Gonzalo Cisternas
www.gonzalo-cisternas.org
November 1st, 2020

Research Statement

Summary

I am a microeconomic theorist with a strong interest in industrial organization and applications of game theory. Most of my work has focused on dynamic games of incomplete information and imperfectly observed actions, which I have developed to examine questions of competition, reputation formation, and information transmission in markets and organizations. To provide a sense of my most representative work, let me offer some examples:

- A worker produces output that depends on her ability, her effort, and shocks to production. Her ability is unobserved and must be inferred from the output that she produces, but only the worker observes her effort choices. Effort by the worker can then be used as a substitute for ability to generate high levels of output in an attempt to mislead employers into believing that her true ability is higher than it is.

- In oligopolistic markets, firms hold private information about their costs. The market price depends on total production and price shocks, so firms can use price realizations to make inferences about their rivals' efficiency. A low-cost firm wants to reveal itself as such. Because individual production is often not directly observable, a high-cost firm may want to produce a large quantity to mislead its rivals into believing that it is more efficient, and thus gain a larger market share in the future.

- An organization's performance depends on its adaptation to the economic environment and on the coordination of its activities. The leader of the organization has information about a set of performance-improving practices, but such knowledge is difficult to communicate. The leader thus implements changes intended to signal her knowledge, but the organization members' inferences are subjective and hence private. Changing the organization's practices improves its position in the industry, but it creates coordination costs due to mutual uncertainty about each party's knowledge.

In all these settings there is a variable—a worker's skills, a firm's costs, or a management practice—that is payoff-relevant. The true value of each of these is, however, hidden from some or all economic agents. Employers, firms and organizations then form beliefs about such variables that end up dictating these agents' behavior, but the signals that they rely on (output measures, market prices, or subjective interpretations) are often times imperfect. Consequently, two types of phenomena arise: certain types of agents may engage in strategic behavior (e.g., a high-cost firm pretending to be efficient) or may find it impossible to reveal true information (e.g., a leader not transmitting her knowledge). What are the implications of such forms of incomplete information for firms' outcomes such as output, profits, and organizational learning?

Addressing this question requires allowing for features that match the informational environments that economic agents face: enabling them to learn about unobserved variables, to take actions that are hidden from others, and even to see signals privately. The challenge is that these settings naturally lead to agents holding *private beliefs*—i.e., to agents developing a form of evolving private information in addition to having access to unobserved actions—and such

beliefs matter for behavior. From this perspective, my work can be seen as developing a framework where we can successfully handle this type of private information, a distinctive feature being the combination of statistical filtering theory and *continuous-time methods*, i.e., the connections between stochastic calculus, dynamic programming and differential equations.

The next section summarizes five papers on these *stochastic games of learning and signaling*, where the relevant states are the players' beliefs—a schematic representation explaining how the issue of private beliefs progresses through the framework is included. I conclude by discussing other work examining specific markets in detail: procurement auctions and news in social media.

Stochastic Games of Learning and Signaling

The first two papers that I describe explore games of *symmetric incomplete information*: initially (and in equilibrium) agents are equally uninformed about a common source of uncertainty. The remaining three papers examine games of *asymmetric information*: at the outset, some agents possess information that others do not. In all the settings examined, signals are noisy: thus, agents make decisions while they simultaneously and gradually learn about their environments.

The first paper in this area is entitled **"Career Concerns and the Nature of Skills"** (*American Economic Journal: Microeconomics,* Vol. 10(1), 2018) [1]. The starting point of this paper is the vast empirical evidence documenting both that employers learn about workers' productive skills over time and that such skills accumulate over the life cycle. In particular, if market competition forces wages to be based on perceived productivity and such perceptions are based on past performance, workers will find it profitable to take actions to boost their output in an attempt to be perceived as more skilled. In this paper, I ask the following question: how does this reputational mechanism operate when skills accumulate over time via past effort choices?

This question is of particular relevance because (i) information about performance flows easily among employers and (ii) human capital is a key driver of economic growth. Interestingly, the literature on reputation in labor markets—which originates from Holmström's seminal article [8]—has explored the case of exogenous skills, i.e., of productive ability that is unaffected by workers' past actions, despite the real-world importance of on-the-job skill acquisition. In line with this body of work, my analysis shows that when skills are exogenous and effort is an input to production, effort can be inefficiently high at the beginning of a worker's career. However, this is not the case if effort is solely an input to skills (and, hence, if skills are endogenous): there is *underinvestment* in human capital over the entire life cycle.

My model extends Holmström's analysis [8] by allowing effort choices to affect skill acquisition, and this generality is key for uncovering that a statistical identification problem is at the heart of the aforementioned discrepancy. Specifically, how much observed output is due to skills versus noise in the output signal? Two types of *ex-post* errors are made: when a change in output is due to noise, learning attaches a positive probability to that change being due to a change in skills; conversely, if the change in output is due to skills, learning attaches a positive probability to it being due to noise. If effort is only an input to production, the labor market assigns a positive probability to the output resulting from an increase in effort instead being linked to a skills shock. Since skills are persistent, so is the market's belief; hence, effort can generate a strong response in wages despite having only a short-term effect on output. By contrast, when workers

invest in skills, effort has a persistent effect on output, just as shocks to skills do. However, the market now assigns a positive probability to noise being the driver of the output that results from effort choices; effort then has an effect on wages that is always *weaker* than its impact on output.

In **"Two-Sided Learning and the Ratchet Principle"** (*Review of Economic Studies,* Vol. 