in collaborative filtering algorithms. In an empirical simulation exercise, they provide evidence consistent with the effect documented by Schaefer and Sapi (2020).

Bajari et al. (2019) empirically analyze returns from data in the context of Amazon's retail forecast system and find diminishing returns to scale. The authors find no indication for complementarities between different data dimensions in improving the predictive value of the forecast.

Azevedo et al. (2020) analyze optimal experimentation policies in the context of A/B testing. The authors show that trying out more ideas is optimal if the distribution of innovation quality is fat-tailed. One implication of their findings is that returns from additional experimental data may only diminish at a slow rate, which suggests that the marginal value of data is economically significant even for large platforms. Jones and Tonetti (2020) theoretically motivate why the non-rival nature of data might lead to increasing returns to scale and show that allocating data property rights to consumers might be socially optimal as a consequence.

3 Markets for information

So far, we focused on the bilateral trade of information between a seller and a buyer. This allowed us to concentrate on the fundamental properties of information trade. In this section, we shall focus on market interactions where there is competition for information at least on one side of the market. The competition may arise as the information product of one seller is a substitute for the information product from another seller, or because the information privileges one market participants over another in a downstream market interaction.

We begin this section by describing how information may be sold by a monopolist when there are many competing buyers for the information. This may represent the practice of selling financial information. We then ask whether competing firms may have an incentive to share information. Finally, we discuss the role of information intermediaries such as data brokers and digital platforms.

In several markets, information is sold not only directly, but also indirectly in the form of customized goods and services. The case of carefully selected consumer segments is probably the best-known example of such a transaction. Consider the market for sponsored search advertising, e.g., on Google or Bing. The information held by the search engine consists first and foremost of the search query entered by the user. In practice, the search engine appends some of its own data too. The search engine then provides a prediction to the advertisers about the user's preferences, typically in form of an expected click-through rate. Aggregating over multiple users, this could be viewed as purchasing an original list of selected consumers. Of course, search engines adopt a different, more profitable strategy for selling their information: they grant access to the targeted population by selling advertising slots on specific keyword searches.

The indirect sale of information is not limited to advertising markets, either. Consider a monopolist seller of financial data, as in Admati and Pfleiderer (1990). As the sole owner of the information, the seller can either provide potential investors with informative signals about a stock, or she can construct a portfolio on the basis of her information. In both cases, the seller follows Blackwell's key insight, that data is only valuable insofar as it enables better decision making. The former is a direct sale, as the data buyers can buy the stock themselves. The latter is an indirect sale, because the data are never transferred and the data buyers must invest in the seller's portfolio instead. In other words, the seller can enable the buyer to take a better action without giving away the data.

3.1 Value of information in a normal quadratic environment

The literature that we cover in this section largely assumes a stochastic environment in which the fundamentals as well as the signals are given by multivariate normal distributions. In addition, the utility function of the agents are assumed to be either quadratic or exponential, which leads frequently to explicit characterizations and detailed comparative statics. It is therefore helpful to briefly introduce this environment before we discuss the results.

Accordingly, the willingness to pay of each data buyer is given by w_i :

$$w_i \triangleq \theta + \theta_i,$$

where the willingness to pay w_i of the data buyer i is the sum of a common and idiosyncratic component, θ and θ_i respectively.

Each data buyer maximizes a quadratic utility function:

$$u(w_i, q_i, p) \triangleq w_i q_i - p q_i - \frac{1}{2} q_i^2.$$

In the context of financial markets that we discuss next, the common component may represent the value of the asset and the idiosyncratic shock may represent the liquidity shock of trader i . The quantity q_i is then the quantity of the asset bought. In the context of information sharing among oligopolists, which we discuss next, the signs may be reversed, and w_i may represent the cost of production, and $p q_i$ may represent the revenue from selling the product.

At the outset, each data trader may not observe his true willingness to pay, but rather receives a noisy signal s_i . The privately observed signal s_i can include a common and an idiosyncratic shock, which we denote by ε and ε_i , respectively:

$$s_i = \theta + \varepsilon + \theta_i + \varepsilon_i, \quad (8)$$

and all the variables are jointly normally distributed:

$$\begin{pmatrix} \theta \\ \theta_i \\ \varepsilon \\ \varepsilon_i \end{pmatrix} \sim N \left(\begin{pmatrix} \mu_\theta \\ \mu_{\theta_i} \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_\theta^2 & 0 & 0 & 0 \\ 0 & \sigma_{\theta_i}^2 & 0 & 0 \\ 0 & 0 & \sigma_\varepsilon^2 & 0 \\ 0 & 0 & 0 & \sigma_{\varepsilon_i}^2 \end{pmatrix} \right). \quad (9)$$

The data structure as represented by (8) and (9) is sufficiently rich to represent the significant aspects of learning from social data. In particular, the noise in the signal s_i of each individual given by (8) reflects the importance of social learning from others, as might be enabled by recommender, rating, and search engines. This data environment has two important features. First, any demand information beyond the common prior comes from the signals of the individual consumers. Second, with any amount of noise in the signals (i.e., if $\sigma > 0$), each consumer can learn more about her own demand from the signals of the other consumers.

The following leading examples illustrate two ways in which data sharing can help each consumer learn more about her individual willingness to pay. In Example 1 a new product has a common value that consumers are imperfectly informed about.

Example 1 (Common Preferences). Fundamentals w_i are perfectly correlated and errors ε_i are independent: $s_i = w + \sigma \cdot \varepsilon_i$.

In this case, data sharing helps to filter out the idiosyncratic error terms. In particular, the average signal across all consumers identifies the common willingness to pay as N becomes large.

In Example 2, individual consumers have independent values, but are all exposed to a common shock.

Example 2 (Common Experience). Errors e_i are perfectly correlated and fundamentals w_i are independent: $s_i = w_i + \sigma \cdot \varepsilon$.

Under this structure, the average signal identifies the common error component ε as N becomes large. All market participants can then precisely estimate each w_i from the difference between individual and average signal.

As we shall see, while information sharing enables learning in both examples, the actions of consumers $-i$ impact the surplus of consumer i quite differently in the two cases, which has implications for the equilibrium price of data. More generally, the data structure will determine how to separate the individual and the aggregate information, see Bergemann et al. (2021a).

The joint prior distribution is commonly known by all market participants. The data seller collects the information in the form of the signals s_i from the individual traders, and then aggregates or sells it to the data buyers. The data buyers then use the information to improve their price and quantity policy. Sometimes the data seller possesses the information at the outset, and sometimes the data seller is more appropriately modeled as the data intermediary. In this case, the intermediary does not initially possess any information on her own, but rather collects the data from the data traders and then redistributes it among the data buyers.

3.2 Selling information to competing firms

An earlier literature, beginning with the seminal contribution of Admati and Pfleiderer (1986), directly started with a model where traders buy information from a monopolistic seller. From the outset, the data seller is assumed to be in possession of the information and hence in complete control of the entire database. Initially, the traders all share a common prior regarding the value of the asset. Each trader can acquire additional information regarding the value of the asset from the monopolistic seller. There is a continuum of traders, and each trader submits his demand function in response to his realized private information. The equilibrium price of the asset is determined in a speculative market formalized as a noisy rational expectations equilibrium. The true value of the asset is common to all the traders. The information seller therefore faces possible dilution in the value of information due to its leakage through informative prices. This leads to the impossibility of informationally-efficient markets, as argued by Grossman and Stiglitz (1980), who claim that if prices were to perfectly reflect all available information, then there would be no profit from gathering additional information.

The first set of results concerns the optimal selling policy of the information monopolist. The seller can (i) restrict access to the information, or not and (ii) can add

noise to the information, or not. Admati and Pfleiderer (1986) present conditions under which each one of these four possible information policies that arise from the combination of the above alternatives can be optimal. Then they consider the personalized sale of information. Here, the seller is allowed to add idiosyncratic noise to the common value signal for each trader. They show that the seller of information may prefer to sell noisier versions of the information than he actually has. Moreover, to obtain higher profits, it is desirable for the seller to sell different signals to different traders, so that the added noise realizations do not affect equilibrium prices. One way of doing so, which does not require discrimination, is to sell identically distributed personalized signals to each of a large number of traders.

Bergemann and Morris (2013) analyze the information structure that guarantees the highest industry profit in an oligopoly setting with incomplete information. Similar to Admati and Pfleiderer (1986), they find that if the strategic substitutes are sufficiently strong, then a noisy signal in which each firm learns the common value subject to idiosyncratic noise sustains the largest possible level of industry profits. In Admati and Pfleiderer (1986), the monopolistic seller in turn extracts the value of the industry profits by charging the individual traders for their private information.⁷

Admati and Pfleiderer (1990) extend their analysis to allow for two distinct methods of selling information. As before, they allow for the direct sale of information to the investors, but now they also allow the seller of the information to bundle the information with a product, in particular a portfolio whose composition depends on the available information. The analysis mostly considers a linear pricing policy for the portfolio and compares the revenue from a direct and indirect sale of information. They find that indirect sale is more profitable when the externality in the valuation of information is relatively intense.⁸

In an extension, they also consider the possibility that the seller can use a two-part tariff. Now, the indirect sale always dominates the direct sale. In an interesting discussion at the end of their paper, they also consider the possibility that the traders have different private information. In this case, the direct sale of information can improve the revenue as the seller can unbundle the initial information of the trader and the supplemental information.

Finally, Admati and Pfleiderer (1988) allow the seller of information to trade strategically on his own accounts as well. The information seller can now either trade, sell his information, or both. In either case, the seller commits to a policy in advance. Admati and Pfleiderer (1988) show that the optimal policy depends on the degree of risk aversion of the information buyers and of the information seller. In particular, if the buyer's risk aversion increases, the value of trading on the information decreases, and the value of selling information directly increases.

⁷ See also Bimpikis et al. (2019) on the nature of downstream competition and its implications for selling information in oligopolies.

⁸ As mentioned in the Introduction, the distinction between direct and indirect sale is similar to the distinction between pure information intermediaries and search engines or social platforms that jointly price information and access to the consumer.

3.3 Information sharing among competing firms

There is a large literature on information sharing among oligopolists, whose main results are succinctly presented in Raith (1996). The main question of this literature, which began with Novshek and Sonnenschein (1982), Clarke (1983), and Vives (1988) is whether competing firms, all with partial information, may have an incentive to share information through an intermediary, such as a trade association. Relative to this literature, the model of information markets we presented above has two important features. First, in the earlier models the information was collected and shared by an intermediary, such as a trade association, that merely organized and facilitated the exchange between the oligopolists, but had no genuine interest or market power. Second, the firms had all the information to begin with and did not have to collect the information from the consumer.

This earlier literature on information sharing leaves a limited role for information design. In particular, while the firms were allowed to add noise to their private information, the intermediary was restricted to simply aggregate and report the received information in the same format to all of the firms. The restrictiveness of this analysis was documented in Bergemann and Morris (2013). They investigated the role that private information of the competing firms can play for the realization of equilibrium values, prices and quantities, and the welfare of the market participants. Among other results, Bergemann and Morris (2013) identify the information structure that maximizes the industry profits as a function of the demand and supply conditions in the market. Similar to the earlier results of Admati and Pfleiderer (1986), they show that the optimal information structure has each individual firm receive private information with idiosyncratic noise that limits the correlation in the quantity choices by the firms.

There remain many interesting questions to be pursued. Even if jointly, but not individually, the firms already have all the relevant demand information, one might ask under which conditions could an intermediary profitably collect and redistribute information among the competing firms. In this respect, the credit rating and monitoring agencies serve the role of information intermediaries. The credit rating agencies both collect information about the borrowers and lenders from a given bank, as well as provide the bank with additional information about the credit worthiness of a new or established client. Thus, it both collects and redistributes demand information among the financial institutions.

3.4 Data intermediation

The role of data intermediaries has become more prominent with the rise of large data brokers and digital platforms. Data brokers are firms such as Acxiom, Nielsen, and Oracle that sell information about a consumer (or a group of consumers) to downstream data buyers, such as advertisers or retailers. An increasingly prominent use of third-party data occurs in display advertising, whether in (mobile) browsing, social media, or various video and formats. The display advertising market is a two-sided market where the publisher forms a match between competing advertisers (the

bidders) and a viewer as the object of the auction. A match between viewer and advertiser creates an *impression* (or search result) on the publisher’s website. The seller (a *publisher*) platform sells the attention (“eyeball”) of the viewer to competing advertisers. The viewer is thus the object of the digital advertising auction, see Levin and Milgrom (2010) and Bergemann et al. (2021c). In between, intermediaries facilitate the match between advertisers and publishers by managing data and providing algorithms to serve ads. Besides demand-side and supply-side platforms that provide the infrastructure for buying, selling, and serving ads, data intermediaries employ consumer data to efficiently target digital ads. This allows either advertisers or publishers (or possibly both) to augment their proprietary data with additional data, frequently acquired from a variety of sources. The data intermediaries, which in this context are often referred to as data management platforms, import data from multiple sources to then inform downstream allocation and bidding decisions; see Choi et al. (2020) for a recent survey of the display advertising market.

A conventional view of mechanism design with private information is that the principal elicits the private information from the agents and then implements an allocation. In independent private value settings, where each individual has private information about his value, the main goal is then to balance the individual objectives optimally.

A novel view of mechanism design is that the principal herself is in possession of some information regarding the fundamentals of each individual. The issue is then how to integrate this information into the mechanism. We may refer to these as augmented mechanisms; Hartline et al. (2019) use the evocative title of “dashboard mechanism.”

Emek et al. (2012) introduced a probabilistic single-item auction given by a matrix $V(i, j)$ which describes the valuation of bidder i for good j

$$V = \begin{array}{c|ccc} v_{ij} & g_1 & \cdots & g_J \\ \hline b_1 & v_{11} & \cdots & v_{1J} \\ \vdots & \vdots & & \vdots \\ b_I & v_{I1} & \cdots & v_{IJ} \end{array} . \quad (10)$$

The allocation problem is addressed by a second-price auction without reserve prices. Emek et al. (2012) initially assume that the valuation of each bidder is known but the identity of the object is not known. The question then arises of how to bundle the objects that appear stochastically to maximally raise the revenue of the seller.

Emek et al. (2012) refer to (10) as the known valuation case and to the case when there are finitely many matrices V_1, \dots, V_n as the Bayesian setting. The optimal information structure frequently bundles the objects in a stochastic manner. Ghosh et al. (2009) analyzed this problem under a restriction to deterministic allocation, referring to the deterministic bundling as clustering.

An alternative model of objects and preferences may be given by:

$$W = \begin{array}{c|ccc} v_{ij} & g_1 & \cdots & g_J \\ \hline w_1 & w_{11} & \cdots & w_{1J} \\ \vdots & \vdots & & \vdots \\ w_I & w_{I1} & \cdots & w_{IJ} \end{array} \quad (11)$$

where $w_{i \cdot}$ are the preferences of type i and $w_{\cdot j}$ are the possible values for object (or state) j . This interpretation was suggested by Bergemann et al. (2021c) to determine the allocation of informative signals from seller to bidder. Here, the bidder knows his type and the seller knows the identity of the object. These two models of preferences and objects given by (10) and (11) are closely related to the model introduced earlier in (1)–(3) regarding the value of information. The difference is that now the actions take a specific interpretation of an object allocation, and the choice of the information structure supports the allocation of the object in an auction mechanism. Bergemann and Pesendorfer (2007) investigated the joint determination of an optimal auction and an optimal information structure.

3.5 Data markets and data externalities

There is growing literature on data markets, some of it discussed in the previously mentioned survey, Bergemann and Bonatti (2019). The role of data externalities in the diffusion of personal data is a central concern in Choi et al. (2019) and Acemoglu et al. (2019). Bergemann et al. (2021a) develop a framework to evaluate the flow and allocation of individual data in the presence of data externalities. Their model focuses on three types of economic agents: consumers, firms, and data intermediaries. These agents interact in two distinct but linked markets: a *data market* and a *product market*.

In the product market, each consumer (she) determines the quantity that she wishes to purchase, and a single producer (he) sets the unit price at which he offers a product to the consumers. Initially, each consumer has private information about her willingness to pay for the firm's product. This information consists of a signal with two additive components: a *fundamental* component and a *noise* component. The fundamental component represents her willingness to pay, and the noise component reflects that her initial information might be imperfect. Both components can be correlated across consumers: in practice, different consumers' preferences can exhibit common traits, and consumers might undergo similar experiences that induce correlation in their errors. The social dimension of the data—whereby a consumer's data are also predictive of the behavior of others—is central to understanding the consumer's incentives to participate in the data market.

In the data market, a monopolist intermediary acquires demand information from the individual consumers in exchange for a monetary payment. The intermediary then chooses how much information to share with the other consumers and how much information to sell to the producer. In particular, sharing data with each consumer is analogous to providing a personalized purchase recommendation on the basis of

other consumers' signals. Selling data to the producer enables him to choose more precise, potentially personalized prices. Thus, the data intermediary has control over the volume and the structure of the information flows across all of the product market participants.

Choi et al. (2019) introduce information externalities into a model of monopoly pricing with unit demand. Each consumer is described by two *independent* random variables: her willingness to pay for the monopolist's service and her sensitivity to a loss of privacy. The purchase of the service by the consumer requires the transmission of personal data. From the collected data, the seller gains additional revenue, depending on the proportion of units sold and the volume of data collected. The total nuisance cost paid by each consumer depends on the total number of consumers sharing their personal data. In consequence, the optimal pricing policy of the monopolist yields excessive loss of privacy relative to the social welfare maximizing policy.

Acemoglu et al. (2019) also analyze data acquisition in the presence of information externalities. As in Choi et al. (2019), they consider a model with many consumers and a single data-acquiring firm. Acemoglu et al. (2019) propose an explicit statistical model for their data; the model allows them to assess the loss of privacy for the consumer and the gains in prediction for the firm. They analyze how consumers with heterogeneous privacy concerns trade information with a data platform. They derive conditions under which the equilibrium allocation of information is (in)efficient.

Galperti et al. (2021) consider a platform that mediates trade between a seller and a population of buyers using individual data records of their personal characteristics. After formulating this as an information-design problem, they use linear-programming duality to characterize the unit value that the platform derives from each buyer's specific record.

The low level of compensation that users command for their personal data is discussed in Arrieta-Ibarra et al. (2018), who propose sources of countervailing market power. An early and influential paper on consumer privacy is Taylor (2004)—to be discussed in more detail in the next section—which analyzes the sales of consumer purchase histories without data externalities.⁹ Fainmesser et al. (2020) provide a digital privacy model in which data collection improves the service provided to consumers. However, as the collected data can also leak to third parties and thus impose privacy costs, an optimal digital privacy policy must be established. Similarly, Jullien et al. (2020) analyze the equilibrium privacy policy of websites that monetize information collected from users by charging third parties for targeted access. Gradwohl (2017) considers a network game in which the level of beneficial information sharing among the players is limited by the possibility of leakage and a decrease in informational interdependence. Ali et al. (2019) study a model of personalized pricing with

⁹ Acquisti et al. (2016) provide a recent literature survey of the economics of privacy, and Goldfarb and Tucker (2019) discuss privacy in the broader context of digital economics. Fudenberg and Villas-Boas (2006) focus on behavior-based price discrimination.

disclosure by an informed consumer, and they analyze how different disclosure policies affect consumer surplus. Liang and Madsen (2020) investigate how data policies can provide incentives in principal-agent relationships. They emphasize the structure of individual data and whether individual signals are substitutes or complements determines the impact of data on incentives. Ichihashi (2020a) considers a single data intermediary and asks how the complement- or substitute-nature of the consumer signals affects the equilibrium price of the individual data. The notion of information complements and substitutes was introduced earlier by Börgers et al. (2013).

4 Instruments to trade and monetize information

We begin this section by discussing price discrimination as a commonly used instrument to monetize the value of information. This then brings us to the limits of how much information can be traded, in particular when consumers must be given incentives to generate or reveal information without direct monetary transfers for their data. Section 4.2 describes the ratchet effect and the problem of sourcing information from the consumer's actions. Section 4.3 illustrates how the use of ratings, recommender systems, and information aggregators determines the market's ability to obtain new information from consumers.

4.1 Third-degree price discrimination

Price discrimination is an important and central use of additional information about demand. We shall focus our discussion on third-degree price discrimination.¹⁰ The large literature on third-degree price discrimination starting with the classic work of Pigou (1920) examines what happens to prices, quantities, and various measures of welfare as the market is segmented. As every segment is offered a different price, there is scope for the producer to extract more surplus from the consumer. Yet if the producer can tailor the price to each segment, more consumers might be reached and there might be less exclusion. With the increase in available information about consumer demand comes increasing flexibility in the ensuing market segmentation: the platform that provides the data or the product seller can to a large extent determine how to optimally segment a given aggregate demand.

Bergemann et al. (2015) analyze the limits of price discrimination. They show that the segmentation and pricing induced by the additional information can achieve every combination of consumer and producer surplus such that: (i) consumer surplus is nonnegative, (ii) producer surplus is at least as high as profits under the uniform

¹⁰ A seller engages in *third-degree price discrimination* if she uses information about consumer characteristics to offer different prices to different market segments. If a seller has complete information about the buyer's willingness to pay, then she can engage in perfect or *first-degree* price discrimination. The seller can also offer a menu of choices, in terms of quality or quantity, to screen among different segments of the market, and this process is referred to as *second-degree* price discrimination.

monopoly price, and (iii) total surplus does not exceed the surplus generated by the efficient trade. The setting is as follows. A monopolist sells a good to a continuum of consumers, each of whom demands one unit. The total mass of consumers is one and the constant marginal cost of the good is normalized to zero. There are K possible values $V \triangleq \{v_1, \dots, v_k, \dots, v_K\}$, with $v_k \in \mathbb{R}_+$:

$$0 < v_1 < \dots < v_k < \dots < v_K.$$

A *market* x is a distribution over the K valuations, with the set of all markets being:

$$X \triangleq \Delta(V) = \left\{ x \in \mathbb{R}_+^K \mid \sum_{k=1}^K x(v_k) = 1 \right\}.$$


This set can be identified with the $(K - 1)$ -dimensional simplex, and to simplify notation we will write x_k for $x(v_k)$, which is the proportion of consumers who have valuation v_k . Thus, a market x corresponds to a step demand function, where $\sum_{j \geq k} x_j$ is the demand for the good at any price in the interval $(v_{k-1}, v_k]$ (with the convention that $v_0 = 0$).

A *segmentation* is a division of the aggregate market into different markets. Thus, a segmentation σ is a simple probability distribution on X , with the interpretation that $\sigma(x)$ is the proportion of the population in market x . A segmentation can be viewed as a two-stage lottery on outcomes in V whose reduced lottery is x^* :

$$\Sigma = \left\{ \sigma \in \Delta(X) \mid \sum_{x \in \text{supp } \sigma} \sigma(x) \cdot x = x^* \right\}.$$

Now, if the monopolist has no information beyond the prior distribution of valuations, there will be no segmentation. The producer charges the optimal monopoly price and gets the associated monopoly profit, and consumers receive a positive surplus; this is marked by point A in Fig. 1. If the monopolist has complete information, then he can charge each buyer his true valuation, i.e., engage in perfect or *first-degree price discrimination*; this is marked by point B. The point marked C is where consumer surplus is maximized; the outcome is efficient and the consumer gets all the surplus gains over the uniform monopoly profit. At the point marked D, social surplus is minimized by holding producer surplus down to uniform monopoly profits and holding consumer surplus down to zero.

A simple example will help to illustrate these results. There are three valuations, $v \in V = \{1, 2, 3\}$, which arise in equal proportions, and there is zero marginal cost of production. The feasible social surplus is $w^* = (1/3)(1 + 2 + 3) = 2$. The uniform monopoly price is $v^* = 2$. Under the uniform monopoly price, profit is $\pi^* = (2/3) \times 2 = 4/3$ and consumer surplus is $u^* = (1/3)(3 - 2) + (1/3)(2 - 2) = 1/3$. A segment x is a vector of probabilities of each valuation, thus $x = (x_1, x_2, x_3)$, and by $\sigma(x)$ we denote the total mass of a segment x . A *segmentation* of the market is therefore a collection of segments $x \in X$ and a probability distribution $\sigma(\cdot)$ over the

**FIGURE 1**

The triangle of consumer surplus and producer surplus attainable in third-degree price discrimination.

segments. We give an example of a segmentation below. In the example, there are three segments and each segment is identified by its support on the valuations indicated by the set $\{\cdot\}$ in the superscript. The frequency of each segment x is given by $\sigma(x)$:

Segment	x_1	x_2	x_3	$\sigma(x)$
$x^{\{1,2,3\}}$	$\frac{1}{2}$	$\frac{1}{6}$	$\frac{1}{3}$	$\frac{2}{3}$
$x^{\{2,3\}}$	0	$\frac{1}{3}$	$\frac{2}{3}$	$\frac{1}{6}$
$x^{\{1\}}$	0	1	0	$\frac{1}{6}$
x^*	$\frac{1}{3}$	$\frac{1}{3}$	$\frac{1}{3}$	1

(12)

The particular segmentation has a number of interesting properties. First, in each segment, the seller is *indifferent* between charging as the price *any* valuation that is in the support of the segment. Second, the uniform monopoly price, $p^* = 2$ is in the support of every segment. Thus this particular segmentation preserves the uniform

monopoly profit. If the monopolist charges the uniform monopoly price on each segment, we get point A. If he charges the lowest value in the support of each segment (which is also an optimal price, by construction), we get point C; and if he charges the highest value in the support, we get point D. The choice of the optimal monopoly price is of course a special mechanism design problem with a single agent. It is therefore an important next step to determine how the distribution of private information affects the distribution of surplus with many agents, thus bidding in an auction environment. Bergemann et al. (2017) characterize for a given symmetric and arbitrarily correlated prior distribution over values, the lowest winning-bid distribution that can arise across all information structures and equilibria. Bergemann et al. (2019) rank auctions according to their revenue guarantees, i.e., the greatest lower bound of revenue across all informational environments, where the distribution of bidders' values is fixed. If all equilibria and all information structures are considered, then first-price auctions have a greater revenue guarantee than all other auctions considered.

Roesler and Szentes (2017) analyze the impact of information in a variant of the above environment of optimal monopoly pricing. They propose a model where the buyer's valuation for the object is uncertain and she can commit to an optimal information structure that in turn affects the price-setting behavior by the seller. They show that the resulting outcome leads to efficient trade under unit-elastic demand.

Ichihashi (2020a) considers a richer model of consumer privacy and price discrimination featuring a multi-product seller. The seller has k products in his portfolio but is required to make a single product and price offer to each consumer. Although the consumer may benefit from accurate recommendations, the seller may use information about the willingness to pay to price discriminate. At the beginning of the game, the consumer chooses a disclosure rule, which determines the information that the seller learns about the consumer's values about the k products. He compares two settings that differ in the timing at which the seller sets prices. Under the no-commitment regime, the seller sets prices after learning about the consumer's values. Under the commitment regime, the seller sets prices up front without observing the consumer's information disclosure. The commitment regime captures the seller's commitment to not use consumer information for pricing. The main finding of Ichihashi (2020a) is that the seller is better off by committing to not use the consumer's information for pricing. By making such a commitment, the seller can induce the consumer to disclose more information, which enables the seller to make more accurate recommendations. This shifts up the consumer's demand for recommended products and increases revenue. His analysis provides a complementary argument to Bergemann et al. (2021a) regarding why personalized price discrimination seems to occur infrequently in e-commerce, see Iordanou et al. (2017).

The size of the possible gains, for both consumer and producer surplus, relative to the uniform pricing rule suggests that there is substantial scope for the provision of additional information. The large range of feasible pairs of consumer and producer surplus implies that there may be many possible business models for data intermediaries to cater in various degrees to producers or consumers. The potential for individualized, personalized pricing was recognized earlier by Shapiro and Varian

(1999) and is reviewed in a survey by Fudenberg and Villas-Boas (2012). A recent report by the Council of Economic Advisers (2015) offers largely negative conclusions regarding consumer welfare.

Dube and Misra (2017) consider the empirical implications of price discrimination using high-dimensional data from a large, digital firm. They run a large, randomized price experiment with a high-dimensional vector of customer features that is observed prior to price quotes. The outcomes of the price experiment are used to train the demand model. Then they conduct an optimal third-degree price discrimination exercise on the basis of the observable variables. Even the optimal uniform price substantially increases profits relative to the current price policy of the firm. They estimate that the third-degree price discrimination policy delivers further increases in the profits without affecting the consumer surplus by much. The social welfare increases as more than two-thirds of the consumers face lower prices than under the optimal uniform price.

Shiller (2020) analyzes personalized pricing for Netflix's movie rentals by mail with data from 2006. He distinguishes between demographic data and web-browsing data and finds small incremental gains from using price discrimination that relies on demographic data and substantially more significant gains based on web-browsing data.¹¹ With the abundance of personal and aggregate data available, it remains a puzzle why firms—in particular digital platforms—remain reluctant to implement personalized pricing at a large scale. Bergemann et al. (2021a) argue that the cost of using individual data might exceed the benefits of personalized pricing and note that group-, location-, or time-based price discrimination may approximate the gains from personalized pricing at a lower cost. Shiller (2021) provides evidence that firms frequently tailor a posted price to the arriving consumer but change prices at low enough frequency across consumers so that the resulting pricing strategy of almost personalized pricing is indistinguishable from traditional dynamic pricing.

Shiller (2020) establishes similar differences for the effectiveness of targeting advertisement based on these two different sources of data. Goldfarb and Tucker (2011) investigated the impact of a law limiting web tracking based on stated purchase intention among banner advertising viewers. They find a substantial reduction in advertising effectiveness in the absence of web-tracking data.

Dube et al. (2017) consider the value of one piece of information for targeting policies, namely the GPS data of a consumer as conveyed by her mobile phone. In a field experiment, they test mobile targeting based on consumers' real-time and historic locations, allowing them to evaluate popular mobile coupon strategies in a competitive market. They find substantial profit gains from price discrimination in a competitive environment.

¹¹ In related work, Reimers and Shiller (2019) examine the impact of telematics-based monitoring on price discrimination in automobile insurance. Here, the monitoring data are obtained by using telematics devices—which records information about driving behavior, including how fast you drive, how fast you brake, and the distance you drive—allows the insurer to more finely target premiums to monitored consumers.

4.2 Ratchet effect and privacy

The profitability of trading consumer information to facilitate price discrimination raises the issue of the endogenous availability of such information. In particular, information is rarely purchased directly from a consumer in exchange for a monetary payment, a practice far more common in business-to-business transactions. Instead, it is often the case that information must be sourced indirectly, by recording the consumer's actions, e.g., their purchase histories. The expected use of this information influences a consumer's willingness to reveal information through their behavior. In other words, ratcheting forces determine the level of the *indirect compensation* that the consumer requires for the information they generate.

It is well known that in dynamic principal-agent models, private information held by the agent requires the optimal contract to be a long-term arrangement, one that cannot be replicated by a sequence of short-term contracts. This is due the “ratchet effect,” see the seminal work of Freixas et al. (1985) and Hart and Tirole (1988). Feldman et al. (2013) offer an extension for digital markets with multiple goods offered via a sequence of first-price auctions when the bidders have additive or gross-substitute valuations.

A leading example of the ratchet effect emerges in the problem of selling a number of goods over time as the demand for the goods evolves. The dynamics that arise in such problems pertain both to the evolution of the willingness-to-pay as well as to the set of feasible allocations over time—through a variety of natural channels; see Bergemann and Said (2011) and Bergemann and Välimäki (2019) for surveys regarding these dynamic and intertemporal problems.

An interesting set of issues arises when the mechanism itself can govern only some of the relevant economic transaction. A specific setting where this occurs is optimal mechanism in the initial stage of the game is affected by the interaction in the resale market, see for example Calzolari and Pavan (2006), Dworzak (2016), Carroll and Segal (2019), and Bergemann et al. (2020). In particular, the information that is generated by the mechanism may affect the nature of the subsequent interaction, and thus tools from information design and mechanism design jointly yield interesting new insights.

In the context of price discrimination, such indirect compensation often takes the form of more favorable terms (e.g., a lower purchase price) for transactions that are likely to be recorded and subsequently used against a consumer. For example, a sophisticated consumer may become wary of purchasing unhealthy foods or tobacco products if that information impacts their health insurance premium.¹²

Taylor (2004) develops the first analysis of such a scenario in a two-period model of price discrimination, showing how tracking and selling a consumer's purchase

¹² Information about a consumer's preferences may also be used in their favor, e.g., through the customization of product characteristics. De Corniere and de Nijs (2016), Hidir and Vellodi (2018), and Ichihashi (2020b) analyze different aspects of the tradeoff between content personalization and price discrimination.

history introduces the need to compensate a sophisticated consumer for their first-period actions. Overall, the transmission of information may benefit a sophisticated consumer, while unambiguously hurting a naive consumer. However, even a sophisticated consumer is hurt by any adverse (e.g., discriminatory) use of information that is not collected in the context of a monetary transaction. For example, if a consumer's browsing (not purchasing) history affects future prices, the scope for compensating them for the data generated is greatly diminished. Conitzer et al. (2012) consider a two-period model in which a monopoly seller can use the information gained from the purchase history yet the buyer can pay to anonymize his purchase history against a fixed cost. The endogenous anonymization decision then leads to segmentation of the buyers according to their willingness-to-pay.

Hoffmann et al. (2020) suggest a subtle channel through which even light-touch privacy regulation requiring consumer consent may backfire and harm consumers. In the absence of privacy regulation, sellers have an incentive to acquire information about the individual preferences of consumers and then to selectively disclose the features of their products that consumers find most attractive. The incentive to selectively disclose, however, may be self-defeating for sellers when consumers rationally anticipate selective disclosure. Privacy regulation that requires sellers to obtain consent from consumers (before acquiring personal information that enables selective disclosure) may then enable firms to commit not to selectively disclose. Overall, Hoffmann et al. (2020) argue that privacy regulation increases consumer welfare if and only if (i) competition among sellers is limited or asymmetric, (ii) consumers are unaware of selective disclosure, and (iii) sellers are able to price discriminate.

It remains an open question how privacy policies should be set to find the optimal balance between privacy concerns (whether they are intrinsic or instrumental) and possible gains from improved targeting and matching in the product market. Recent empirical work on the effects of privacy regulation, such as the European Union's General Data Protection Regulation (e.g., Aridor et al. (2020) and Johnson et al. (2020)), indicates that data externalities and associated loss of data privacy are relevant for consumers' and businesses' decisions to share their data.¹³

Importantly, a compensatory channel in the trade of data is present even if the participating firms do not benefit, on aggregate, from participating in the market for information. Calzolari and Pavan (2006) establish this result in a two-period, two-firm model with general mechanisms. Conversely, exogenous (e.g., regulatory) limits to the available contractual instruments may reduce the firms' ability to extract surplus through price discrimination. In this case, the transmission of information can benefit firms and/or consumers.

Along these lines, Bonatti and Cisternas (2020) study how aggregating the information about purchase histories into a *consumer score* impacts the ratchet effect.

¹³ In the United States, legislators are also increasingly aware of the consequences of data externalities. In particular, the U.S. House Committee on the Judiciary (2020) reports that "[...] the social data gathered through [a platform's] services may exceed their economic value to consumers."

They do so in a continuous-time model with a changing consumer type and discriminatory, but linear prices. Thus, the information environment is high-dimensional, as signals arrive dynamically over time. A consumer score is modeled as a linear aggregate of past quantities with exponential decay. One specific instance of a score is given by the posterior mean belief about the consumer's type, given the equilibrium strategy and the entire history of past quantities.

A monopolist data intermediary constructs the consumer score and sells it to a sequence of short-run firms who use it to set prices. As information collection is free, the intermediary is always able to extract a positive price from the sellers. Bonatti and Cisternas (2020) further show that by increasing the persistence of the consumer's score relative to the Bayesian benchmark the intermediary is able to mitigate the ratchet effect. This allows her to collect more informative signals from the consumer, which are in turn more valuable for the sellers.

Finally, Ball (2020) introduces a high-dimensional model of scoring. A receiver seeks to predict a sender's quality. An intermediary observes multiple features of the sender and aggregates them into a score. Based on the score, the receiver makes a decision. The sender wants the most favorable decision, and can distort each feature at a privately known cost. Ball (2020) characterizes the most accurate scoring rule. This rule underweights some features to deter sender distortion, and overweights other features so that the score is correct on average. The receiver prefers this score to full disclosure because the aggregated information mitigates the commitment problem. Here, the richness of information is due to the fact that the agent has a multidimensional type vector, yet only one dimension of the type is relevant for the decision-maker.

4.3 Ratings and recommender systems

The sale of consumer scores for marketing purposes is but one instance of a market for aggregated information. For example, consider FICO credit scores for individual consumers and Moody's, Standard & Poor's, or Fitch credit ratings for corporate and sovereign debt. These ratings compress the high-dimensional information about an entire financial history into a single dimension that facilitates the coordination of actions, such as lending or investment.

More generally, all ratings and recommender systems are means to induce an appropriate course of action. As such, any rating raises the issue of incentive compatibility, as the use of past information determines the rated agent's incentives to undertake specific actions. For example, in the career concerns model of Hörner and Lambert (2021), a rating is used to aggregate a worker's past performance, and to convey a productivity estimate (and hence, the correct level of pay) to the market. At the same time, ratings are "motivational," since they affect the worker's incentives to boost current performance, and thus future wages.

Recommenders systems employ a variety of algorithms to predict users' preferences based on their consumption histories, their demographic profiles, and their search and click behaviors, see Avery et al. (1999), Schafer et al. (1999), and Bergemann and Ozmen (2006) for descriptions of collaborative filtering. Mansour et al.

(2015) consider a social planner, who by means of carefully designed information disclosure can incentivize the agents in a recommender-like setting to balance the exploration and exploitation so as to maximize social welfare. They formulate a multi-armed bandit problem under incentive-compatibility constraints induced by the agents' Bayesian priors. They establish that an incentive-compatible bandit algorithm exists for which the regret of the social planner is asymptotically optimal among all bandit algorithms (incentive-compatible or not). Miller et al. (2005) devise a scoring system that induces honest reporting of feedback. Each rater merely reports a signal, and the scoring system applies proper scoring rules to the implied posterior beliefs about another rater's report.

Incentive-compatibility constraints can also affect the very ability of the market to generate new information. Several online platforms (e.g., the traffic navigation software Waze or the reviews site Tripadvisor) incentivize social experimentation (e.g., trying a new route connecting two points or a new hotel), illustrating how the use of information influences a consumer's incentives to generate data in the first place. Related to this problem, Kremer et al. (2014) and Che and Hörner (2017) analyze the information design problem of a benevolent planner who wishes to induce a sequence of uninformed, short-lived agents to engage in socially useful (but privately costly) experimentation. In the example of navigation software, experimentation entails recommending to some users a route that has not yet been taken. In both these papers, commitment power is required to dynamically use past information in a way that makes it worthwhile for consumers to follow the platform's current recommendation.

Recommender systems, as well as analytics services that leverage Artificial Intelligence (AI), can also be seen as mechanisms for selling information in the form of predictions. This feature is somewhat related to the question of how to measure information, see Frankel and Kamenica (2019), and closely related to the optimal pricing of information. On this point, Agrawal et al. (2018) argue that firms that own considerable data on users' preferences online can use AI as means to sell information indirectly: instead of distributing unique datasets, providers such as Google, Facebook, Microsoft, and Amazon can bundle a prediction ("consumer i is high-value for firm j ") and a product (e.g., an advertising slot or product recommendation).

The distinction between selling information and selling access to a consumer has important implications for the price of information in a dynamic environment. With direct sales of information, buyers can either retain the data, and hence use stale, old predictions as an outside option, or hold and retain the original contact. In both cases, the value added of an information seller is to keep the buyer up to date. In particular, as long as the buyer retains the possibility of taking an informed action (e.g., contact a consumer), the data broker will only be able to charge for the *innovation* component of her data. If, on the contrary, an AI provider offers *exclusive access* to qualified prospects, it will be able to repeatedly charge for the full (flow) value of her information over time. The potential value of a market for insights—actionable recommendations that do not require distributing raw data—is also discussed in Dahleh (2018).

4.4 Certification and expert markets

Lizzeri (1999) analyzes the role of intermediaries who gather information of privately informed parties. The intermediaries then choose what to reveal to uninformed parties. Focusing on a monopoly intermediary, in a class of environments he shows that the optimal choice is to reveal only whether quality is above some minimal standard. While there is only minimal information transmission, the intermediary can capture a large share of the surplus. By contrast, competition among the intermediaries can lead to full information revelation. This environment is related to the analysis in Bergemann et al. (2021a), who focus on the sequential optimal solution for the intermediary, but in an extension discuss the nature of the commitment solution. Lizzeri (1999) has a number of distinct features. First, in Lizzeri (1999) the private information is held by a single agent, and multiple downstream firms compete for the information and for the object offered by the agent. Second, the privately informed agent enters the contract after she has observed her private information; thus, an interim perspective is adopted. The shared insight is that the intermediary *with* commitment power might be able to extract a rent without any further influence on the efficiency of the allocation. Dranove and Jin (2010) provide a comprehensive survey on the theoretical as well as the empirical literature on quality disclosure and certification, with a particular focus on healthcare, education, and finance.

In the absence of an intermediary, experts may still be limited in their ability to communicate their private information and their expertise. Morris (2001) considers a repeated cheap talk game. An informed advisor wishes to convey her valuable information to an uninformed decision maker with identical preferences. Thus she has a current incentive to truthfully reveal her information. But there is uncertainty about the type of the advisor. If the advisor is “good,” she has the same preferences as the decision maker. If the advisor is “bad,” she always wants as high an action as possible. The state is realized (and publicly observed) after the decision maker’s action is chosen. Thus, if the decision maker thinks that the advisor might be biased in favor of one decision, and the advisor does not wish to be thought of as biased, the advisor has a reputational incentive to lie. If the advisor is sufficiently concerned about her reputation, no information is conveyed in equilibrium. In a repeated version of this game, the advisor cares (instrumentally) about her reputation simply because she wants her valuable and unbiased advice to have an impact on future decisions.

In markets for financial products and health services, the advice of expert intermediaries is influenced by inducements paid by sellers. Inderst and Ottaviani (2012a) investigate how competition among service providers for the opinion of experts endogenously determines the objectives of the intermediaries. They compare the welfare properties of disclosed commissions with hidden kickbacks, which result in stronger incentive payments but may make the advisor more responsive to the cost of the different services on offer. In addition, they characterize conditions under which commonly adopted policies such as mandatory disclosure and caps on commissions have unintended welfare consequences. See Inderst and Ottaviani (2012b) for a broader overview of the role of advice in markets, both from a positive and a normative perspective.

Santosh et al. (2017) provide rich empirical evidence to understand the nature of advice by commission-motivated agents. They conduct a series of field experiments to evaluate the quality of advice provided by life insurance agents in India. They find that agents overwhelmingly recommend unsuitable, strictly dominated products that provide high commissions to the agent. In fact, they find that agents appear to focus on maximizing the amount of premiums (and therefore their commissions) that customers pay, as opposed to focusing on how much insurance coverage customers need. They further establish that the agents cater to the beliefs of uninformed consumers, even when the beliefs held by the consumers are wrong. They also find that requiring disclosure of commissions for a specific product, but not for others, results in agents recommending the products with high commissions but no disclosure requirement.

Robles-Garcia (2019) analyzes the role of brokers in mortgage markets through a structural model applied to loan-level data from the universe of UK mortgage originations. The model features households' demand for mortgage products and broker services, lenders' optimal pricing decisions, and broker-lender bilateral bargaining over commission rates. Mortgage brokers often receive commission payments from lenders. The view of the role of brokers that emerges is nuanced: (1) brokers increase upstream competition by facilitating sales of small, lower-cost lenders, and (2) commission rates distort brokers' advice and generate an agency problem with households. A counterfactual exercise evaluates the impact of policies restricting brokers' services and remuneration. A ban on commission payments would lead to a 25% decrease in consumer welfare, supporting the competition enhancing role of the brokers, whereas caps can increase consumer surplus by at least 10%, suggesting that some regulation of the commission payments may be welfare enhancing.

5 Forecasting and aggregation of predictions

We now turn to institutional arrangements for collecting information and aggregating expert judgments for the purpose of improving decision making. Expert forecasts play a crucial role in shaping decisions in areas ranging from meteorology to medicine, politics, and economics. Firms, for instance, decide whether they should hire workers based on their prospects for future demand and choose their financing strategies using projected long-term interest rates. But should we expect forecasters to truthfully report their expectations? What are the properties of the predictions that result from different mechanisms?

Section 5.1 introduces forecasting nonmarkets, a catchall for different mechanisms where experts interact in a non-anonymous way and receive rewards, possibly in terms of reputation and prizes. Section 5.2 turns to prediction markets where experts interact anonymously and prices are formed.

5.1 Forecasting nonmarkets

Section 5.1.1 begins by introducing incentive schemes, known as scoring rules, designed to give forecasters incentives to report their predictions truthfully. We then

overview the reporting incentives by forecasters who either compete in contests (Section 5.1.2) or are motivated by their reputation for expertise (Section 5.1.3). Our presentation is brief to avoid overlap with Marinovic et al.'s (2013) survey of the literature of forecasting nonmarkets.

5.1.1 Scoring rules

Consider an expert forecaster, such as a meteorologist, with private information about the state of the world. An evaluator may wish to measure the accuracy of his predictions, for example to compare different forecasting techniques or to reward exceptional forecasters. Meteorologists refer to this task as *forecast verification*. Let Ω be a finite set of mutually exclusive outcomes, for example rain and no-rain. Given a forecast $p \in \Delta\Omega$ and the actual realization $\omega^* \in \Omega$, a scoring rule $S(\Delta\Omega, \Omega) \rightarrow \mathbb{R} \cup \{-\infty, \infty\}$ assigns a numerical score, $S(p, \omega^*)$, based on the forecast and on the actual event that materialized. Brier (1950) observes that then-standard methods for assessing weather forecasts provided bad incentives as forecasters might be tempted to “hedge” or “play the system” and to forecast something other than what they think will occur. He then proposes a way to overcome the problem, thus anticipating the notion of incentive compatibility.¹⁴

Suppose that the forecaster's best judgment is $q \in \Delta\Omega$. Let

$$S(p, q) = \mathbb{E}_q [S(p, \cdot)] = \sum_{\omega \in \Omega} q(\omega) S(p, \omega)$$

be the expected value of $S(p, \cdot)$ under q , or the expected score. A scoring rule is said to be *proper* if the forecaster maximizes the expected score by being truthful: for all p and q , $S(q, q) \geq S(p, q)$. A scoring rule is *strictly proper* if q is the unique forecast that maximizes the expected value of S . Classic examples of strictly proper scoring rules include Brier's (1950) quadratic scoring rule

$$S(p, \omega^*) = 2p(\omega^*) - \sum_{\omega \in \Omega} p(\omega)^2$$

and Good's (1952) logarithmic scoring rule,

$$S(p, y) = \ln p(y).$$

The quadratic and logarithmic scoring rules are part of a larger family of proper probability scoring rules, axiomatized by McCarthy (1956), De Finetti (1962), and Savage (1971).¹⁵

¹⁴ We find it curious that Leonid Hurwicz from 1942 to 1944 was a member of the faculty of the Institute of Meteorology at the University of Chicago.

¹⁵ For developments see Gneiting and Raftery (2007), Osband (1989), and Schlag et al. (2015). Lambert (2011) characterizes the set of properties of the state distribution for which a proper scoring rule can be devised.

5.1.2 Forecasting contests

In a *Nature* article titled *Vox Populi*, octogenarian polymath Francis Galton (1907b) reports his seminal investigation “into the trustworthiness and peculiarities of popular judgments” as follows¹⁶:

“A weight-judging competition was carried on at the annual show of the West of England Fat Stock and Poultry Exhibition recently held at Plymouth. A fat ox having been selected, competitors bought stamped and numbered cards, for 6d. each, on which to inscribe their respective names, addresses, and estimates of what the ox would weigh after it had been slaughtered and ‘dressed.’ Those who guessed most successfully received prizes. About 800 tickets were issued, which were kindly lent me for examination ... Now the middlemost estimate is 1207 lb., and the weight of the dressed ox proved to be 1198 lb.; so the vox populi was in this case 9 lb., or 0.8 per cent of the whole weight too high.”

Galton’s competition is an example of a winner-take-all forecasting contest, modeled by Ottaviani and Sørensen (2006c) as a game of incomplete information among forecasters who each observe a noisy signal. More generally, the model captures forecasters’ concern for relative accuracy. In a setting with normally distributed signals and a common prior, forecasters have an incentive to report predictions that are more extreme than their “honest” posterior expectations when they have a convex preference for ranking. Intuitively, by moving away from the common prior and exaggerating the weight on their private signals, forecasters decrease the chance of having to share the prize. The level of exaggeration increases in the number of forecasters, as shown by Lichtendahl et al. (2013) and Banerjee (2021).

5.1.3 Reputational forecasting

Next, consider forecasters who are concerned about their reputation for being well-informed, rather than about the quality of the decisions made on the basis of their recommendations. For example, managers and business consultants with better reputation can charge more for their services. Forecasters who are believed to have access to more accurate information are more influential. Politicians deemed to be better informed are more likely to be re-elected and, thus, to obtain increased private benefits from holding office.

One may think that forecasters would maximize the accuracy of their reports by reporting truthfully their best predictions.¹⁷ However, rather than reporting truthfully, forecasters concerned about their reputation may be tempted to manipulate their predictions strategically, as noticed by Bayarri and DeGroot (1988, 1989) in the Bayesian statistics literature.

¹⁶ Re-examining Galton (1907a), Wallis (2014) makes some small corrections.

¹⁷ As often argued informally, analysts’ projections should be truthful because their livelihood depends on their accuracy; for example, see Keane and Runkle (1998).

To analyze the resulting equilibrium outcome, Ottaviani and Sørensen (2001) formulate a model of reputational cheap talk that works as follows. First, an expert forecaster privately observes a signal about a state of the world and reports a forecast to an evaluator. The forecaster can be more or less informed. In the baseline model, however, the forecaster's informativeness is initially unknown to both the forecaster and the evaluator.¹⁸ The evaluator assesses the informativeness of the forecaster on the basis of the forecast and the realized state of the world.¹⁹ The objective of the forecaster is to maximize the reputation for being well-informed, according to the assessment made by the evaluator.

As shown by Ottaviani and Sørensen (2006a,b,c), a reputation-driven forecaster who is believed to report truthfully generically has an incentive to misreport. Under natural conditions on the information structure, the forecaster benefits from strategically shading the prediction toward the prior consensus to avoid unfavorable publicity when wrong. In equilibrium, the forecaster is able to credibly communicate only part of their information. For an update on this literature and for experimental evidence we refer to Meloso et al. (2018) and references therein. In an interesting recent development, Deb et al. (2018) propose a mechanism design approach to evaluating forecasters.

5.1.4 Expert aggregation

Given the many practical applications of forecasting to decision making across different domains, forecast aggregation has attracted considerable attention from scholars at the interplay between econometrics, operations management, and decision sciences. The empirical starting point is the observation that forecast combinations tend to be more accurate on average than any single forecast, even by the most accurate forecaster. A classic approach is to compute the average forecast, the so-called *consensus forecast*. To improve on the arithmetic average, Bates and Granger (1969) combine forecasts by assigning weights to individual forecasters inversely proportional to their past errors. Pushing these ideas further, techniques have been developed to derive the optimal weights to be assigned to different forecasters by exploiting the covariance structure in past data.²⁰

To aggregate the predictions of different experts and obtain a more accurate crowd estimate, as a preliminary step it is key to properly adjust for the common information possessed by the experts. To this end, Prelec (2004) suggests asking experts to report not only their judgments, but also their guesses of how other experts will respond.

¹⁸ An incentive to exaggerate emerges if the forecasters privately know the informativeness of the signal; see Levy (2004).

¹⁹ The forecast is a cheap talk message, as in Crawford and Sobel (1982). Ex post observation of the state allows for some communication to be transmitted in equilibrium. Gentzkow and Shapiro (2006) model communication by media as a reputational cheap talk game in which the state of the world is observed with positive probability. See also Prat (2005) for the impact of transparency on the equilibrium outcome.

²⁰ See Winkler (1981) for an early contribution, Armstrong (2001) and Timmermann (2006) for reviews, and McGraw and Tetlock (2005) for an overview from a behavioral psychology angle.

Based on this idea, Prelec et al. (2017) develop a nonparametric algorithm that selects the “surprisingly popular” answer, i.e., the answer that maximizes the difference between the fraction of experts that indicate that answer and the experts’ expectation about the popularity of that answer. Along these lines, Palley and Soll (2019) propose an elicitation method for assessing the amount of information shared by forecasters by asking them to provide the average forecast they expect to be provided by the other forecasters.²¹

By and large, however, the literature on forecast aggregation—traditionally grounded in operations research—has been abstracting away from strategic behavior of forecasters, with the notable exception of the early contribution by Miller et al. (2005). More recent work incorporates incentive compatibility, see Cvitanic et al. (2019), Baillon and Xu (2021), Dai et al. (2021), and Chen et al. (2021). We expect future research to finesse aggregation techniques to account more explicitly for incentives and behavioral biases of forecasters in realistic settings. This is a fruitful and exciting research area with a wide range of applications that we expect to blossom on solid foundations. For practical adoption, however, it is important to develop elicitation methods that are simple and not cognitive demanding.

5.2 Prediction markets

Prediction markets are trading mechanisms that produce forecasts by aggregating expectations of traders.²² The assets traded in prediction markets are contingent-claim contracts that yield payments based on the outcome of future events. Leveraging the information aggregation properties of markets,²³ prediction market prices can be used to predict the underlying future events on which payments are made. Homer already recognized that wagering gives incentives to truthfully report beliefs. In Book XXIII of the *Iliad*, two illustrious spectators of a chariot race enter into an acrimonious dispute about which horses are in the lead. Idomeneus invites Ajax Oileus to wager a tripod or a cauldron, so that “you shall know when you pay.”²⁴

In the Iowa Electronic Markets (IEM), the prototypical example of prediction markets introduced to the research community by Forsythe et al. (1992), traders can take positions on whether an event E will be realized or not.²⁵ For instance, event E

²¹ Gaba et al. (2019) advance the methodology for assessing uncertainty from point forecasts.

²² See Wolfers and Zitzewitz (2004) for an introduction to the economics of prediction markets. Arrow et al. (2008) advocate that forecasts from prediction markets can have social value.

²³ Hayek (1945, pg. 526) envisions the ability of markets to aggregate information this way: “The mere fact that there is one price for any commodity ... brings about the solution which (it is just conceptually possible) might have been arrived at by one single mind possessing all the information which is in fact dispersed among all the people involved in the process.”

²⁴ For additional historical and institutional discussion see Orr (2014) and Pradier (2019) and references therein.

²⁵ Wagering on outcomes of US presidential elections has a long history; see Rhode and Strumpf (2004). According to *Fortune* magazine, the amount wagered in the 2020 presidential race surpasses the record figure of \$1 billion, resulting in predictions from betting markets closer to the mark than polls (Rey

might correspond to a Democrat winning the presidential race. For this binary market, there are two Arrow-Debreu contingent-claim assets associated with the two possible realizations. Asset E pays out 1 currency unit if event E is realized and 0 otherwise, while asset E^c pays out 1 if the complementary event E^c is realized and 0 otherwise. The occurrence of the event is verified in an ex post liquidation stage after the market is closed.

Consider a binary prediction market for two securities $\theta \in \{E, E^c\}$ with a continuum of traders indexed by $i \in [0, 1]$.²⁶ Traders have state-independent preferences over wealth in the two states of the world, E and E^c . Assume that these preferences satisfy Savage's subjective expected utility (SEU) axioms, so that trader i with subjective belief π_i about state E maximizes subjective utility $\pi_i u_i(w(E)) + (1 - \pi_i) u_i(w(E^c))$. Suppose that beliefs π_i are distributed in the population according to a continuous distribution function G over $[0, 1]$.²⁷

5.2.1 Equilibrium with heterogeneous beliefs and limited wealth

In a pioneering analysis of betting markets with heterogeneous beliefs, Ali (1977) characterizes the equilibrium in a binary contingent-claim market populated by risk-neutral traders, $u_i(w) = w$, with heterogeneous beliefs and bounded wealth equal to \bar{w} .²⁸ Ali's (1977) original model considers a pari-mutuel market in which the total amount of money wagered by traders on all outcomes is shared among the traders who wager on the outcome that is realized. However, when all money is returned to the bettors (the track take is zero), the equilibrium conditions that result in pari-mutuel markets are the same as in contingent-claim prediction markets such as the IEM.

When entering the market, traders exchange their wealth for an equal endowment in the two assets equal to \bar{w} . Replicating Ali's (1977) original formulation, traders have heterogeneous beliefs and make no inference from the equilibrium price. It follows that each bettor wagers all the available wealth on either of the two events. The competitive equilibrium is characterized by the indifference threshold belief $\bar{\pi}$ at which the expected payoff from placing the budget on either outcome is equalized. Traders with a belief above (below) threshold $\bar{\pi}$ place their entire budget on event E (E^c).

Ali's (1977) key observation is that the equilibrium price of the asset is equal to the belief of the median trader only when the median belief is exactly fair $\pi_m =$

Mashayekhi, "Betting markets called the presidential election more accurately than polls," November 19, 2020).

²⁶ While prediction markets typically trade binary contracts, the outcome does not have to be binary. There can easily be N binary contracts for each of $N \geq 2$ outcomes.

²⁷ The analysis reported below considers special versions of Ottaviani and Sørensen's (2015) model.

²⁸ For example, in the Iowa Electronic Markets traders are allowed to wager up to \$500. Equivalently, traders might allocate a limited wealth to the prediction market as a result of a two-stage budgeting process. The model can be easily extended to allow traders to have a heterogeneous wealth bounding the amount they can wager, with distribution H . In the more general model with risk aversion reported below the positions taken by traders are endogenously bounded, so that we can allow \bar{w} to be unbounded.

1/2. But if the median belief is favorable, $\pi_m > 1/2$, the equilibrium price p is still favorable but lower than that of the belief,

$$p \in (1/2, \pi_m) \text{ where } \pi_m = G^{-1}(1/2) > 1/2. \quad (13)$$

The asset price that equalizes supply and demand is equal to the fraction of the overall budget that is placed on the corresponding outcome E . To see this, note that a risk-neutral trader who holds a belief higher than the market price optimally wagers the entire budget on that outcome. Thus, at price p the demand for asset E is equal to $[1 - G(p)]\bar{w}$ divided by the price p , while its supply is equal to \bar{w} . So, it must be that the fraction of traders who have subjective beliefs above p is equal to the market probability:

$$1 - G(p) = p. \quad (14)$$

This market-clearing condition implies that the equilibrium price is unique and generically equals the $(1 - p)$ -quantile of the distribution of beliefs. By definition of the median π_m , half of the bettors have beliefs above it:

$$1 - G(\pi_m) = 1/2. \quad (15)$$

Hence, (13) follows from (14) and (15), given that the belief distribution G is increasing.

Result (13) follows from competitive equilibrium behavior regardless of the pari-mutuel market structure considered by Ali (1977), provided that traders are allowed to wager a limited amount of money. Intuitively, when more traders have beliefs in favor of an outcome, the corresponding asset price increases. Hence, traders who desire to hold that asset can acquire a lower amount of that asset, given the fixed wealth that traders are allowed to invest. For the market to equilibrate, the favorite must attract some traders with beliefs below the median. For this to happen, the price of the favorite (respectively longshot) outcome with price $p > 1/2$ (respectively $p < 1/2$) must be below (respectively above) the belief of the median trader.

Ali's (1977) result is compatible with the favorite-longshot bias, a widely-documented empirical regularity according to which favorite (respectively longshot) outcomes turn out to be empirically even more (respectively less) frequent than indicated by the market price.²⁹ The key identifying assumption for result (13) to explain the favorite-longshot bias is that the belief of the median bettor is equal to the empirical frequency of the outcome. However, this assumption is problematic because if the belief of the median trader were to contain information about the empirical

²⁹ See Ottaviani and Sørensen (2008) for an overview of the favorite-longshot bias and a discussion of the main alternative explanations that have been proposed in the literature. Additional explanations are given by Kajii and Watanabe (2017) (longshots are overvalued in the long run because they give a better chance of getting ahead with rare but large gains, thus allowing survival for longer than favorites) and Green et al. (2020) (racetracks have incentives to promote predictions that overestimate the chances of longshots so that naive bettors are induced to overbet longshots).

frequency of the outcome, traders should take this information into account when taking their positions—something Ali’s (1977) model does not contemplate. Like a lot of work before the information economics revolution in finance, associated with Grossman (1976), and the notion of rational expectations equilibrium (REE), Ali’s (1977) approach mixes beliefs and information in an unsatisfactory way.

Manski bounds

Manski (2006) proposes a resolution of this tension based on the following important observation. Without making Ali’s (1977) strong—and problematic—identifying assumption, from the market-clearing equation we can only conclude that when the market price is p , the mean belief in the population of traders satisfies

$$E[\pi] \in (p^2, 2p - p^2). \quad (16)$$

To see this, note that the most favorable distribution of beliefs consistent with price p assigns mass $1 - p$ to belief $\pi = 1$ and mass p to belief $\pi = p$, resulting in average belief $E[\pi] = p^2 + (1 - p)p = 2p - p^2$; similarly, the most unfavorable distribution of beliefs assigns mass $1 - p$ to belief $\pi = 0$ and mass p to belief $\pi = p$, resulting in $E[\pi] = p^2$. Manski’s (2006) upper and lower bounds cannot be achieved, but any point in the interior of the interval (16) is clearly achievable.

Overall, this observation cautions against over-interpreting prices as belief aggregators. At the same time, the tension is not fully removed because the market price contains some information about the average belief—and if the average belief is somewhat relevant to estimate the probability of the underlying event, traders should take this information into account when taking their positions—contrary to what the model assumes. To this end, however, it is necessary to explicitly model information, and the inference traders should draw from equilibrium prices, as done in the literature on rational expectations discussed in Section 5.2.3.

5.2.2 Equilibrium with heterogeneous beliefs and risk-averse traders

To reach a deeper understanding of the conditions under which market prices underreact to information, we now turn to a more general specification of the model with risk-averse traders; i.e., with $u'' < 0$. The assumption of risk aversion is natural in financial markets. To streamline the notation, we lift the exogenous bound on the amount traders can wager.³⁰ As in the baseline specification considered in the previous section, beliefs are heterogeneous and traders make no inference of information from the market price. The presentation follows Gjerstad and Hall (2005) and Wolfers

³⁰ As shown by Ottaviani and Sørensen (2015), the results reported below hold more generally also when traders have bounded wealth, for example, because of financial frictions.

and Zitzewitz (2006), but also relates to Wilson's (1968) result on the aggregation of prior beliefs with CARA preferences.³¹

Consider a trader i with initial wealth level w_0 facing price p for state E and price $1 - p$ for state E^c , where w_0 follows distribution H .³² The trader sells px_i units of the E^c asset at price $1 - p$ and invests the proceeds $(1 - p)px_i$ to purchase $(1 - p)x_i$ units of the E asset, obtaining final wealth $w_0 + (1 - p)x_i$ in state E and $w_0 - px_i$ in state E^c . The trader problem is

$$\max_{x_i} \pi_i u(w_0 + (1 - p)x_i) + (1 - \pi_i) u(w_0 - px_i)$$

with first-order condition

$$(1 - p)\pi_i u'(w_0 + (1 - p)x_i) = (1 - \pi_i) pu'(w_0 - px_i). \quad (17)$$

Logarithmic preferences

Assuming logarithmic utility, $u(w) = \ln(w)$, the demand function that solves (17) is $x_i = (\pi_i - p)w_0/[p(1 - p)]$. Market clearing requires that the excess demand across traders is zero

$$\int_0^\infty \left[\int_0^1 \frac{\pi_i - p}{p(1 - p)} g(\pi) d\pi \right] w dH(w) = \frac{E[\pi] - p}{p(1 - p)} E[w] = 0$$

so that the unique equilibrium price is equal to the mean of the beliefs of the traders

$$p = \int_0^1 \pi g(\pi) d\pi = E[\pi].$$

Gjerstad and Hall (2005) and Wolfers and Zitzewitz (2006) also derive the demand when traders have constant relative risk aversion, $u(w) = w^\rho/(1 - \rho)$, for $\rho \geq 0$; logarithmic utility obtains for $\rho \rightarrow 1$.

CARA preferences

With exponential utility $u_i(w) = -\exp(-w)$, demand is

$$x_i = \log\left(\frac{1 - p}{p}\right) + \log\left(\frac{\pi_i}{1 - \pi_i}\right),$$

which does not depend on initial wealth level w_0 . From the market-clearing condition we obtain that the log-likelihood ratio of the equilibrium price is the average of the

³¹ Luenberger (1995) reports similar results; see Exercise 11.11 for the case with logarithmic preferences and Exercise 11.5 for CARA. Rubinstein (1974) extends aggregation à la Wilson (1968) results to allow for heterogeneity in endowments.

³² For simplicity we consider the case where beliefs and wealth are independently distributed.

traders' log-likelihood ratios

$$\log \frac{p}{1-p} = \int_0^1 \log \left(\frac{\pi}{1-\pi} \right) dG(\pi). \quad (18)$$

5.2.3 Aggregation of information and beliefs

Building on this framework, suppose now that the belief π_i of trader i results from the combination of the prior belief q_i , possibly heterogeneous across traders with distribution G , and the observation of information summarized by the likelihood ratio $L = \ln \frac{f(s|\theta=E)}{f(s|\theta=E^c)}$ contained in the signal s . Initially suppose that all traders publicly observe the same signal s , following Ottaviani and Sørensen (2015). We further assume that beliefs are *concordant* in the sense of Milgrom and Stokey (1982): traders interpret information in the same way. After demonstrating that the equilibrium price under-reacts to public information, we show that the result carries over to the case in which the information is privately held and dispersed among traders—because all relevant private information is revealed by the equilibrium price that results under rational expectations.

Bounded wealth

In the model with risk aversion and bounded wealth, equilibrium condition (14) is replaced by

$$p = 1 - G \left(\frac{p}{(1-p)L + p} \right). \quad (19)$$

The competitive equilibrium price $p(L)$ is the unique solution to this equation—it is a strictly increasing function of the information realization L . Inverting Bayes' rule, the market price can be interpreted as the posterior belief of a hypothetical individual with initial belief

$$q(L) := \frac{p(L)}{[1-p(L)]L + p(L)}.$$

This initial belief $q(L)$ could be interpreted as an aggregate of the heterogeneous subjective prior beliefs of the individual traders. According to (19), this individual is the marginal trader. However, this way of aggregating subjective priors cannot be separated from the realization of information. Given that $q(L)$ is strictly decreasing in L , this initial belief of the marginal trader moves systematically against the public information available to traders. In this sense, the market price under-reacts to information.

Decreasing absolute risk aversion

More generally, Ottaviani and Sørensen (2015) show that prices under-react to information in the model with risk aversion and heterogeneous prior beliefs under the realistic assumption that traders have decreasing absolute risk aversion (DARA). With constant absolute risk aversion (CARA) preferences, the heterogeneous prior

beliefs are aggregated into a market prior belief

$$\log \frac{q}{1-q} = \int_0^1 \log \frac{q_i}{1-q_i} dG(q_i). \quad (20)$$

Without wealth effects, the market prior is equal to a weighted geometric average of the traders' prior beliefs, according to Wilson's (1968) classic result. In this boundary CARA case, the equilibrium price reacts like a Bayesian posterior belief computed with a market prior equal q . Under DARA preferences, under-reaction results because an increase in price induced by the realization of favorable information generates a negative wealth effect that induces traders to buy less. Thus, revelation of favorable information reduces (or increases) the weight assigned to more optimistic (or pessimistic) traders, resulting in under-reaction. As prior beliefs become more heterogeneous through a median preserving spread, the price under-reacts more to information.

Revelation of private information under rational expectations

While in the presentation so far the information L is publicly revealed, next we consider the case in which this information is initially dispersed among market participants, as is natural in prediction markets. When information is private rather than public, the same outcome results under the unique fully revealing rational expectations equilibrium (REE) à la Grossman (1976) and Radner (1979), as we show below.³³ For this reinterpretation, suppose that $s = \{s_i\}_{i \in I}$, where signal s_i is privately observed by trader i before trading. The model (i.e., preferences, prior beliefs, and signal distributions) and the rationality of all traders are common knowledge.

Before presenting the result, note that the informational requirements for REE equilibrium are more demanding than those for the notion of competitive equilibrium. As also argued by Morris (1995), the REE concept in general relies on strong assumptions in terms of common knowledge. The learning that is necessary for strategic equilibrium play to become sensible would also eliminate heterogeneity in prior beliefs (see Dekel et al. (2004)). In spite of this caveat, REE equilibrium is a useful benchmark to capture the fact that traders make an inference about the state from the realized price when information is asymmetric.

A related issue is whether the equilibrium is implementable. As discussed by Palfrey and Srivastava (1986), when agents are informationally small, they have no incentive to falsely report their information. We provide here an adaptation of their result to our setting. Suppose the trader population is divided into N classes of traders, $n = 1, \dots, N$. Each class n consists of at least two traders denoted by $i_n \in I_n$. Traders in class n have an aggregate endowment of each asset equal to $\tau_n > 0$, such that $\sum_{n=1}^N \tau_n = 1$. Within class n , $G_n(q) \in [0, 1]$ denotes the share of its assets initially held by individuals with subjective prior belief less than or equal to q .

³³ We focus on the necessary properties of this equilibrium, while Radner (1979) discusses sufficient conditions for its existence.

Traders in class n all observe the common signal s_n . This observation is private to the class. Let $f_n(s_n|\omega)$ denote the conditional probability density, with a likelihood ratio $L_n(s_n)$. We assume that signals are independently distributed across classes. Then, $s = (s_1, \dots, s_N)$ has conditional density $f(s|\theta) = \prod_{n=1}^N f_n(s_n|\theta)$.

Consider the following mechanism. Each trader is asked to report the signal. If all traders in each group send identical reports, then the mechanism designer computes the competitive equilibrium $p(L)$ that would result if s were publicly observed, with $L = f(s|E)/f(s|E^c)$. Traders are then allowed to exchange assets at price $p(L)$, subject to their possible wealth bound. If some trader-pair from some group sends different reports, all traders stay with their initial asset allocation. This mechanism admits a Bayesian Nash equilibrium in which every player reports truthfully.³⁴

The proof works as follows. Suppose that all traders but i_n send truthful reports. Suppose that i_n has observed signal s_n . Reporting s_n results in the competitive equilibrium allocation given s . Deviating results in the initial endowment. In competitive equilibrium, trader i_n can always choose to keep the initial endowment of the assets, so the competitive allocation is weakly preferred by revealed preference. Hence, truthtelling is optimal and the mechanism implements the competitive equilibrium as if s were publicly observable.

Because the assumption that the market reaches a fully revealing equilibrium is not warranted for some market rules, this is the most optimistic scenario for information aggregation.³⁵ Relaxing the REE assumption, Ottaviani and Sørensen (2009, 2010) analyze a game-theoretic model of pari-mutuel betting where traders have a common prior but are unable to condition their behavior on the information that is contained in the equilibrium price.³⁶ They find that the price under-reacts or over-reacts to information depending on the amount of information relative to noise that is present in the market. In contrast, in the model presented here the information possessed by the traders is revealed in equilibrium, but the price under-reacts to information nevertheless when the priors are heterogeneous and wealth effects are present.

Link to no-trade theorem

It is worth relating this analysis to the no-trade theorem, according to which rational traders cannot take positions against one another on the basis of differences of information alone.³⁷ In the static model presented here, trade takes place only if traders have heterogeneous beliefs; different private information alone would not be enough

³⁴ This result and the surrounding discussion is based on an unpublished appendix of the working paper version of Ottaviani and Sørensen (2015). We thank Peter Norman Sørensen for collaboration in proving the result.

³⁵ Plott and Sunder (1982) and Forsythe and Lundholm (1990) experimentally investigate the conditions leading to equilibrium in settings with differential private information.

³⁶ See also Shin (1991, 1992) for a derivation of the favorite-longshot bias when prices are quoted by an uninformed monopolist bookmaker.

³⁷ See Rubinstein (1975), Kreps (1977), Tirole (1982), Milgrom and Stokey (1982), and Morris (1994).

for trade. Embedding this model in a multi-period framework with additional concordant information arriving in each period, no more trade occurs from the second period, consistent with Milgrom and Stokey's (1982) general framework. As shown by Ottaviani and Sørensen (2015), price changes are positively correlated, consistent with the empirical pattern of momentum and post-earning announcement drift.

Two key deviations from Grossman's (1976) classic model of REE are essential to obtain under-reaction: heterogeneous priors and wealth effects. The finance literature before REE mostly focused on heterogeneous priors models (see Miller (1977) and Mayshar (1983)). Since Grossman (1976) and Grossman and Stiglitz (1980), for roughly twenty years REE models dominated the scene in the finance literature, in various guises and forms, but almost always with common prior and CARA preferences.³⁸ However, a bifurcation occurred in the finance literature in the last twenty years. Because the volume of trade is hard to reconcile with predictions of models with a common prior, a large part of the more applied literature returned to heterogeneous beliefs, but abstracted away from information (see Hong and Stein (2007)). By combining heterogeneous beliefs with information, the synthesis proposed by the model reported here suggests a way forward that retains the basic lesson of the information economics revolution, without throwing away the baby (information) with the bath water (REE).

A more realistic model would also relax the assumption of concordant beliefs, consistent with mounting evidence, at least since Kandel and Zilberfarb (1999), that information is processed and interpreted differently across agents. A key difficulty is the limited analytical tractability of models with heterogeneous interpretation of information where Milgrom and Stokey's (1982) characterization does not apply,³⁹ especially when forward-looking traders operate in a dynamic environment.⁴⁰

5.3 Automated market maker

A common implementation of prediction markets relies on determining the price through an automated market maker. An attractive feature of this design is that trading participants can instantaneously take positions against the automated market maker at any time the market is open, without having to wait for another trader who is willing to take the other side of the trade.

The automated market maker operates according to a cost function $C : \mathcal{R}^{\Theta} \mapsto \mathcal{R}$ that maps a vector of bets \mathbf{b} to a scalar representing how much money has been paid into the system. Each entry b_{θ} in the vector \mathbf{b} represents how much money is to be

³⁸ In a notable exception, Varian (1989) relaxes common prior, but sticks to CARA preferences to retain tractability in combination with normality assumptions for the prior and the signals. Without wealth effects, the price still reacts one for one to information. The model with binary state presented here can handle general risk preferences—and thus relax CARA—in a tractable way.

³⁹ See Harris and Raviv (1993) for a model of trading with differential interpretation of information across market participants.

⁴⁰ See, for example, Banerjee and Kremer (2010).

paid out if event θ is realized. For instance, $\Theta = \{E, E^c\}$ when the market is over a binary event. If the market maker has taken bets that sum to paying out b_E if E is realized and b_{E^c} if E^c is realized, then $\mathbf{b} = (b_E, b_{E^c})$. The market maker charges $C(b_E + 1, b_{E^c}) - C(b_E, b_{E^c})$ to a trader that makes an additional bet, in exchange for committing to pay one dollar under E to the trader.

Hanson (2003, 2007) proposed a commonly adopted design for the automated market maker based on a logarithmic market scoring rule (LMSR), with cost function

$$C(\mathbf{b}) = \beta \log \left[\sum_{\theta \in \Theta} \exp(b_\theta / \beta) \right],$$

where $\beta > 0$ is a market liquidity parameter set by the market administrator. When the level of β is higher, the market is less affected by small bets. The LMSR has a worst-case loss of $\beta \log n$, which corresponds to the highest amount the participating traders can win from the market maker. Prices are defined by the gradient of the cost function, so that the price of event θ with LMSR is

$$p_\theta(\mathbf{b}) = \frac{\exp(b_\theta / \beta)}{\sum_{\theta \in \Theta} \exp(b_\theta / \beta)}.$$

These prices define a probability distribution over the event space (they are non-negative, sum to one, and exist for any set of events) and thus can be thought of as event probabilities. For the binary market example sketched above, starting from the vector of bets (b_E, b_{E^c}) , when a trader places an additional bet on asset $\theta = E$ the price increases from $p_\theta(b_E, b_{E^c})$ to $p_\theta(b_E + 1, b_{E^c})$. Participating traders who maximize the expected value of their wins without capital constraints should interact with the LMSR until the market maker's prices reflect their beliefs.

When participants have a common prior and are allowed to trade sequentially for infinite turns, and prices are publicly announced after each period, Ostrovsky (2012) shows that the market price of each asset converges in probability to its expected value, conditional on the traders' pooled information. Under common prior, we know from Aumann (1976) that posterior beliefs must be identical if they are public information, and from Geanakoplos and Polemarchakis (1982) that repeated announcements of beliefs generically lead to belief convergence. Sethi and Vaughan (2016) investigate the performance of the automated market maker design when traders have heterogeneous priors.

5.4 Performance of prediction markets

Market manipulation

Prediction and betting markets are not immune to manipulation. Ernest Hemingway, in Chapter 20 of *Farewell to Arms*, gives a literary account of market manipulation at Milan's San Siro racetrack. Having heard rumors that Italian horse racing is crooked, the American protagonists notice a purplish-black horse they believe has been dyed

to cover its true color. Hoping the horse is a champion in disguise, the protagonists bet on this horse that displays very long provisional odds of thirty-five to one. The horse goes on to win soundly, but just before the race begins the pari-mutuel odds drastically adjust, because insiders heavily back the horse. The protagonists end up collecting much less than initially expected.

In the baseline model presented above, the state of the world is exogenous and cannot be affected by the traders. This assumption is realistic for prediction markets on economic statistics, such as non-farm payroll employment. In the context of corporate decision making, however, traders might be able to manipulate the outcome. The positions taken in the prediction market may actually enhance the manipulation incentives. Building on this model, Ottaviani and Sørensen (2007) analyze the incentives of traders to *manipulate the outcome*. Focusing on the case with CARA preferences, they show that more optimistic traders have an incentive to bet on success while pessimists bet against success. Because of the positions they take in the market, optimists have an incentive to improve the chances of project success, while pessimists manipulate the chances in the opposite direction. The manipulation game among traders leads to wasteful activities, in which participants partly offset each others' manipulations. Overall, manipulations of optimists and pessimists typically do not perfectly offset each other. Introducing the prediction market helps the market designer become more informed about the chance of project success, but it also directly influences this chance.⁴¹

Traders in prediction markets also have incentives to spread rumors or false information for the purpose of manipulating the market price and then take advantage of the induced mispricing.⁴² *Information-based manipulation* is a public good among manipulators, who would individually prefer to free-ride on the attempts by others to manipulate the market. Free riding increases as the market becomes more liquid. However, this kind of manipulation is a serious concern for illiquid prediction markets. In a field experiment to verify the possibility of information-based manipulation in pari-mutuel betting markets, Camerer (1998) manages to move odds visibly by taking, and then strategically canceling, positions. However, in spite of the additional bets attracted toward the horse that was temporarily bet, the final net effect was close to zero. In Hanson et al.'s (2006) lab experiment, attempts to manipulate the market also failed and ended up increasing liquidity, thus improving information aggregation.⁴³

Finally, consider *trade-based manipulation* analyzed by Goldstein and Guembel (2008) in the context of regular asset markets where equilibrium prices incorporate the value of information created by the improved decision. In that context, traders

⁴¹ Wolfers (2006) analyzes empirically the interplay between point shaving and betting in basketball. See Bag and Saha (2011) for a model of match fixing in competitive bookmaking.

⁴² See Benabou and Laroque (1992) for an insightful analysis of manipulation through information in regular asset markets.

⁴³ Deck et al.'s (2013) lab experiment gives more negative results. They find that well-funded manipulators can impact the predictive ability of markets and mislead users of market predictions.

may have an incentive to manipulate the market in order to affect the decisions made on the basis of market outcomes. The scope for this kind of manipulation is moot for prediction markets in which the state of the world is exogenous to the trading process. In political prediction markets, however, the state can become endogenous to the trading process. Manipulating political prediction markets can then be a more cost-effective way to affect public opinion than political campaigning. The possibility of trade-based manipulation then arises if traders are allowed to take large positions; Rothschild and Sethi (2016) report evidence suggestive of manipulation by a single large trader. Trade-based manipulation strategies are also relevant for idea markets or, more generally, for decision markets built for the purpose of improving decision making, as discussed below.

Feedback effects

Lieli and Nieto-Barthaburu (2020) extend the baseline prediction market model to analyze feedback, whereby a decision maker can intervene and change the probability of the underlying event in response to the information revealed by the market price.⁴⁴ They show that the action taken by the decision maker substantially affects how equilibrium prices respond to information. In the presence of feedback, an REE may not exist and there may be partially revealing REE. They characterize conditions under which a fully revealing REE exists, in which the equilibrium price declines in response to information that is ex-ante more favorable to the occurrence of the underlying event. In some cases, there is an equilibrium that does not reveal any information, resulting in a zero price for the asset regardless of the information held by traders.⁴⁵

Empirical analyses of betting markets

Like prediction markets, betting markets feature a specific termination point at which uncertainty is resolved. The outcome is also broadly exogenous to trade, making betting markets natural laboratories for testing economic theories, as argued by Thaler and Ziemba (1988). A large empirical literature has attempted to test the so-called efficient market hypothesis, EMH, according to which competitive asset prices reflect all relevant information. With a lineage that can be traced back at least to Bachelier, the EMH broadly means that it is not possible to outperform the market in a systematic fashion. The various forms of the EMH were crystallized by Fama (1970). According to the weak form, the current price of an asset traded in a competitive market reflects the information contained in past prices. The semi-strong form predicates that an asset's current price should reflect all public information on past, current, and future events. The strong form states that prices reflect all public and private information available to traders.

Betting markets, the most common prediction markets, have long been a testing ground for the EMH. In many jurisdictions, betting is subject to government regula-

⁴⁴ Bond et al. (2010) and Siemroth (2021) analyze similar models and obtain closely related insights.

⁴⁵ See also Meirowitz and Tucker (2004) and Boleslavsky et al. (2017).

tion that restricts gambling and it is typically allowed in association to recreational events or sport games.⁴⁶ Strict government restrictions are also imposed on the operation of betting markets. For example, the US allows bets on the outcome of horse races to be placed only through the pari-mutuel system. UK punters, instead, can choose whether to bet with the (pari-mutuel) tote or with competing bookmakers, known as bookies. Bookies' odds are also called fixed odds, in contrast with pari-mutuel odds which instead vary depending on the bets placed overall on the different outcomes.⁴⁷

Griffith (1949), a psychologist, is credited with the first academic paper comparing the market odds for different horses and the associated empirical frequencies. The early empirical literature was mostly based on aggregate betting data with no information about bets placed by single market participants. More recent work has started using individual-level data. Jullien and Salanié (2000) develop a methodology to estimate the preferences of a representative agent, while allowing for heterogeneity among races. Jullien and Salanié (2008) overview the large empirical literature on betting markets. Here, we highlight some more recent developments.

Comparing pari-mutuel odds for bets on outcomes of single races and on exotic bets corresponding to the combination of outcomes across a number of races, Snowberg and Wolfers (2010) argue that misperceptions of probabilities, rather than risk-loving attitudes, drive the favorite-longshot bias. Also analyzing a large data set from pari-mutuel markets, Gandhi and Serrano-Padial (2015) conclude that the belief heterogeneity model empirically outperforms existing preference-based explanations of the favorite-longshot bias. Chiappori et al. (2019) estimate risk preferences of bettors within a theoretical framework that does not allow for heterogeneous beliefs. Their results point to the importance of nonlinear probability weighting and suggest that several dimensions of heterogeneity may be present.

Empirical performance of prediction markets

Page and Clemen (2013) document that prices in prediction markets with relatively short time expiration are reasonably well-calibrated, but prices for markets predicting events further in the future exhibit a more pronounced favorite-longshot bias. This result is consistent with an increase in the dispersion of beliefs for events in the more distant future.

Prediction market prices cannot contain more information than the traders possess in the first place. As argued by Grossman and Stiglitz (1980), given that market prices reveal information, traders should have limited incentives to acquire information. In a prediction market lab experiment, Page and Siemroth (2017) find that traders tend to acquire more information than is individually optimal in equilibrium. This finding is consistent with the high level of forecasting accuracy observed in prediction markets. In another experimental study, Page and Siemroth (2020) find that prices

⁴⁶ For a primer and references see Pradier (2019).

⁴⁷ Ottaviani and Sørensen (2005) compare the equilibrium in parimutuel markets and fixed-odds in a unified model.

incorporate less than 50% of the private information held by traders, but reflect almost completely public information. This evidence points to semi-strong, rather than strong-form, informational efficiency.⁴⁸

Gauriot and Page (2020) compare the evolution of market prices for bets on football matches, exploiting the occurrence of random information shocks based on whether shots that hit the post end up going in the goal or bouncing off the goal. They find that market prices react surprisingly quickly and efficiently to the arrival of information, with no evidence of average abnormal returns. This null result is precisely estimated, given their large sample size. Nonetheless, when the information shocks are large, toward the end of the markets' life, some significant under-reaction occurs. This result could be due to budget constraints faced by traders, consistent with the theory presented above.

Performance comparison to polls

Morgan and Stocken (2008) propose a model of a poll with strategic survey respondents. As in Crawford and Sobel's (1982) model of strategic information transmission, respondents have incentives to misreport their information in order to influence the outcome of the poll. Kawamura (2013) shows that an increase in sample size dampens the individual influence on the decision and hence worsens the incentive of respondents to exaggerate their reports. The quality of communication with each respondent may improve as the sample size becomes smaller, resulting in a trade-off between the quality and quantity of communication.

A number of empirical studies contrast the performance of prediction markets with polls or surveys of experts. Compared to 964 national polls over 1988–2004, Berg et al. (2008) document that IEM vote-share predictions for presidential elections are closer to the eventual outcome 74% of the time. They also find that prediction markets significantly outperform polls in every election when forecasting more than 100 days in advance. With German data on the premier soccer league, Spann and Skiera (2009) find that prediction markets and betting odds have comparable forecast accuracy, but both methods strongly outperform tipsters. For the 2008 presidential election, Rothschild (2009) finds superior performance of properly de-biased prediction market results from Intrade compared to the victory probability published by Nate Silver's *FiveThirtyEight*, based on an aggregation of de-biased poll results. In a series of high-profile forecasting tournaments at the Intelligence Advanced Research Project Agency, Atanasov et al. (2017) and Dana et al. (2019) find that polls based on self reports—once proper scoring feedback, collaboration features, and statistical aggregation are applied—can be as accurate as prediction-market prices in forecasting a wide range of geopolitical events. Combining forecasts obtained from the two methods further improved accuracy.

⁴⁸ See also Ngangoué and Weizsäcker (2021) for experimental evidence about the difficulty of inferring information from prices that are determined from simultaneous trading rules. Findings along these lines might explain empirical deviations from rational expectations equilibria.

Corporate prediction markets

Plott and Chen (2002) pioneer the implementation of corporate prediction markets at Hewlett-Packard Corporation. In spite of the difficulties in ensuring participation, they found that market forecasts improve over more traditional methods for forecasting sales. Cowgill and Zitzewitz (2015) analyze a wide array of internal markets at Google, Ford Motor Company, and an anonymous basic materials conglomerate, including markets that forecast sales, product quality, project deadlines being met, and external events, as well as with different design and pools of participants, from either the entire company or specific divisions. They broadly conclude that corporate prediction markets outperform the expert forecasts that had been previously used as a basis for decisions. In an HBS teaching case on prediction markets at Google, Coles et al. (2007) stress the difficulties in maintaining interest by informed market participants, as well as in finding effective use of the information collected by management. Also, the initiative risks being boycotted by parties in the organization who do not benefit from increased information transparency.

These early prediction markets were organized as web-based double auctions for binary securities, following the pioneering design of the IEM described by Berg et al. (2008). The double auction design typically works well when it attracts sustained interest and active participation by a large number of traders. However, limited participation often makes prediction markets illiquid, resulting in intermittent trading for long stretches of time and wide bid-ask spreads. To deal with this problem, a number of designers adopted automated market maker rules, such as Hanson's (2003) LMSR. Automated market makers can be easily adapted to forecast events with many possible realizations, such as project completion dates. Othman and Sandholm (2013) report the experience gained running an LMSR market to predict the opening day of a new computer science complex at Carnegie Mellon University.

Yet another design exploits the *pari-mutuel* method for aggregating bets, first analyzed by Borel (1938) and more recently by Lambert et al. (2015). Gillen et al. (2017) report on a long-running field experiment at Intel testing the performance of a *pari-mutuel* information aggregation mechanism for forecasting sales.

Dianat and Siemroth (2020) design laboratory experiments to compare the performance of different market designs. They contrast a straightforward market design for a project deadline market with an asset that pays out 1 if the deadline is met and 0 otherwise. As shown theoretically by Siemroth (2021), with such a *deadline asset* design, employees predicting a missed deadline suffer from a self-defeating prophecy. The manager response triggered by the prediction falsifies the prediction, thus destroying equilibria in which workers reveal their information.⁴⁹ A theoretically preferable market design involves an *action asset* that pays off 1 if the manager assigns additional resources, and 0 if not. For this asset a revealing equilibrium exists, whereby employees trade at high prices in the bad state (and the manager invests

⁴⁹ If employees bet on a missed deadline in the bad state, the price is driven down. This price reaction prompts the manager to assign additional resources. This response by the manager leads to the deadline being met, reducing the profit for the trader. Thus, employees suffer from revealing their information.

additional resources) and at low prices in the good state (suggesting to the manager not to invest); however, in this case there are multiple equilibria, which makes coordination difficult. Dianat and Siemroth (2020) find that the theoretically superior action asset design leads to improved manager decisions compared to the deadline asset only if it is accompanied by advice explaining how the “good” revealing equilibrium strategy works. Overall, this experiment shows the importance and subtlety involved in running effective corporate prediction markets.

Idea markets

Idea markets pit competing projects within a company. In an extended case study, Soukhoroukova et al. (2012) stress the effectiveness of idea markets in assisting management in the process for developing new products. Spears and LaComb (2009) describe an early implementation at General Electric (GE). Company researchers and product developers working in competing teams propose ideas worth financing. An asset is associated to each idea. Ideas are initially screened by a committee of experts and then evaluated according to market prices resulting from trading among a broad set of GE employees, including the proponents of the ideas. During the trading process, market participants can also give feedback to the different ideas. Given that the performance of the projects associated to the ideas is typically unobserved, the market is liquidated on the basis of the final asset prices resulting when the market is terminated. At that point, GE allocates a research budget to the ideas with the highest prices. The proposers of the winning ideas thus obtain highly visible rewards. Incentives for accuracy in the evaluation process are bolstered by allocating small prizes to top traders, who are rewarded based on the value of their portfolio when the market is closed. Ottaviani (2009) discusses how the design of GE idea markets can be improved to contain market manipulation.

Idea markets resemble regular financial markets because asset values are endogenous to the trading process, rather than exogenous as in classic prediction markets. Information about the ex post performance of the underlying idea becomes available, if at all, only for projects that are undertaken, and often with a long delay. Without the ex post validation that characterizes prediction markets, idea market participants are rewarded for their ability to accurately predict the choices of other market participants. Pure idea markets are beauty contests, rife with multiple equilibria. See Marinovic et al. (2013) for a theoretical analysis in the context of a model based on Morris and Shin (2002).⁵⁰

Compared to regular prediction markets, idea markets are more prone to manipulation. Given that the information revealed affects the value of the underlying assets corresponding to the different ideas, idea markets are also subject to the risk of trade-based manipulation, in addition to outcome-based and information-based manipulation. As Goldstein and Guembel (2008) show in a microstructure model

⁵⁰ On the role of incentives in idea generation platforms see also Toubia (2006) and Luo and Toubia (2015).

populated by informed and uninformed traders, an uninformed trader can profit by establishing a short position in the asset and subsequently driving down the price by further short sales. The lower price that is induced is taken as negative evidence inefficiently driving down the investment, thereby reducing the value of the asset and thus allowing the speculator to cover the short position at a lower cost and make a profit.

6 Science as nonmarket

The failures of the market for information identified by Arrow (1962), as reported in the Introduction, are particularly serious for basic research whose input is used in subsequent, more applied, innovations.⁵¹ Given the non-rival nature of knowledge and innovation, governments have long recognized the social value of subsidizing research through a number of tools, including research grants, procurement contracts, R&D subsidies, support to universities, and research prizes. Most work in industrial organization has focused on the market for technology transfer and patents, where markets and explicit incentives naturally play an important role.⁵² However, science largely operates outside the market and so far has been studied mostly by historians, sociologists, and philosophers of science.⁵³

We believe that the marginal value of economic analysis in science is high. So far, science has attracted limited attention from economists, in spite of Stigler's (1983) invitation. The organization of science is one of the central problems in modern societies. Over the last decades, rich data on scientists, publications, and citations has become available, spurring the developments of interdisciplinary fields in the area of information science. Journals and funding agencies are increasingly willing to share data to improve the design of their procedures for selecting articles for publication and research projects for funding. We believe the time is ripe to deploy the tools of information economics, agency theory, and mechanism design to solve the incentive problems that are rife in science.

After surveying some classic contributions to the sociology of science, this section presents some recent developments in the economics of science, while emphasizing some important open questions. We refer to Dasgupta and David (1994) and Stephan (2012) for broader introductions to the area.

⁵¹ According to Arrow (1962), the market system is unfit to yield an efficient production of new information for the following reasons (see also the introduction): First, adding to the stock of knowledge is a business fraught with great risks, which are mostly uninsurable due to moral hazard problems. Second, on the supply side, once produced, information can be replicated at zero (or very low) cost, so it should be made widely available to achieve efficiency. Third, turning to the demand side, the value of information becomes known only after purchase, so that trading becomes particularly difficult.

⁵² See Scotchmer (2004) for an overview of the economics of innovation.

⁵³ Godfrey-Smith (2009) presents an accessible introduction to philosophy of science.

6.1 Organization of science

Merton norms

In a defining contribution to the sociology of science, Merton (1942) introduces the institutional imperatives that characterize the ethos of modern science: communalism, universalism, disinterestedness, and organized skepticism. According to the communalism norm, scientific ideas should be common property, as “the substantive findings of science are the product of collaboration and are assigned to the community.”⁵⁴ Universalism prescribes that scientific ideas have universal value, independent of the personal attributes and social background of the proponents. Scientists should also be disinterested and act for the benefit of the common scientific enterprise, rather than for personal gain. Finally, organized skepticism refers to the community-wide norm of openly challenging and testing ideas, instead of accepting them on trust. Merton’s four principles are often summarized as the CUDO norms, or CUDOS.⁵⁵

History of science

When modern science developed as a blend of experimentalism and Renaissance mathematics in the sixteenth and seventeenth centuries, aristocratic patrons had to be convinced to finance scientists rather than artists, who could offer immediate ornamental value. As David (1998) argues, scientists started disclosing their discoveries to peers who could vouch for the quality of the contributions. This hypothesis views openness as a certification mechanism.

Shapin (1994) documents how in seventeenth-century England the modern scientific culture emerged in gentlemen’s circles around Robert Boyle, one of the founders of the Royal Society, which started publishing *Philosophical Transactions*, the world’s first science journal. As reported by Crosland (2002), in the first half of the nineteenth century the French Académie des sciences pioneered the adoption of a number of modern practices for evaluating and funding research, such as refereeing submitted articles, awarding small research grants (“encouragements”) to allow leading scientists to carry out promising projects, and measuring citations to evaluate candidates. As explained by Ben-David (1970), around 1840 France lost its lead in research to England and then Germany.⁵⁶

⁵⁴ Merton (1942): “Property rights in science are whittled down to a bare minimum by the rationale of the scientific ethic. The scientist’s claim to ‘his’ intellectual ‘property’ is limited to that of recognition and esteem which, if the institution functions with a modicum of efficiency, is roughly commensurate with the significance of the increments brought to the common fund of knowledge. Eponymy—for example, the Copernican system, Boyle’s law—is thus at once a mnemonic and a commemorative device.”

⁵⁵ Ziman (2000) conveniently splits Merton’s (1942) fourth norm into originality and skepticism. Scientists should strive to make original contributions, while expecting to defend their claim to fame against the attacks of skeptical peers.

⁵⁶ Ben-David (1977) documents how modern Western institutions for the production and diffusion of academic knowledge evolved historically in England, France, Germany, and the United States.

Science reward system

Mertonian ideals depict a somewhat utopian world of open science.⁵⁷ There is an unresolved tension between these Mertonian prescriptive norms and the individual motivations of actual scientists, who might not always embrace these noble ideals. In another classic contribution, Merton (1957) describes the reward system in science. Recognition for being the first to make a discovery is the basic currency for scientific reward. This recognition is private property of the discoverer. However, according to the norm of communalism, common scientific property is established once an idea is published. In the best-case scenario, a scientist is rewarded for naming the idea—however, recognition creates conflict among multiple discoverers. Even if scientific fraud is rare, adherence to the norms is likely to be far from perfect. Lacetera and Zirulia (2011) develop an economic model of scientific misconduct. For empirical evidence see Fanelli (2009) on fraud and Azoulay et al. (2017) on retractions.

Stern (2004) quantifies the wage scientists are willing to give up (compensating differential) for the freedom to pursue their own individual research agendas, as well as for the ability to publish their work. Aghion et al. (2008) model academia as a commitment to allow individual scientists to pursue their own preferred research strategies.⁵⁸ Because scientists' motivations are largely intrinsic, to understand the organization of science it is important to take into account the interaction of social preferences with explicit incentives, following the lead of Kreps (1997), Sobel (2005), Besley and Ghatak (2005), Bénabou and Tirole (2006), and Ellingsen and Johannesson (2008). Cohen et al. (2020) evaluate empirically how scientists in different fields trade off academic against commercial returns.

Priority, disclosure, and citations

Given that scientific knowledge is a public good, according to a long-standing argument (see Arrow, 1962), we should expect market provision to be inefficient. Priority can be seen as an effective mechanism to incentivize the private production of scientific knowledge. As stressed by Merton (1988), priority awards a prize that consists in the public recognition of the first individual who discovers and clearly explains a new idea. As Dasgupta and David (1987) explain: “Priority creates a privately-owned asset—a form of intellectual property—from the very act of relinquishing exclusive possession of the new knowledge.”

Hull (1988) highlights that to further link the priority prize to the user value of the idea, the reward is made proportional to the number and quality of citations. Scientists

⁵⁷ Merton (1942): “The institutional conception of science as part of the public domain is linked with the imperative for communication of findings. Secrecy is the antithesis of this norm; full and open communication its enactment. The pressure for diffusion of results is reinforced by the institutional goal of advancing the boundaries of knowledge and by the incentive of recognition which is, of course, contingent upon publication.”

⁵⁸ Craig and Vierø (2013) analyze an equilibrium model where heterogeneous researchers sort into working for universities or an outside sector. Researchers care about peer effects, their relative status within the university, as well as salary compensation.

operate within a social structure that leverages their individual curiosity to discover new knowledge. Each scientist, motivated by personal curiosity and advancement, builds and expands on the ideas and methods laid out by previous generations of scientists in the field. Of course, the really successful scientists manage to revolutionize the field (see Kuhn, 1962), but even revolutionary work takes full stock of the state of the art in the field. Individual scientists operate within a system of cooperation and trust. To develop their own ideas, scientists use the work of others in ways that provides support for what they are doing. The lifeblood and currency of science is the scientific credit on which the entire reputation system is built. By citing relevant previous work, authors reassure users that they build on sound foundations. Thus, scientists trade credit for support, in the hope that others will do the same for them. In a sense, this reciprocation originates from, or at least is broadly compatible with, self-interest. According to this invisible hand argument, the disclosure, priority, and citation system directs the individual curiosity and motivation of individual scientists toward what is useful for the scientific community as a whole.⁵⁹

This way, the science reward system can be seen as an incentive system made up of a number of complementary elements, as in Holmstrom and Milgrom (1994). It would be useful to formalize this interesting argument and address questions like: When does the introduction of commercial incentives worsen the performance of the priority system? Can the secrecy associated with commercial science threaten the CUDOS norms? Which alternative mechanisms could perform better?

6.2 Measuring influence

The availability of publication data is closely tied to the development of scientometrics, a meta-research field spearheaded by Garfield (1955) and De Solla Price (1965). Rooted in library and archival science, scientometrics and information science combine network theory with data analysis with the aim of developing effective measures of the impact of publications.⁶⁰ Work in this area has been blossoming with improvements in information technology in the last decades and with transdisciplinary collaborations among social, natural, and physical scientists.⁶¹

So far, scientometrics has had rather limited points of contact with economics. After using citation data to measure research productivity within economics (see, e.g., Liebowitz and Palmer, 1984), economists started contributing insights that are more generally applicable to the evaluation of science (see Stigler et al., 1995; Ellison, 2002, 2013). In a contribution at the interplay of economics and scientometrics, Palacios-Huerta and Volij (2004) provide an axiomatic characterization of the invariant method that assigns to each journal a weighted sum of its citations, with weights equal to the scores of the citing journals divided by their respective reference in-

⁵⁹ See also Kitcher (1995), Strevens (2003), Hull (1997), and Kleinberg and Oren (2011) for other contributions along these lines.

⁶⁰ See Moed (2006) for a systematic overview of the main issues and techniques of citation analysis.

⁶¹ Fortunato et al. (2018) introduce the frontier in the science of science.

tensities. The invariant method implemented through the article influence score (see eigenfactor.org, West et al. (2010)) is gaining ground over Garfield's (1972) impact factor, equal to the normalized average number of citations obtained by the articles published in a journal.

In this line of research, Demange's (2014) handicap method associates to each journal two values: (i) a score, which measures the impact of the journal through citations it receives from other journals; and (ii) an incompetence weight, capturing the influence the journal obtains from the other journals it cites. Scores and weights are simultaneously determined: a journal's score is the sum of its citations weighted by the citing journal's competence, while the incompetence of a journal is the weighted sum of its references, where the weights are the inverse of the score of the cited journal. Perry and Reny (2016) study the axiomatic properties of measures of intellectual influence of scholars who publish a number of articles. We refer to Palacios-Huerta and Volij (2014) for an overview of the literature on axiomatic properties of measures of intellectual influence.

6.3 Selection, publication, and funding

Science can be a great laboratory for studying information collection and aggregation. As richer data sets are shared by repositories and funding agencies, it should be possible to dissect the layers of selection behind the production of scientific knowledge:

- Individuals self-select into becoming scientists depending on their interests.
- Scientists select their field of study, the topic to work on, and the site where they undertake their investigation.
- Research funding agencies select which scientists to fund after soliciting opinions from selected reviewers.
- Editors select reviewers and eventually select which papers to publish among the many that are submitted.
- Readers select articles they find worth reading.

These different levels of selection determine the information that is created by the scientific community and that is eventually added to the body of knowledge of society at large. Selection in research can be unpacked into a downstream and an upstream level: (i) At the downstream level individuals make choices, such as deciding whether to become researchers, what to work on, where to conduct the experiment, which data sample to analyze, and which articles to read. (ii) At the upstream level, organizations such as funding agencies, universities, and private foundations support research by managing selection mechanisms and funding schemes. The evolution of knowledge in society depends on the overall balance of incentives in the system. Science policy designs these incentives.

In a seminal paper, Zuckerman and Merton (1971) compared acceptance rates at leading scholarly journals across academic disciplines. They hypothesized that the higher rejection rates in the social sciences and humanities than in the physical sci-

ences is associated with a lower level of consensus. While Hargens (1975, 1988) and Cole (1983) find mixed evidence, the time might be ripe to revisit the question of differential selection by explicitly modeling the self-selection by authors into journal submissions, as well as the noisy evaluation by reviewers.

The fledgling literature on research and the editorial process leading to publication started investigating issues such as: the role of submission fees (Leslie (2005) and Cotton (2013)), anonymity of reviews (Taylor and Yildirim (2011)), sequential submissions (Muller-Itten (2019)), and the role of editors (Krieger et al. (2021)). Building on contest theory, Bobtcheff et al. (2017) model incentives in a priority race in which researchers trade off the benefit of marginal improvement to their innovation with the risk of being scooped by competing researchers. Hill and Stein (2021b) find that scooped scientists still obtain a non-negligible fraction of credit and citations. Hill and Stein (2021a) quantify the reduction in quality that results when competition among researchers becomes fiercer in the field of structural biology, where macromolecules can be assigned objective quality measures.

The related problem of the organization and design of research funding is now attracting renewed attention; see Azoulay and Li (2020) for an overview of the recent empirical literature. Among recent papers, we single out Li (2017) estimating the tradeoff between bias and informativeness of reviewers at the National Institutes of Health (NIH) and Azoulay et al. (2019) quantifying the impact of NIH grants on pharmaceutical and biotechnological innovation and patents. On the theory front, Ottaviani (2020) develops a foundational model of grantmaking; the model highlights how the method used to apportion grant prizes across different fields of research interacts with the incentives to submit grant applications, thus determining the overall performance of the resulting allocation.

While we are still far from having a framework encompassing the different layers of selection listed at the beginning of this section, Carnehl and Schneider (2021) propose a tractable model of how the organization of science affects the social accumulation of knowledge. The framework could also be used to assess recent trends and calls for increased openness in science; on the emergence of crowd science see, e.g., Franzoni and Sauermaun (2014).

6.4 Economics of statistical inference

Statistical decision theory, the foundation on which modern statistics is built, sees the statistician as a decision maker who aims at establishing the truth or intends to make an optimal decision. In reality, however, empirical research is carried out by researchers who may pursue their own objectives and communicate their research findings to other decision makers with different objectives. A more realistic theory of statistics should take into account the social environment in which evidence is collected and communicated.

As applications of statistics to demography became *à la mode* in France in the first half of the nineteenth century—Cournot (1843, §103) warns “against its premature and excessive applications which can discredit it for a while and delay the so

desirable epoch when the materials of the experience will serve as a certain basis for all the theories aimed at the diverse parts of the social organization.” In Chapters 9 and 10 of his book on the theory of chance and probability, Cournot presents a pioneering account and perceptive analysis of statistical inference in the presence of strategic sample selection. Cournot questioned the interpretation of results about the anomalous preponderance of male sex at birth when the data is sliced along specific dimensions, such as the season of birth, whether the child is legitimate or illegitimate, as well as the residence, age, profession, wealth, and religion of the parents:

“§101 ... A person not knowing how the data were analysed and whom the experimenter told the result of that analysis concerning the system ... but not how many attempts he made to achieve that result, is unable to judge with a determined chance of error whether the chances ... are equal or not... §111 ... for the statistician occupied by grouping and comparisons, the probability that a given difference is not attributable to anomalies of chance will take very different values depending on the number of groups tested before encountering that difference. ... However, unsuccessful tests usually leave no traces; the public only knows the results which the experimenter thought to be deserving notice. It follows that a person alien to the testing is absolutely unable to regulate bets on whether the result is, or is not attributable to anomalies of chance.”

Cournot’s (1843) analysis has remained largely unnoticed. Fisher (1936) voices preoccupation with data manipulation and fraud in his forensic investigation into Mendel’s empirical analysis of the laws of inheritance. Bonferroni (1936) proposes a methodology for explicitly adjusting for testing of multiple hypotheses. As argued by Di Tillio et al. (2017), robustness to manipulation by researchers played a key role in the early development and adoption of randomized controlled trials. While randomization is generically suboptimal for a Bayesian decision maker, it becomes optimal when facing an adversarial audience, as shown by Banerjee et al. (2020), provided that the sample size is large and robustness to the audience’s beliefs is an important concern.

Strategic sample selection

Di Tillio et al. (2021) characterize the welfare impact of anticipated sample selection for a decision maker. Are the highest realizations of a sample more or less informative than the same number of random realizations? For illustration, consider simple hypothesis testing where acceptance is optimal in state $\theta = \theta_H$ and rejection is optimal in state in $\theta = \theta_L$. For simplicity, take sample size $n = 1$ and compare a random location experiment $x = \theta + \varepsilon$ with noise distribution $\varepsilon \sim F$ to a selected experiment with noise distribution $\varepsilon_{(k)} \sim F^k$ resulting from a researcher who reports the highest realization observed from a presample of size $k > n = 1$. As Di Tillio et al. (2021) show, an increase in presample size k results in an increase in the payoff of the decision maker if and only if the reverse hazard function of the noise distribution, $-\ln F$, is logconcave, as with normal or logistic noise. In the boundary case with noise drawn from a Gumbel extreme value distribution, selection shifts up the location, but does

not change the shape, of the noise distribution. Deviating from Gumbel, maximal selection not only induces a first-order stochastic dominance upshift in the realizations, but also changes the shape of the distribution. The pushed-up realizations become less (respectively more) dispersed when the noise distribution is a concave (respectively convex) transformation of, and hence has a thinner (respectively thicker) top tail than, a Gumbel distributed random variable. Selection damages the decision maker when the reverse hazard function of the noise distribution is logconvex, as with exponential noise. For a general sample size n and general experiment, selection increases (respectively, decreases) information accuracy if the reverse hazard rate of the noise distribution $f(x|\theta)/F(x|\theta)$ is log-supermodular (respectively, log-submodular).

Research bias and the approval process

Henry and Ottaviani (2019) model research as a sequential information acquisition process by a biased agent (such as a pharmaceutical company running clinical trials) who aims at persuading an evaluator (such as the Food and Drug Administration) to approve a treatment. Information acquisition is costly and the agent discloses the information acquired when submitting an application for approval to the evaluator. A number of organizational structures are considered depending on how the commitment power is distributed between the agent and the evaluator.

When the agent can commit not to undertake additional research if the evaluator rejects the application, as the cost of acquiring information converges to zero, the equilibrium converges to the Bayesian persuasion solution characterized by Kamenica and Gentzkow (2011). More generally, Henry and Ottaviani (2019) characterize how the evaluator can curb persuasion bias by committing to raise the standard for approval. This way false positives are reduced at the expense of false negatives—given that the agent is discouraged from acquiring information when the balance of the evidence collected is sufficiently unfavorable. As they show, delegating authority to the informer is socially optimal when information acquisition is sufficiently costly. This model can be applied to approval regulation as well as the publication process.

Publication bias and the file drawer

In recent years, the credibility of results published in academic journals has been called into question across a number of scientific fields. When researchers write up and submit and/or journals publish only studies containing statistically significant findings, an unspecified number of negative results tend to remain unpublished in the file drawer, as noted by Rosenthal (1979). The null hypothesis could well be false when the 5% of studies that constitute Type I errors are published, while the 95% that show non-significant results are not published. By comparing published and unpublished social science studies resulting from an NSF-sponsored program, Franco et al. (2014) document that strong results are 40% more likely to be published than null results and 60% more likely to be written up.

In recent years awareness about publication bias has been rising in a number of disciplines, starting from medicine and psychology, but now also in economics (see Ioannidis, 2005; Simonsohn et al., 2014; and Brodeur et al., 2016). As argued by

Gall et al. (2017) and Di Tillio et al. (2017), this is not only a challenge but also a great research opportunity for economists who have the tools to analyze statistical inference and experimental design taking into account the incentives of researchers. A particularly interesting and open area is the design of research procedures that take into account the incentives of researchers to select questions, collect data, perform analysis, and communicate results. Christensen and Miguel (2018) overview the burgeoning literature on research credibility and publication bias.

Andrews and Kasy (2019) advance the econometric methodology for identifying and correcting for publication bias with data from either replication studies or meta studies. Replication studies involve a new sample from the same population and the same experimental protocol as the original study. The key assumption on which the methodology relies is that publication decisions depend only on the original estimates—this assumption is naturally satisfied for systematic replication. It is then possible to nonparametrically identify the conditional publication selection function exploiting asymmetries in the joint distribution of original and replication estimates. This powerful approach is based on the following simple intuition: in the absence of publication selection bias, the probability of observing a certain realization of z statistic in the original studies and statistic z' in the replication studies should be the same as the probability of observing z' in the original studies and z in the replication studies, because the data-generating process is symmetric. If these probabilities are different, their ratio identifies the ratio of the probabilities of the publication selection function $p(z)/p(z')$.⁶²

Meta studies collect estimates and standard errors from multiple published studies. The approach to correct for selection effect relies on a (rather natural) independence assumption between the estimands and the standard errors, so that studies with smaller standard errors do not have systematically different estimands. If publication is more likely when the results adjusted for standard error are above the significance threshold, the estimated results tend to be higher for larger standard errors. The publication selection function is identified by the way in which the conditional distribution of the estimated results varies with the standard error. Intuitively, under no selectivity it should be possible to obtain the distribution of estimates for studies with larger standard errors by adding noise to the distribution for studies with lower standard errors. Deviations from this prediction identify conditional publication probabilities.

Analyzing data from Camerer et al. (2016) and Open Science Collaboration (2015), Andrews and Kasy (2019) find strong evidence of selective publication. Even though the replication and meta-studies methodologies are based on different identifying assumptions, they deliver similar estimates of selectivity: results significant at the 5 percent level are over 30 times more likely to be published compared to

⁶² Through this methodology, for example, it is possible to estimate the ratio of [the probability that a significant result z is published] to [the probability that an insignificant result z' is published] based on the empirical ratio of [the frequency of the joint outcome corresponding to a significant original z and an insignificant replication z'] to [the frequency of an insignificant original z' and a significant replication z].

insignificant results. Applying the meta-analysis methodology to the impact of minimum wage on employment, they find that results indicating a negative and significant effect of minimum wages on employment are about three times more likely to be published than insignificant results. Focusing on the reporting of significant results, they also find that results showing a positive impact of minimum wages on employment are less likely to be published than results showing a negative impact.

Frankel and Kasy (2021) propose a normative framework for evaluating publication bias. Suppose that publishing results is costly, for example, because the audience of journals has limited attention. Considering publication cost is reasonable because in the absence of any attention cost by the audience all data should be published, leaving no role for journals. If publication is costly, it is not worth publishing null-results that do not change the default decision that would result without publication. Surprising results, which are highly valuable to decision makers, instead should be published.

Building the literature on disclosure, Bertomeu (2021) proposes a positive model of the review process where some reviewers are biased either in favor or against the submitter. He shows that when the type distribution of submitters has heavy tails, the probability of acceptance due to the bias is bounded away from zero, even if almost all reviewers are unbiased.

Registration

Adda et al. (2020) systematically analyze the p-values reported to the *ClinicalTrials.gov* registry. First, they find no cluster of results just above the significance threshold, in sharp contrast with findings for p-values obtained from results reported in journal publications across a number of disciplines. Second, they document that research sponsors with more favorable phase II results are more likely to continue to phase III investigations. Thus, incentives matter for determining the investigators' decision to proceed to further investigation. Finally, they are unable to explain part of the excess amount of phase III significant results relative to phase II for small industry sponsors. This pattern is consistent with some selective reporting and suggests the need for tightening enforcement of registration regulations, especially for small industry sponsors.

Exploiting newly assembled data on psychiatric randomized controlled trials (RCTs), Oostrom (2021) analyzes the results of clinical trials conducted on the *same* drugs, depending on the financial interests of the parties sponsoring the trials. She estimates that a drug is about 0.15 standard deviations more effective when the RCT is sponsored by the drug maker. Half of this effect is due to selective publication—sponsored studies without positive results are likely to remain unpublished. The effect decreases over time as pre-registration requirements are met. This evidence suggests that preregistration is at least partially effective in overcoming sponsorship bias.

Abrams et al. (2021) evaluate the registration of economics experiments at the AEA RCT Registry since its inception in 2013. Similar to what has been found for clinical trials, they document that mandatory registration before publication at the AEA journals is enforced only partially, even though the registry does not require

preregistration before the experiments are conducted. They argue in favor of tightening the standards for registration, for example by not allowing late registration.

Careful theoretical modeling of equilibrium reaction of researchers in settings with private information reveals subtle results. Once the incentives to acquire information by researchers are taken into account, Herresthal (2021) shows that a hidden testing regime (in which the researcher voluntarily discloses findings) Pareto dominates observable testing (with mandatory disclosure) provided the researcher is sufficiently biased. In a contest model, Bobtcheff et al. (2021) show that under winner-take-all competition mandating the disclosure of negative results can backfire. These results underlie the importance of finding ways to properly design the objectives and incentives of researchers.

Controlling p-hacking

Echenique and He (2021) propose an interesting mechanism for controlling p-hacking. They show that an evaluator is able to effectively screen researchers by adding noise to the data before dissemination to the research community. Intuitively, researchers who p-hack perform more poorly on the full data than honest researchers who do not manipulate the inference. This approach for controlling p-hacking is practical and low cost, given that dissemination noise is a tool already used by statistical agencies to protect privacy.

Learning and paradigm selection

Does the scientific process result in the selection of superior paradigms? This question can be addressed in the context of the Sethi and Yildiz (2012) model of learning among a number of informed individuals with heterogeneous prior beliefs. They show that if individuals sequentially and publicly announce their posterior beliefs about the state, as in Geanakoplos and Polemarchakis (1982), the information the individuals possess is not fully aggregated when the prior beliefs are independently distributed. Sethi and Yildiz (2016), instead analyze the dynamic evolution of beliefs when individuals, each with heterogeneous priors and varying but publicly observed information quality, choose which other individual to consult in each period. They show that if the variance of the distribution of prior beliefs is sufficiently high, a trade-off emerges between familiarity (i.e., information about the prior acquired in the past) and quality of information, leading to information segregation. Individuals who do not have the highest information quality end up becoming opinion leaders. This mechanism generates own-field bias and large-field dominance.

Pushing further insights from earlier work by Kitcher (1995) and Brock and Durlauf (1999), Akerlof and Michailat (2018) show that when evaluators are biased in favor of (or against) candidates with similar (or different) beliefs, beliefs converge, but not necessarily to the superior paradigm.

Replication markets

To assess the reproducibility of published results, large-scale replication projects have been conducted in a number of academic fields. However, replications are time con-

suming and costly. Can prediction markets help improve the allocation of resources into replication?

Dreber et al. (2015) pioneer the application of prediction markets to the assessment of whether published results are replicable. The idea is to ask market participants to bet on the binary event corresponding to whether or not a study (or hypothesis) will replicate, where successful replication is typically defined as a result in the same direction and with the same significance threshold as the original published study. Do prediction market prices indicate accurately whether studies replicate? Prediction market participants were endowed with a small initial budget (USD 100) and allowed to trade for roughly two weeks. According to the mean prediction market price, about half of the studies were expected to replicate. From the 41 studies investigated, which were sampled from three top journals in psychology in conjunction with Open Science Collaboration's (2015) "Reproducibility Project: Psychology", 16 (39%) replicated and 25 (61%) did not replicate. The markets correctly predicted 29 of 41 replications (71%), where success is based on whether the price indicates a replication probability over 50%.

Camerer et al. (2016) replicate results for 11 out of 18 experimental economics studies published in the *American Economic Review* and the *Quarterly Journal of Economics* in 2011–2014. The prediction markets they run, mostly with experimental economists participating, reveal an optimistic faith in replication (all prices above 50, with the mean 75%). Broadening the fields investigated, Camerer et al. (2018) replicate only 62% of results from twenty-one social science experiments originally published in *Nature* and *Science* in 2010–2015, even though they used much larger sample sizes than the original studies and disclosed their replication plan to the authors of the original experiments. Prediction markets, run with a mix of 206 social scientists who were not involved with the replication exercise, correctly predicted all the studies that replicated.

In parallel to prediction markets, some of these replication projects run surveys to elicit beliefs about replication probabilities. Prior to participating in the prediction markets, for each study participants were asked to indicate the probability that the replication would yield a statistically significant effect in the same direction as the original study. Pooling the data from a number of projects, Gordon et al. (2020) find that the prediction surveys and prediction markets have hit rates of 66% and 73%, respectively. Forsell et al. (2019) find that prediction markets correctly predicted 75% of the replication outcomes, performing better (respectively worse) than survey data in predicting replication outcomes (respectively effect size). Given these promising results, it might be possible to use prediction markets to decide which studies should be replicated, thus saving on replication costs.

Pre-research information

More generally, DellaVigna et al. (2019) argue that it is valuable to collect predictions of a study's expected results before the research is carried out. Predictions collected from the scientific community, policy-makers, as well as the general public, can serve as a useful benchmark against which to compare ex post results. Thanks to these

forecasts, the research community should be able to make research results more informative in a number of ways:

- First, by making it possible to relate ex post research findings to prior expectations as captured by pre-study forecasts, it becomes possible to measure how surprising results are relative to previous public opinion. The results from replication markets summarized above indicate that prediction market participants understand pretty well which studies are likely to replicate.
- Second, by allowing the comparison of results against the average prediction of experts, rather than the null hypothesis of no effect, ex ante forecasts can be used to counteract publication bias.
- Third, the experiment itself could benefit from the forecasts, for example by estimating the value of information for different treatment arms.

DellaVigna et al. (2020) present the design of www.socialscienceprediction.org, a new online platform designed to allow researchers to collect forecasts of the effects of social programs before testing. Rather than prediction markets, this platform makes use of incentivized surveys.

Economic design of experiments

Applying the powerful tool of mechanism design, economists have started to analyze the optimal design of experiments taking into account incentive problems. In this promising new vein, Chassang et al. (2012) develop a technique to improve the estimate of heterogeneous treatment effects by eliciting information from subjects through selective trials that are differently appealing for subjects.

Narita (2021) proposes an experimental design that incorporates the welfare of the experimental subjects. The optimal procedure endows the subjects with a common artificial budget. Subjects can then use the budget to purchase their most preferred bundle of treatment-assignment probabilities. Personalized prices are set so as to make a treatment cheaper for subjects with better predicted treatment effects. In addition to being Pareto optimal, this procedure results in an incentive-compatible preference elicitation when the number of subjects is large, and gives unbiased estimates of the causal effect that can be estimated with standard RCTs.

7 Conclusion

Over fifty years ago, Marschak (1968) suggested a useful classification by dividing the economics of information into three subproblems that he referred to as the economics of inquiring, communicating, and deciding. The literature since then made substantial progress in many of these areas. Arguably, the economics of communicating remains perhaps the least developed. And while as researchers we are keenly aware of the limits of communication, in economics we have not yet developed many models on the constraints and costs of communicating and transmitting information. Nowhere is this perhaps more easy to illustrate than in the theory of mechanism

design, where the revelation principle as a statement of unconstrained communication and transmission of information has not yet found a corresponding principle where those constraints are acknowledged and made operational. Specifically, when the amount of private information (the type space) of the agents is large, the direct revelation mechanism requires the agents to have abundant capacity to communicate with the principal and the principal to have abundant capacity to process information. Thus we may ask what happens when the agents can communicate only limited amounts of information, and/or the principal can process only limited amounts of information. Wilson (1989), McAfee (2002), and Bergemann et al. (2021d) begin to address this issue in the context of nonlinear pricing and matching environments where the agents only have access to finite coded information, even though the underlying information source is represented by a continuum.

Marschak (1968) suggested that: “The informational revolution is upon us, and the manipulation of symbols dominates our lives more and more.” He concluded: “I hope we shall soon understand how to harness and benefit from those trends in our culture.” This survey indicates that we have just started to scrape the surface and much work is left to understand the broad implications of the informational revolution for the organization of economic activity.

References

- Abrams, E., Libgober, List, J., 2021. Research Registries and the Credibility Crisis: an Empirical and Theoretical Investigation. Discussion paper. Chicago and USC.
- Acemoglu, D., Makhdoumi, A., Malekian, A., Ozdaglar, A., 2019. Too Much Data: Prices and Inefficiencies in Data Markets. Discussion paper. MIT.
- Acquisti, A., Taylor, C., Wagman, L., 2016. The economics of privacy. *Journal of Economic Literature* 54 (2), 442–492.
- Adda, J., Decker, C., Ottaviani, M., 2020. P-hacking in clinical trials and how incentives shape the distribution of results across phases. *Proceedings of the National Academy of Sciences* 117 (24), 13386–13392.
- Ades, A., Lu, G., Claxton, K., 2004. Expected value of sample information calculations in medical decision modeling. *Medical Decision Making* 24, 207–227.
- Admati, A.R., Pfleiderer, P., 1986. A monopolistic market for information. *Journal of Economic Theory* 39 (2), 400–438.
- Admati, A.R., Pfleiderer, P., 1988. Selling and trading information in financial markets. *The American Economic Review* 78 (2), 96–103.
- Admati, A.R., Pfleiderer, P., 1990. Direct and indirect sale of information. *Econometrica* 58 (4), 901–928.
- Aghion, P., Dewatripont, M., Stein, J.C., 2008. Academic freedom, private-sector focus, and the process of innovation. *The Rand Journal of Economics* 39 (3), 617–635.
- Agrawal, A., Gans, J., Goldfarb, A., 2018. *Prediction Machines*. Harvard Business Review Press.
- Akerlof, G.A., Michailat, P., 2018. Persistence of false paradigms in low-power sciences. *Proceedings of the National Academy of Sciences* 115 (52), 13228–13233.

- Ali, M.M., 1977. Probability and utility estimates for racetrack bettors. *Journal of Political Economy* 85 (4), 803–815.
- Ali, N., Lewis, G., Vasserman, S., 2019. Voluntary Disclosure and Personalized Pricing. Discussion paper. Pennsylvania State University.
- Andrews, I., Kasy, M., 2019. Identification of and correction for publication bias. *The American Economic Review* 109 (8), 2766–2794.
- Anton, J., Yao, D., 1994. Expropriation and inventions: appropriable rents in the absence of property rights. *The American Economic Review* 84 (1), 190–209.
- Aridor, G., Che, Y.-K., Salz, T., 2020. The Economic Consequences of Data Privacy Regulation: Empirical Evidence from GDPR. NBER Working Paper, 26900.
- Armstrong, J.S., 2001. *Principles of Forecasting: A Handbook for Researchers and Practitioners*. Springer.
- Arrieta-Ibarra, I., Goff, L., Jimenez-Hernandez, D., Lanier, J., Weyl, G., 2018. Should we treat data as labor? Moving beyond ‘free’. *AEA Papers and Proceedings* 108, 38–42.
- Arrow, K., 1962. Economic welfare and the allocation of resources for invention. In: Nelson, R. (Ed.), *The Rate and Direction of Inventive Activity*, pp. 609–626.
- Arrow, K.J., Forsythe, R., Gorham, M., Hahn, R., Hanson, R., Ledyard, J.O., Levmore, S., Litan, R., Milgrom, P., Nelson, F.D., et al., 2008. The promise of prediction markets. *Science* 320 (5878), 877–878.
- Atanasov, P., Rescober, P., Stone, E., Swift, S.A., Servan-Schreiber, E., Tetlock, P., Ungar, L., Mellers, B., 2017. Distilling the wisdom of crowds: prediction markets vs. prediction polls. *Management Science* 63 (3), 691–706.
- Aumann, R.J., 1976. Agreeing to disagree. *The Annals of Statistics* 4 (6), 1236–1239.
- Avery, C., Resnick, P., Zeckhauser, R., 1999. The market for evaluations. *The American Economic Review* 89 (3), 564–584.
- Azevedo, E., Deng, A., Olea, J.M., Rao, J., Weyl, G., 2020. A/B testing with fat tails. *Journal of Political Economy* 128 (12), 4614–4672.
- Azoulay, P., Bonatti, A., Krieger, J.L., 2017. The career effects of scandal: evidence from scientific retractions. *Research Policy* 46 (9), 1552–1569.
- Azoulay, P., Graff Zivin, J.S., Li, D., Sampat, B.N., 2019. Public R&D investments and private-sector patenting: evidence from NIH funding rules. *The Review of Economic Studies* 86 (1), 117–152.
- Azoulay, P., Li, D., 2020. Scientific Grant Funding. NBER Working Paper, 26889.
- Babaioff, M., Kleinberg, R., Paes Leme, R., 2012. Optimal mechanisms for selling information. In: *Proceedings of the 13th ACM Conference on Electronic Commerce. EC '12*, Valencia, Spain, pp. 92–109.
- Bag, P.K., Saha, B., 2011. Match-fixing under competitive odds. *Games and Economic Behavior* 73 (2), 318–344.
- Baillon, A., Xu, Y., 2021. Simple bets to elicit private signals. *Theoretical Economics* 16 (3), 777–797.
- Bajari, P., Chernozhukov, V., Hortascu, A., Suzuki, J., 2019. The impact of big data on firm performance: an empirical investigation. *AEA Papers and Proceedings* 109, 33–37.
- Ball, I., 2020. Scoring Strategic Agents. Discussion paper. Yale University.
- Banerjee, A.V., Chassang, S., Montero, S., Snowberg, E., 2020. A theory of experimenters: robustness, randomization, and balance. *The American Economic Review* 110 (4), 1206–1230.
- Banerjee, S., 2021. Does competition improve analysts’ forecast informativeness? *Management Science* 67 (5), 3219–3238.

- Banerjee, S., Kremer, I., 2010. Disagreement and learning: dynamic patterns of trade. *The Journal of Finance* 65 (4), 1269–1302.
- Bates, J.M., Granger, C.W., 1969. The combination of forecasts. *Journal of the Operational Research Society* 20 (4), 451–468.
- Bayarri, M., DeGroot, M., 1988. Gaining weight: a Bayesian approach. In: *Bayesian Statistics 3: Proceedings of the Third Valencia International Meeting*, Vol. 3, pp. 25–44.
- Bayarri, M., DeGroot, M., 1989. Optimal reporting of predictions. *Journal of the American Statistical Association* 84 (405), 214–222.
- Ben-David, J., 1970. The rise and decline of France as a scientific centre. *Minerva* 8 (2), 160–179.
- Ben-David, J., 1977. *Centers of Learning. Britain, France, Germany, United States*. McGraw-Hill.
- Benabou, R., Laroque, G., 1992. Using privileged information to manipulate markets: insiders, gurus, and credibility. *The Quarterly Journal of Economics* 107 (3), 921–958.
- Bénabou, R., Tirole, J., 2006. Incentives and prosocial behavior. *The American Economic Review* 96 (5), 1652–1678.
- Berg, J.E., Nelson, F.D., Rietz, T.A., 2008. Prediction market accuracy in the long run. *International Journal of Forecasting* 24 (2), 285–300.
- Bergemann, D., Bonatti, A., 2015. Selling cookies. *American Economic Journal: Microeconomics* 7 (3), 259–294.
- Bergemann, D., Bonatti, A., 2019. Markets for information: an introduction. *Annual Review of Economics* 11, 85–107.
- Bergemann, D., Bonatti, A., Gan, T., 2021a. The Economics of Social Data. *The Rand Journal of Economics*. Forthcoming.
- Bergemann, D., Bonatti, A., Smolin, A., 2018. The design and price of information. *The American Economic Review* 108 (1), 1–48.
- Bergemann, D., Brooks, B., Morris, S., 2015. The limits of price discrimination. *The American Economic Review* 105 (3), 921–957.
- Bergemann, D., Brooks, B., Morris, S., 2017. First-price auctions with general information structures: implications for bidding and revenue. *Econometrica* 85 (1), 107–143.
- Bergemann, D., Brooks, B., Morris, S., 2019. Revenue guarantee equivalence. *The American Economic Review* 109 (5), 1911–1929.
- Bergemann, D., Brooks, B., Morris, S., 2020. Countering the winner’s curse: optimal auction design in a common value model. *Theoretical Economics* 15, 1399–1434.
- Bergemann, D., Cai, Y., Velez, G., Zhao, M., 2021b. Is Selling Complete Information (Approximately) Optimal? Discussion paper. Yale University.
- Bergemann, D., Heumann, T., Morris, S., 2021c. Selling Impressions: Efficiency vs. Competition. Cowles Foundation Discussion Paper 2291.
- Bergemann, D., Morris, S., 2013. Robust predictions in games with incomplete information. *Econometrica* 81 (4), 1251–1308.
- Bergemann, D., Morris, S., 2019. Information design: a unified perspective. *Journal of Economic Literature* 57 (1), 44–95.
- Bergemann, D., Ozmen, D., 2006. Optimal pricing with recommender systems. In: *Proceedings of the 7th ACM Conference on Electronic Commerce*, pp. 43–51.
- Bergemann, D., Pesendorfer, M., 2007. Information structures in optimal auctions. *Journal of Economic Theory* 137 (1), 580–609.
- Bergemann, D., Said, M., 2011. Dynamic auctions: a survey. In: Cochran, J. (Ed.), *Wiley Encyclopedia of Operations Research and Management Science*, No. 1757. Wiley, New York, pp. 1511–1522.

- Bergemann, D., Välimäki, J., 2019. Dynamic mechanism design: an introduction. *Journal of Economic Literature* 57, 235–274.
- Bergemann, D., Yeh, E., Zhang, J., 2021d. Nonlinear pricing with finite information. *Games and Economic Behavior* 130, 62–84.
- Bertomeu, J., 2021. The Editor's Problem. Discussion paper. Washington University in St. Luis.
- Besley, T., Ghatak, M., 2005. Competition and incentives with motivated agents. *The American Economic Review* 95 (3), 616–636.
- Bimpikis, K., Crapis, D., Tahbaz-Salehi, A., 2019. Information sale and competition. *Management Science* 65 (6), 2646–2664.
- Blackwell, D., 1951. Comparison of experiments. In: *Proceedings of the Second Berkeley Symposium in Mathematical Statistics and Probability*. University of California Press, Berkeley, pp. 93–102.
- Blackwell, D., 1953. Equivalent comparison of experiments. *Annals of Mathematics and Statistics* 24 (2), 265–272.
- Blackwell, D., Girshick, M., 1954. *Theory of Games and Statistical Decisions*. John Wiley.
- Bobtcheff, C., Bolte, J., Mariotti, T., 2017. Researcher's dilemma. *The Review of Economic Studies* 84 (3), 969–1014.
- Bobtcheff, C., Levy, R., Mariotti, T., 2021. Negative Results in Science: Blessing or (Winner's) Curse? Discussion paper. University of Toulouse and HEC, Paris.
- Boleslavsky, R., Kelly, D.L., Taylor, C.R., 2017. Selloffs, bailouts, and feedback: can asset markets inform policy? *Journal of Economic Theory* 169, 294–343.
- Bonatti, A., Cisternas, G., 2020. Consumer scores and price discrimination. *The Review of Economic Studies* 87 (2), 750–791.
- Bond, P., Goldstein, I., Prescott, E.S., 2010. Market-based corrective actions. *The Review of Financial Studies* 23 (2), 781–820.
- Bonferroni, C., 1936. *Teoria Statistica Delle Classi e Calcolo Delle Probabilita*. Pubblicazioni del R Istituto Superiore di Scienze Economiche e Commerciali di Firenze 8, 3–62.
- Borel, É., 1938. Sur Le Pari Mutuel. *Comptes Rendus Hebdomadaires des Séances de L'Académie des Sciences* 207 (3), 197–200.
- Börgers, T., Hernando-Veciana, A., Kräbmer, D., 2013. When are signals complements or substitutes. *Journal of Economic Theory* 148, 165–195.
- Brier, G.W., 1950. Verification of forecasts expressed in terms of probability. *Monthly Weather Review* 78 (1), 1–3.
- Brock, W.A., Durlauf, S.N., 1999. A formal model of theory choice in science. *Economic Theory* 14 (1), 113–130.
- Brodeur, A., Lé, M., Sangnier, M., Zylberberg, Y., 2016. Star wars: the empirics strike back. *American Economic Journal: Applied Economics* 8 (1), 1–32.
- Calzolari, G., Pavan, A., 2006. Monopoly with resale. *The Rand Journal of Economics* 37 (2), 362–375.
- Camerer, C.F., 1998. Can asset markets be manipulated? A field experiment with racetrack betting. *Journal of Political Economy* 106 (3), 457–482.
- Camerer, C.F., Dreber, A., Forsell, E., Ho, T.-H., Huber, J., Johannesson, M., Kirchler, M., Almenberg, J., Altmejd, A., Chan, T., et al., 2016. Evaluating replicability of laboratory experiments in economics. *Science* 351 (6280), 1433–1436.
- Camerer, C.F., Dreber, A., Holzmeister, F., Ho, T.-H., Huber, J., Johannesson, M., Kirchler, M., Nave, G., Nosek, B.A., Pfeiffer, T., et al., 2018. Evaluating the replicability of social science experiments in nature and science between 2010 and 2015. *Nature Human Behaviour* 2 (9), 637–644.

- Carnehl, C., Schneider, J., 2021. A Quest for Knowledge. Discussion paper. Bocconi University and Universitat Carlos III de Madrid.
- Carroll, G., Segal, I., 2019. Robustly optimal auctions with unknown resale opportunities. *The Review of Economic Studies* 86 (4), 1527–1555.
- Chassang, S., Miquel, P.I., Snowberg, E., 2012. Selective trials: a principal-agent approach to randomized controlled experiments. *The American Economic Review* 102 (4), 1279–1309.
- Che, Y.-K., Hörner, J., 2017. Recommender systems as mechanisms for social learning. *The Quarterly Journal of Economics* 133 (2), 871–925.
- Chen, Y.C., Mueller-Frank, M., Pai, M.M., 2021. The wisdom of the crowd and higher-order beliefs. arXiv preprint. arXiv:2102.02666.
- Chiappori, P.-A., Salanié, B., Salanié, F., Gandhi, A., 2019. From aggregate betting data to individual risk preferences. *Econometrica* 87 (1), 1–36.
- Chiou, L., Tucker, C., 2017. Search Engines and Data Retention: Implications for Privacy and Antitrust. NBER Working Paper, 23815.
- Choi, H., Mela, C., Balseiro, S., Leary, A., 2020. Online display advertising markets: a literature review and future directions. *Information Systems Research* 31, 556–575.
- Choi, J., Jeon, D., Kim, B., 2019. Privacy and personal data collection with information externalities. *Journal of Public Economics* 173, 113–124.
- Christensen, G., Miguel, E., 2018. Transparency, reproducibility, and the credibility of economics research. *Journal of Economic Literature* 56 (3), 920–980.
- Clarke, R., 1983. Collusion and the incentives for information sharing. *Bell Journal of Economics* 14, 383–394.
- Claussen, J., Peukert, C., Sen, A., 2019. The editor vs. the algorithm: returns to data and externalities in online news. SSRN. <https://ssrn.com/abstract=3479854>.
- Cohen, W.M., Sauermann, H., Stephan, P., 2020. Not in the job description: the commercial activities of academic scientists and engineers. *Management Science* 66 (9), 4108–4117.
- Cole, S., 1983. The hierarchy of the sciences? *American Journal of Sociology* 89 (1), 111–139.
- Coles, P.A., Lakhani, K.R., McAfee, A., 2007. Prediction Markets at Google. Harvard Business School, Boston, MA.
- Committee on the Judiciary, 2020. Investigation of Competition in Digital Markets. Discussion paper. United States House of Representatives.
- Conitzer, V., Taylor, C.R., Wagman, L., 2012. Hide and seek: costly consumer privacy in a market with repeat purchases. *Marketing Science* 31 (2), 277–292.
- Cotton, C., 2013. Submission fees and response times in academic publishing. *The American Economic Review* 103 (1), 501–509.
- Council of Economic Advisers, 2015. Big Data and Differential Pricing. Executive Office of the President of the United States.
- Cournot, A.A., 1843. Exposition de la Théorie des Chances et des Probabilités. L. Hachette.
- Cowgill, B., Zitzewitz, E., 2015. Corporate prediction markets: evidence from Google, Ford, and Firm X. *The Review of Economic Studies* 82 (4), 1309–1341.
- Craig, A., Vierø, M.-L., 2013. Academia or the private sector? Sorting of agents into institutions and an outside sector. *The B.E. Journals in Theoretical Economics* 13 (1), 303–345.
- Crawford, V.P., Sobel, J., 1982. Strategic information transmission. *Econometrica* 50 (6), 1431–1451.
- Crèmer, J., de Montjoye, Y.-A., Schweitzer, H., 2019. Competition Policy for the Digital Era—Final Report. Discussion paper. European Commission.
- Crosland, M., 2002. Science Under Control: The French Academy of Sciences 1795-1914. Cambridge University Press.

- Cvitanić, J., Prelec, D., Riley, B., Tereick, B., 2019. Honesty via choice-matching. *American Economic Review: Insights* 1 (2), 179–192.
- Dahleh, M., 2018. Why the Data Marketplace of the Future Will Sell Insights, Not Data. Discussion paper. MIT.
- Dai, M., Jia, Y., Kou, S., 2021. The wisdom of the crowd and prediction markets. *Journal of Econometrics* 222 (1), 561–578.
- Dana, J., Atanasov, P., Tetlock, P., Mellers, B., 2019. Are markets more accurate than polls? The surprising informational value of ‘just asking’. *Judgment and Decision Making* 14 (2), 135–147.
- Dasgupta, P., David, P.A., 1987. Information disclosure and the economics of science and technology. In: *Arrow and the Ascent of Modern Economic Theory*. Springer, pp. 519–542.
- Dasgupta, P., David, P.A., 1994. Toward a new economics of science. *Research Policy* 23 (5), 487–521.
- Daskalakis, C., Deckelbaum, A., Tzamos, C., 2017. Strong duality for a multiple good monopolist. *Econometrica* 85 (3), 735–767.
- David, P.A., 1998. Common agency contracting and the emergence of ‘open science’ institutions. *The American Economic Review* 88 (2), 15–21.
- de Corniere, A., de Nijs, R., 2016. Online advertising and privacy. *The Rand Journal of Economics* 47 (1), 48–72.
- De Finetti, B., 1962. The scientist speculates: an anthology of partly-baked ideas. In: Boskin, M. (Ed.), *It Make Sense to Speak of ‘Good Probability Appraisers’*. Wiley, New York, pp. 257–364.
- De Solla Price, D.J., 1965. Networks of scientific papers. *Science* 149 (3683), 510–515.
- Deb, R., Pai, M.M., Said, M., 2018. Evaluating strategic forecasters. *The American Economic Review* 108 (10), 3057–3103.
- Deck, C., Lin, S., Porter, D., 2013. Affecting policy by manipulating prediction markets: experimental evidence. *Journal of Economic Behavior & Organization* 85, 48–62.
- Dekel, E., Fudenberg, D., Levine, D., 2004. Learning to play Bayesian games. *Games and Economic Behavior* 46 (2), 282–303.
- DellaVigna, S., Otis, N., Vivalt, E., 2020. Forecasting the Results of Experiments: Piloting an Elicitation Strategy. *w110*. pp. 75–79.
- DellaVigna, S., Pope, D., Vivalt, E., 2019. Predict science to improve science. *Science* 366 (6464), 428–429.
- Demange, G., 2014. A ranking method based on handicaps. *Theoretical Economics* 9 (3), 915–942.
- Di Tillio, A., Ottaviani, M., Sørensen, P.N., 2017. Persuasion bias in science: can economics help? *The Economic Journal* 127 (605), F266–F304.
- Di Tillio, A., Ottaviani, M., Sørensen, P.N., 2021. Strategic sample selection. *Econometrica* 89 (2), 911–953.
- Dianat, A., Siemroth, C., 2020. Improving decisions with market information: an experiment on corporate prediction markets. *Experimental Economics* 24, 143–176.
- Digital Competition Expert Panel, 2019. *Unlocking Digital Competition*. Discussion paper.
- Drakopoulos, K., Randhawa, R., 2020. Why Perfect Tests May Not Be Worth Waiting for: Information as a Commodity. Discussion paper. University of Southern California.
- Dranove, D., Jin, G., 2010. Quality disclosure and certification: theory and practice. *Journal of Economic Literature* 48 (4), 935–963.
- Dreber, A., Pfeiffer, T., Almenberg, J., Isaksson, S., Wilson, B., Chen, Y., Nosek, B.A., Johannesson, M., 2015. Using prediction markets to estimate the reproducibility of scientific research. *Proceedings of the National Academy of Sciences* 112 (50), 15343–15347.

- Dube, J., Fang, Z., Fong, N., Luo, X., 2017. Competitive price targeting with smartphone coupons. *Marketing Science* 36 (6), 944–975.
- Dube, J., Misra, S., 2017. Scalable Price Targeting. Discussion paper. University of Chicago Booth School of Business.
- Dworczak, P., 2016. Mechanism Design with Aftermarkets: Cutoff Mechanisms. Discussion paper. Stanford University.
- Echenique, F., He, K., 2021. Screening p-hackers: dissemination noise as bait. arXiv preprint. arXiv:2103.09164.
- Ellingsen, T., Johannesson, M., 2008. Pride and prejudice: the human side of incentive theory. *The American Economic Review* 98 (3), 990–1008.
- Ellison, G., 2002. The slowdown of the economics publishing process. *Journal of Political Economy* 110 (5), 947–993.
- Ellison, G., 2013. How does the market use citation data? The Hirsch index in economics. *American Economic Journal: Applied Economics* 5 (3), 63–90.
- Ely, J., Galeotti, A., Jann, O., Steiner, J., 2021. Optimal test allocation. *Journal of Economic Theory*, 105236.
- Emek, Y., Feldman, M., Gamzu, I., Leme, R.P., Tennenholtz, M., 2012. Signaling schemes for revenue maximization. In: *ACM EC* 12.
- Eső, P., Szentes, B., 2007a. Optimal information disclosure in auctions. *The Review of Economic Studies* 74 (3), 705–731.
- Eső, P., Szentes, B., 2007b. The price of advice. *The Rand Journal of Economics* 38 (4), 863–880.
- Fainmesser, I., Galeotti, A., Momot, R., 2020. Digital Privacy. Discussion paper. Johns Hopkins University.
- Fama, E., 1970. Efficient capital markets: a review of theory and empirical work. *The Journal of Finance* 25 (2), 383–417.
- Fanelli, D., 2009. How many scientists fabricate and falsify research? A systematic review and meta-analysis of survey data. *PLoS ONE* 4 (5), e5738.
- Federal Trade Commission, 2014. Data brokers: a call for transparency and accountability. <https://www.ftc.gov/reports/data-brokers-call-transparency-accountability-report-federal-trade-commission-may-2014>.
- Feldman, M., Lucier, B., Syrgkanis, V., 2013. Limits of efficiency in sequential auctions. In: *WINE 2013*. In: *Lecture Notes in Computer Science*. Springer Verlag, pp. 160–173.
- Fisher, R.A., 1936. Has Mendel’s work been rediscovered? *Annals of Science* 1 (2), 115–137.
- Forsell, E., Viganola, D., Pfeiffer, T., Almenberg, J., Wilson, B., Chen, Y., Nosek, B.A., Johannesson, M., Dreber, A., 2019. Predicting replication outcomes in the many labs 2 study. *Journal of Economic Psychology* 75, 102117.
- Forsythe, R., Lundholm, R., 1990. Information aggregation in an experimental market. *Econometrica* 58 (2), 309–347.
- Forsythe, R., Nelson, F., Neumann, G.R., Wright, J., 1992. Anatomy of an experimental political stock market. *The American Economic Review* 82 (5), 1142–1161.
- Fortunato, S., Bergstrom, C.T., Börner, K., Evans, J.A., Helbing, D., Milojević, S., Petersen, A.M., Radicchi, F., Sinatra, R., Uzzi, B., et al., 2018. Science of science. *Science* 359 (6379), eaao0185.
- Franco, A., Malhotra, N., Simonovits, G., 2014. Publication bias in the social sciences: unlocking the file drawer. *Science* 345 (6203), 1502–1505.
- Frankel, A., Kamenica, E., 2019. Quantifying information and uncertainty. *The American Economic Review* 109 (10), 3650–3680.

- Frankel, A., Kasy, M., 2021. Which findings should be published? *American Economic Journal: Microeconomics*. Forthcoming.
- Franzoni, C., Sauermann, H., 2014. Crowd science: the organization of scientific research in open collaborative projects. *Research Policy* 43 (1), 1–20.
- Freixas, X., Guesnerie, R., Tirole, J., 1985. Planning under incomplete information and the ratchet effect. *The Review of Economic Studies* 52, 173–191.
- Fudenberg, D., Villas-Boas, J.M., 2006. Behavior-based price discrimination and customer recognition. In: *Handbook on Economics and Information Systems*.
- Fudenberg, D., Villas-Boas, M., 2012. Price discrimination in the digital economy. In: Peitz, M., Waldfogel, J. (Eds.), *Oxford Handbook of the Digital Economy*. Oxford University Press, pp. 254–272.
- Gaba, A., Popescu, D.G., Chen, Z., 2019. Assessing uncertainty from point forecasts. *Management Science* 65 (1), 90–106.
- Gall, T., Ioannidis, J.P., Maniatis, Z., 2017. The credibility crisis in research: can economics tools help? *PLoS Biology* 15 (4), e2001846.
- Galperti, S., Levkun, A., Perego, J., 2021. The Value of Data. Discussion paper. UC San Diego and Columbia University.
- Galton, F., 1907a. One vote, one value. *Nature* 75 (1948), 414.
- Galton, F., 1907b. Vox Populi. *Nature* 75 (1949), 450–451.
- Gandhi, A., Serrano-Padial, R., 2015. Does belief heterogeneity explain asset prices: the case of the longshot bias. *The Review of Economic Studies* 82 (1), 156–186.
- Garfield, E., 1955. Citation indexes for science. *Science* 122 (3159), 108–111.
- Garfield, E., 1972. Citation analysis as a tool in journal evaluation. *Science* 178 (4060), 471–479.
- Gauriot, R., Page, L., 2020. How Market Prices React to Information: Evidence from a Natural Experiment.
- Geanakoplos, J., Polemarchakis, H., 1982. We can't disagree forever. *Journal of Economic Theory* 28 (1), 192–200.
- Gentzkow, M., Shapiro, J.M., 2006. Media bias and reputation. *Journal of Political Economy* 114 (2), 280–316.
- Ghosh, A., Nazerzadeh, H., Sundararajan, M., 2009. Computing Optimal Bundles for Sponsored Search. Discussion paper. Yahoo Research.
- Gillen, B.J., Plott, C.R., Shum, M., 2017. A pari-mutuel-like mechanism for information aggregation: a field test inside intel. *Journal of Political Economy* 125 (4), 1075–1099.
- Gjerstad, S., Hall, M., 2005. Risk Aversion, Beliefs, and Prediction Market Equilibrium. *Economic Science Laboratory*, University of Arizona.
- Gneiting, T., Raftery, A.E., 2007. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association* 102 (477), 359–378.
- Godfrey-Smith, P., 2009. *Theory and Reality: An Introduction to the Philosophy of Science*. University of Chicago Press.
- Goldfarb, A., Tucker, C., 2011. Search engine advertising: channel substitution when pricing ads to context. *Management Science* 57 (3), 458–470.
- Goldfarb, A., Tucker, C., 2019. Digital economics. *Journal of Economic Literature* 57 (1), 3–43.
- Goldstein, I., Guembel, A., 2008. Manipulation and the allocational role of prices. *The Review of Economic Studies* 75 (1), 133–164.
- Good, I., 1952. Rational decisions. *Journal of the Royal Statistical Society, Series B* 14 (1), 107–114.

- Gordon, M., Viganola, D., Bishop, M., Chen, Y., Dreber, A., Goldfedder, B., Holzmeister, F., Johannesson, M., Liu, Y., Twardy, C., et al., 2020. Are replication rates the same across academic fields? Community forecasts from the darpa score programme. *Royal Society Open Science* 7.
- Gradwohl, R., 2017. Information sharing and privacy in networks. In: *ACM-EC (Economics and Computation)* 2017.
- Green, E., Lee, H., Rothschild, D.M., 2020. The Favorite-Longshot Midas. Available at SSRN 3271248.
- Griffith, R.M., 1949. Odds adjustments by American horse-race bettors. *The American Journal of Psychology* 62 (2), 290–294.
- Grossman, S., 1976. On the efficiency of competitive stock markets where trades have diverse information. *The Journal of Finance* 31 (2), 573–585.
- Grossman, S., Stiglitz, J., 1980. On the impossibility of informationally efficient markets. *The American Economic Review* 70, 393–408.
- Grundy, P., Healy, M., Rees, D., 1956. Economic choice of the amount of experimentation. *Journal of the Royal Statistical Society* 18, 32–49.
- Guo, Y., Li, H., Shi, X., 2020. Optimal Discriminatory Disclosure. Discussion paper. Northwestern, UBC, and Toronto.
- Hanson, R., 2003. Combinatorial information market design. *Information Systems Frontiers* 5 (1), 107–119.
- Hanson, R., 2007. Logarithmic markets scoring rules for modular combinatorial information aggregation. *The Journal of Prediction Markets* 1 (1), 3–15.
- Hanson, R., Oprea, R., Porter, D., 2006. Information aggregation and manipulation in an experimental market. *Journal of Economic Behavior & Organization* 60 (4), 449–459.
- Hargens, L.L., 1975. Patterns of Scientific Research: A Comparative Analysis of Research in Three Scientific Fields. American Sociological Association.
- Hargens, L.L., 1988. Scholarly consensus and journal rejection rates. *American Sociological Review* 53 (1), 139–151.
- Harris, M., Raviv, A., 1993. Differences of opinion make a horse race. *The Review of Financial Studies* 6 (3), 473–506.
- Hart, O., Tirole, J., 1988. Contract renegotiation and coasian dynamics. *The Review of Economic Studies* 55, 509–540.
- Hartline, J., Johnsen, A., Nekipelov, D., Zoeter, O., 2019. Dashboard mechanisms for online marketplaces. In: *Proceedings of the 20th ACM Conference on Economics and Computation*.
- Hayek, F.A., 1945. The use of knowledge in society. *The American Economic Review* 35 (4), 519–530.
- He, D., Kannan, A., Liu, T., McAfee, R., Qin, T., Rao, J., 2017. Scale Effects in Web Search. Discussion paper.
- Henry, E., Ottaviani, M., 2019. Research and the approval process: the organization of persuasion. *The American Economic Review* 109 (3), 911–955.
- Henry, E., Ponce, C., 2011. Waiting to imitate: on the dynamic pricing of knowledge. *Journal of Political Economy* 119, 959–981.
- Herresthal, C., 2021. Hidden Testing and Selective Disclosure of Evidence. Discussion paper. University of Bonn.
- Hidir, S., Vellodi, N., 2018. Personalization, Discrimination and Information Revelation. Discussion paper.
- Hill, R., Stein, C., 2021a. Race to the Bottom: Competition and Quality in Science. Discussion paper. MIT.

- Hill, R., Stein, C., 2021b. Scooped! Estimating Rewards for Priority in Science. Discussion paper. MIT.
- Hirshleifer, J., 1971. The private and social value of information and the reward to inventive activity. *The American Economic Review* 61 (4), 561–574.
- Hoffmann, F., Inderst, R., Ottaviani, M., 2020. Persuasion through selective disclosure: implications for marketing, campaigning, and privacy regulation. *Management Science* 66 (11), 4958–4979.
- Holmstrom, B., Milgrom, P., 1994. The firm as an incentive system. *The American Economic Review* 84 (4), 972–991.
- Hong, H., Stein, J.C., 2007. Disagreement and the stock market. *The Journal of Economic Perspectives* 21 (2), 109–128.
- Hörner, J., Lambert, N., 2021. Motivational ratings. *The Review of Economic Studies* 88 (4), 1892–1935.
- Hull, D.L., 1988. *Science as a Process: An Evolutionary Account of the Social and Conceptual Development of Science*. University of Chicago Press.
- Hull, D.L., 1997. What's wrong with invisible-hand explanations? *Philosophy of Science* 64, S117–S126.
- Ichihashi, S., 2020a. The Economics of Data Externalities. Discussion paper. Bank of Canada.
- Ichihashi, S., 2020b. Online privacy and information disclosure by consumers. *The American Economic Review* 110 (2), 569–595.
- Inderst, R., Ottaviani, M., 2012a. Competition through commissions and kickbacks. *The American Economic Review* 102 (2), 911–955.
- Inderst, R., Ottaviani, M., 2012b. Financial advice. *Journal of Economic Literature* 50 (2), 494–512.
- Ioannidis, J.P., 2005. Why most published research findings are false. *PLoS Medicine* 2 (8), e124.
- Iordanou, C., Soriente, C., Sirivianos, M., Laoutaris, N., 2017. Who is fiddling with prices? Building and deploying a watchdog service for e-commerce. In: *Proceedings of the Conference of the ACM Special Interest Group on Data Communication*, pp. 376–389.
- Johnson, G., Shriver, S., Goldberg, S., 2020. Privacy and Market Concentration: Intended and Unintended Consequences of the GDPR. Discussion paper.
- Jones, C.I., Tonetti, C., 2020. Nonrivalry and the economics of data. *The American Economic Review* 110 (9), 2819–2858.
- Jullien, B., Lefouili, Y., Riordan, M., 2020. Privacy Protection, Security, and Consumer Retention. Discussion paper. Toulouse School of Economics.
- Jullien, B., Salanié, B., 2000. Estimating preferences under risk: the case of racetrack bettors. *Journal of Political Economy* 108 (3), 503–530.
- Jullien, B., Salanié, B., 2008. Empirical evidence on the preferences of racetrack bettors. In: *Handbook of Sport and Lottery Markets*. Elsevier.
- Kajii, A., Watanabe, T., 2017. Favorite–longshot bias in pari-mutuel betting: an evolutionary explanation. *Journal of Economic Behavior & Organization* 140, 56–69.
- Kamenica, E., Gentzkow, M., 2011. Bayesian persuasion. *The American Economic Review* 101 (6), 2590–2615.
- Kandel, E., Zilberfarb, B.-Z., 1999. Differential interpretation of information in inflation forecasts. *Review of Economics and Statistics* 81 (2), 217–226.
- Kawamura, K., 2013. Eliciting information from a large population. *Journal of Public Economics* 103, 44–54.
- Keane, M.P., Runkle, D.E., 1998. Are financial analysts' forecasts of corporate profits rational? *Journal of Political Economy* 106 (4), 768–805.

- Kitcher, P., 1995. *The Advancement of Science: Science Without Legend, Objectivity Without Illusions*. Oxford University Press.
- Kleinberg, J., Oren, S., 2011. Mechanisms for (mis)allocating scientific credit. In: *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing*, pp. 529–538.
- Kremer, I., Mansour, Y., Perry, M., 2014. Implementing the “wisdom of the crowd”. *Journal of Political Economy* 122 (5), 988–1012.
- Kreps, D., 1977. A note on ‘fulfilled expectations’ equilibria. *Journal of Economic Theory* 14, 32–43.
- Kreps, D.M., 1997. Intrinsic motivation and extrinsic incentives. *The American Economic Review* 87 (2), 359–364.
- Krieger, J., Myers, K., Stern, A.D., 2021. Editors are Gatekeepers of Science, but Individual Editors Don’t Matter Much. Discussion paper. Harvard University.
- Kuhn, T.S., 1962. *The Structure of Scientific Revolutions*. Princeton University Press.
- Lacetera, N., Zirulia, L., 2011. The economics of scientific misconduct. *Journal of Law, Economics, & Organization* 27 (3), 568–603.
- Laffont, J., 1989. *The Economics of Uncertainty and Information*. MIT Press, Cambridge.
- Lambert, N.S., 2011. Elicitation and Evaluation of Statistical Forecasts. Discussion paper. Stanford University.
- Lambert, N.S., Langford, J., Vaughan, J.W., Chen, Y., Reeves, D.M., Shoham, Y., Pennock, D.M., 2015. An axiomatic characterization of wagering mechanisms. *Journal of Economic Theory* 156, 389–416.
- Lee, G., Wright, J., 2021. Recommender systems and the value of user data. National University of Singapore Working Paper. https://836d83c2-c629-40a4-9ca7-b1bf324d720d.filesusr.com/ugd/c6ffe6_e391febde31845a08bc1ef6969182398.pdf.
- Leslie, D., 2005. Are delays in academic publishing necessary? *The American Economic Review* 95 (1), 407–413.
- Levin, J., Milgrom, P., 2010. Online advertising: heterogeneity and conflation in market design. *The American Economic Review: Papers and Proceedings* 100 (2), 603–607.
- Levy, G., 2004. Anti-herding and strategic consultation. *European Economic Review* 48 (3), 503–525.
- Li, D., 2017. Expertise versus bias in evaluation: evidence from the NIH. *American Economic Journal: Applied Economics* 9 (2), 60–92.
- Li, H., Shi, X., 2017. Discriminatory information disclosure. *The American Economic Review* 107, 3363–3385.
- Liang, A., Madsen, E., 2020. Data and Incentives. Discussion paper. Northwestern University.
- Lichtendahl Jr, K.C., Grushka-Cockayne, Y., Pfeifer, P.E., 2013. The wisdom of competitive crowds. *Operations Research* 61 (6), 1383–1398.
- Liebowitz, S.J., Palmer, J.P., 1984. Assessing the relative impacts of economics journals. *Journal of Economic Literature* 22 (1), 77–88.
- Lieli, R.P., Nieto-Barthaburu, A., 2020. On the possibility of informative equilibria in futures markets with feedback. *Journal of the European Economic Association* 18 (3), 1521–1552.
- Lizzeri, A., 1999. Information revelation and certification intermediaries. *The Rand Journal of Economics* 30 (2), 214–231.
- Luenberger, D.G., 1995. *Microeconomic Theory*. McGraw-Hill.
- Luo, L., Toubia, O., 2015. Improving online idea generation platforms and customizing the task structure on the basis of consumers’ domain-specific knowledge. *Journal of Marketing* 79 (5), 100–114.
- Mačkowiak, B., Matějka, F., Wiederholt, M., forthcoming. Rational Inattention: a Review. *Journal of Economic Literature*.

- Manski, C.F., 2006. Interpreting the predictions of prediction markets. *Economics Letters* 91 (3), 425–429.
- Mansour, Y., Slivkins, A., Syrgkanis, V., 2015. Bayesian incentive compatible bandit exploration. In: ACM Conference on Electronic Commerce.
- Marinovic, I., Ottaviani, M., Sorensen, P., 2013. Forecasters' objectives and strategies. In: *Handbook of Economic Forecasting*, Vol. 2. Elsevier, pp. 690–720.
- Marschak, J., 1968. Economics of inquiring, communicating, deciding. *The American Economic Review: Papers and Proceedings* 58 (2), 1–18.
- Maysnar, J., 1983. On divergence of opinion and imperfections in capital markets. *The American Economic Review* 73 (1), 114–128.
- McAfee, P., 2002. Coarse matching. *Econometrica* 70, 2025–2034.
- McCarthy, J., 1956. Measures of the value of information. *Proceedings of the National Academy of Sciences* 42 (9), 654–655.
- McGraw, A.P., Tetlock, P.E., 2005. Taboo trade-offs, relational framing, and the acceptability of exchanges. *Journal of Consumer Psychology* 15 (1), 2–15.
- Meirowitz, A., Tucker, J.A., 2004. Learning from terrorism markets. *Perspective on Politics* 2 (2), 331–336.
- Meloso, D., Nunnari, S., Ottaviani, M., 2018. Looking into Crystal Balls: a Laboratory Experiment on Reputational Cheap Talk. Discussion paper.
- Merton, R.K., 1942. A note on science and democracy. *Journal of Legal and Political Sociology* 1, 115–126.
- Merton, R.K., 1957. Priorities in scientific discovery: a chapter in the sociology of science. *American Sociological Review* 22 (6), 635–659.
- Merton, R.K., 1988. The Matthew effect in science, II: cumulative advantage and the symbolism of intellectual property. *Isis* 79 (4), 606–623.
- Milgrom, P., Stokey, N., 1982. Information, trade and common knowledge. *Journal of Economic Theory* 26 (1), 17–27.
- Miller, E.M., 1977. Risk, uncertainty, and divergence of opinion. *The Journal of Finance* 32 (4), 1151–1168.
- Miller, N., Resnick, P., Zeckhauser, R., 2005. Eliciting informative feedback: the peer-prediction method. *Management Science* 51 (9), 1359–1373.
- Moed, H.F., 2006. *Citation Analysis in Research Evaluation*. Springer.
- Morgan, J., Stocken, P.C., 2008. Information aggregation in polls. *The American Economic Review* 98 (3), 864–896.
- Morris, S., 1994. Trade with heterogeneous prior beliefs and asymmetric information. *Econometrica* 62 (6), 1327–1347.
- Morris, S., 1995. The common prior assumption in economic theory. *Economics and Philosophy* 11 (2), 227–253.
- Morris, S., 2001. Political correctness. *Journal of Political Economy* 109 (2), 231–265.
- Morris, S., Shin, H., 2002. Social value of public information. *The American Economic Review* 92 (5), 1521–1534.
- Muller-Ippen, M., 2019. Gatekeeper Competition. Discussion paper. University of Notre Dame.
- Myerson, R., 1981. Optimal auction design. *Mathematics of Operations Research* 6 (1), 58–73.
- Narita, Y., 2021. Incorporating ethics and welfare into randomized experiments. *Proceedings of the National Academy of Sciences* 118 (1), e200874011.
- New York Times, 2018. As Facebook Raised a Privacy Wall, It Carved an Opening for Tech Giants.
- Ngangoué, K., Weizsäcker, G., 2021. Learning from unrealized versus realized prices. *American Economic Journal: Microeconomics* 13, 174–201.

- Novshek, W., Sonnenschein, H., 1982. Fulfilled expectations Cournot duopoly with information acquisition and release. *Bell Journal of Economics* 13, 214–218.
- Oostrom, T., 2021. Funding of Clinical Trials and Reported Drug Efficacy. Discussion paper. Ohio State University.
- Open Science Collaboration, 2015. Estimating the reproducibility of psychological science. *Science* 349 (6251), aac4716.
- Orr, G., 2014. Betting on elections: history, law and policy. *Federal Law Review* 42 (2), 309–331.
- Osband, K., 1989. Optimal forecasting incentives. *Journal of Political Economy* 97 (5), 1091–1112.
- Ostrovsky, M., 2012. Information aggregation in dynamic markets with strategic traders. *Econometrica* 80 (6), 2595–2647.
- Othman, A., Sandholm, T., 2013. The Gates Hillman prediction market. *Review of Economic Design* 17 (2), 95–128.
- Ottaviani, M., 2009. The design of idea markets: an economist's perspective. *The Journal of Prediction Markets* 3 (1), 41–44.
- Ottaviani, M., 2020. Grantmaking. CEPR Discussion Paper 15389.
- Ottaviani, M., Sørensen, P., 2001. Information aggregation in debate: who should speak first? *Journal of Public Economics* 81 (3), 393–421.
- Ottaviani, M., Sørensen, P.N., 2005. Parimutuel Versus Fixed-Odds Markets. Discussion paper. London Business School and University of Copenhagen.
- Ottaviani, M., Sørensen, P.N., 2006a. Reputational cheap talk. *The Rand Journal of Economics* 37 (1), 155–175.
- Ottaviani, M., Sørensen, P.N., 2006b. The strategy of professional forecasting. *Journal of Financial Economics* 81 (2), 441–466.
- Ottaviani, M., Sørensen, P.N., 2006c. The Timing of Parimutuel Bets. Discussion paper. London Business School and University of Copenhagen.
- Ottaviani, M., Sørensen, P.N., 2007. Outcome manipulation in corporate prediction markets. *Journal of the European Economic Association* 5 (2–3), 554–563.
- Ottaviani, M., Sørensen, P.N., 2008. The favorite-longshot bias: an overview of the main explanations. In: *Handbook of Sports and Lottery Markets*. Elsevier, pp. 83–101.
- Ottaviani, M., Sørensen, P.N., 2009. Surprised by the parimutuel odds? *The American Economic Review* 99 (5), 2129–2134.
- Ottaviani, M., Sørensen, P.N., 2010. Noise, information, and the favorite-longshot bias in parimutuel predictions. *American Economic Journal: Microeconomics* 2 (1), 58–85.
- Ottaviani, M., Sørensen, P.N., 2015. Price reaction to information with heterogeneous beliefs and wealth effects: underreaction, momentum, and reversal. *The American Economic Review* 105 (1), 1–34.
- Page, L., Clemen, R.T., 2013. Do prediction markets produce well-calibrated probability forecasts? *The Economic Journal* 123 (568), 491–513.
- Page, L., Siemroth, C., 2017. An experimental analysis of information acquisition in prediction markets. *Games and Economic Behavior* 101, 354–378.
- Page, L., Siemroth, C., 2020. How much information is incorporated into financial asset prices? Experimental evidence. *The Review of Financial Studies*.
- Palacios-Huerta, I., Volij, O., 2004. The measurement of intellectual influence. *Econometrica* 72 (3), 963–977.
- Palacios-Huerta, I., Volij, O., 2014. Axiomatic measures of intellectual influence. *International Journal of Industrial Organization* 34, 85–90.

- Palfrey, T.R., Srivastava, S., 1986. Private information in large economies. *Journal of Economic Theory* 39 (1), 34–58.
- Palley, A.B., Soll, J.B., 2019. Extracting the wisdom of crowds when information is shared. *Management Science* 65 (5), 2291–2309.
- Pavlov, G., 2011. A property of solutions to linear monopoly problems. *The B.E. Journals in Theoretical Economics* 11.
- Perry, M., Reny, P.J., 2016. How to count citations if you must. *The American Economic Review* 106 (9), 2722–2741.
- Pigou, A., 1920. *The Economics of Welfare*. Macmillan, London.
- Plott, C.R., Chen, K.-Y., 2002. Information Aggregation Mechanisms: Concept, Design and Implementation for a Sales Forecasting Problem. Discussion paper. California Institute of Technology.
- Plott, C.R., Sunder, S., 1982. Efficiency of experimental security markets with insider information: an application of rational-expectations models. *Journal of Political Economy* 90 (4), 663–698.
- Pradier, P.-C., 2019. Betting on sport: history, regulation, measurement. In: *Historical Perspectives on Sports Economics*. Edward Elgar Publishing.
- Prat, A., 2005. The wrong kind of transparency. *The American Economic Review* 95 (3), 862–877.
- Prelec, D., 2004. A Bayesian truth serum for subjective data. *Science* 306 (5695), 462–466.
- Prelec, D., Seung, H.S., McCoy, J., 2017. Solution to the single-question crowd wisdom problem. *Nature* 7638 (541), 532–535.
- Radner, R., 1979. Rational expectations equilibrium: generic existence and the information revealed by prices. *Econometrica* 47 (3), 655–678.
- Raiffa, H., Schlaifer, R., 1967. *Applied Statistical Decision Theory*. Wiley Interscience.
- Raith, M., 1996. A general model of information sharing in oligopoly. *Journal of Economic Theory* 71, 260–288.
- Reimers, I., Shiller, B., 2019. The impact of telematics on competition and consumer behavior in insurance. *The Journal of Law & Economics* 62 (4), 613–631.
- Rhode, P.W., Strumpf, K.S., 2004. Historical presidential betting markets. *The Journal of Economic Perspectives* 18 (2), 127–141.
- Riley, J., Zeckhauser, R., 1983. Optimal selling strategies: when to haggle, when to hold firm. *The Quarterly Journal of Economics* 98, 267–290.
- Robles-Garcia, C., 2019. Competition and Incentives in Mortgage Markets: the Role of Brokers. Discussion paper. Stanford University.
- Roesler, A., Szentes, B., 2017. Buyer-optimal learning and monopoly pricing. *The American Economic Review* 107, 2072–2080.
- Rosenthal, R., 1979. The file drawer problem and tolerance for null results. *Psychological Bulletin* 86 (3), 638.
- Rothschild, D., 2009. Forecasting elections: comparing prediction markets, polls, and their biases. *Public Opinion Quarterly* 73 (5), 895–916.
- Rothschild, D.M., Sethi, R., 2016. Trading strategies and market microstructure: evidence from a prediction market. *The Journal of Prediction Markets* 10 (1), 1–29.
- Rubinstein, M., 1974. An aggregation theorem for securities markets. *Journal of Financial Economics* 1 (3), 225–244.
- Rubinstein, M., 1975. Securities market efficiency in an Arrow-Debreu economy. *The American Economic Review* 65 (5), 812–824.

- Santosh, A., Cole, S., Sarkar, S., 2017. Understanding the advice of commissions-motivated agents: evidence from the Indian life insurance market. *Review of Economics and Statistics* 99 (1), 1–15.
- Savage, L.J., 1971. Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association* 66 (336), 783–801.
- Schaefer, M., Sapi, G., 2020. Learning from Data and Network Effects: the Example of Internet Search. DIW Berlin Discussion Paper No. 1894.
- Schafer, B., Konstan, J., Riedl, J., 1999. Recommender systems in e-commerce. In: *Proceedings of the 1st ACM Conference on Electronic Commerce*, pp. 158–166.
- Schlag, K.H., Tremewan, J., Van der Weele, J.J., 2015. A penny for your thoughts: a survey of methods for eliciting beliefs. *Experimental Economics* 18 (3), 457–490.
- Scotchmer, S., 2004. *Innovation and Incentives*. MIT Press.
- Sethi, R., Vaughan, J.W., 2016. Belief aggregation with automated market makers. *Computational Economics* 48 (1), 155–178.
- Sethi, R., Yildiz, M., 2012. Public disagreement. *American Economic Journal: Microeconomics* 4 (3), 57–95.
- Sethi, R., Yildiz, M., 2016. Communication with unknown perspectives. *Econometrica* 84 (6), 2029–2069.
- Shannon, C., 1948. A mathematical theory of communication. *The Bell System Technical Journal* 27 (3), 379–423.
- Shapin, S., 1994. *A Social History of Truth: Civility and Science in Seventeenth-Century England*. University of Chicago Press.
- Shapiro, C., Varian, H., 1999. *Information Rules: A Strategic Guide to the Network Economy*. Harvard Business School Press, Boston.
- Shiller, B., 2021. Optimized Sticky Targeted Pricing. Discussion paper. Brandeis University.
- Shiller, B.R., 2020. Approximating purchase propensities and reservation prices from broad consumer tracking. *International Economic Review* 61 (2), 847–870.
- Shin, H.S., 1991. Optimal betting odds against insider traders. *The Economic Journal* 101 (408), 1179–1185.
- Shin, H.S., 1992. Prices of state contingent claims with insider traders, and the favourite-longshot bias. *The Economic Journal* 102 (411), 426–435.
- Siemroth, C., 2021. When Can Decision Makers Learn from Financial Market Prices? *Journal of Money, Credit, and Banking* 53 (6), 1523–1552.
- Simonsohn, U., Nelson, L.D., Simmons, J.P., 2014. P-curve: a key to the file-drawer. *Journal of Experimental Psychology: General* 143 (2), 534–547.
- Snowberg, E., Wolfers, J., 2010. Explaining the favorite–long shot bias: is it risk-love or misperceptions? *Journal of Political Economy* 118 (4), 723–746.
- Sobel, J., 2005. Interdependent preferences and reciprocity. *Journal of Economic Literature* 43 (2), 392–436.
- Somerville, P., 1954. Some problems of optimal sampling. *Biometrika* 41 (3/4), 420–429.
- Soukhoroukova, A., Spann, M., Skiera, B., 2012. Sourcing, filtering, and evaluating new product ideas: an empirical exploration of the performance of idea markets. *The Journal of Product Innovation Management* 29 (1), 100–112.
- Spann, M., Skiera, B., 2009. Sports forecasting: a comparison of the forecast accuracy of prediction markets, betting odds and tipsters. *Journal of Forecasting* 28 (1), 55–72.
- Spears, B., LaComb, C., 2009. Examining trader behavior in idea markets: an implementation of Ge’s imagination markets. *The Journal of Prediction Markets* 3 (1), 17–39.
- Stephan, P.E., 2012. *How Economics Shapes Science*. Harvard University Press, Cambridge, MA.

- Stern, S., 2004. Do scientists pay to be scientists? *Management Science* 50 (6), 835–853.
- Stigler, G.J., 1983. Nobel lecture: the process and progress of economics. *Journal of Political Economy* 91 (4), 529–545.
- Stigler, G.J., Stigler, S.M., Friedland, C., 1995. Journals of economics. *Journal of Political Economy* 103 (2), 331–359.
- Stigler Committee on Digital Platforms, 2019. Final Report. Discussion paper. University of Chicago.
- Strevens, M., 2003. The role of the priority rule in science. *The Journal of Philosophy* 100 (2), 55–79.
- Taylor, C., 2004. Consumer privacy and the market for customer information. *The Rand Journal of Economics* 35, 631–651.
- Taylor, C.R., Yildirim, H., 2011. Subjective performance and the value of blind evaluation. *The Review of Economic Studies* 78 (2), 762–794.
- Thaler, R.H., Ziemba, W.T., 1988. Anomalies: parimutuel betting markets: racetracks and lotteries. *The Journal of Economic Perspectives* 2 (2), 161–174.
- Timmermann, A., 2006. Forecast combinations. In: *Handbook of Economic Forecasting*, Vol. 1. Elsevier, pp. 135–196.
- Tirole, J., 1982. On the possibility of speculation under rational expectations. *Econometrica* 50 (5), 1163–1181.
- Tirole, J., 2020. Competition and the Industrial Challenge for the Digital Age. Discussion paper. Toulouse School of Economics.
- Toubia, O., 2006. Idea generation, creativity, and incentives. *Marketing Science* 25 (5), 411–425.
- Varian, H., 2000. Buying, sharing and renting information goods. *Journal of Industrial Economics* 48 (4), 473–488.
- Varian, H.R., 1989. Differences of opinion in financial markets. In: *Financial Risk: Theory, Evidence and Implications*. Springer, pp. 3–37.
- Vives, X., 1988. Aggregation of information in large Cournot markets. *Econometrica* 56 (4), 851–876.
- Wallis, K.F., 2014. Revisiting Francis Galton’s forecasting competition. *Statistical Science* 29 (3), 420–424.
- West, J.D., Bergstrom, T.C., Bergstrom, C.T., 2010. The eigenfactor metrics: a network approach to assessing scholarly journals. *College & Research Libraries* 71 (3), 236–244.
- Wilson, R., 1968. The theory of syndicates. *Econometrica* 36 (1), 119–132.
- Wilson, R., 1989. Efficient and competitive rationing. *Econometrica* 57 (1), 1–40.
- Winkler, R.L., 1981. Combining probability distributions from dependent information sources. *Management Science* 27 (4), 479–488.
- Wolfers, J., 2006. Point shaving: corruption in Ncaa basketball. *The American Economic Review* 96 (2), 279–283.
- Wolfers, J., Zitzewitz, E., 2004. Prediction markets. *The Journal of Economic Perspectives* 18 (2), 107–126.
- Wolfers, J., Zitzewitz, E., 2006. Interpreting Prediction Market Prices as Probabilities. NBER Working Paper, 12200.
- Yang, K., 2020. Selling Consumer Data for Profit. Discussion paper.
- Yoganarasimhan, H., 2020. Search personalization using machine learning. *Management Science* 66 (3), 1045–1070.
- Ziman, J.M., 2000. *Real Science: What It Is and What It Means*. Cambridge University Press.
- Zuckerman, H., Merton, R.K., 1971. Patterns of evaluation in science: institutionalisation, structure and functions of the referee system. *Minerva* 9 (1), 66–100.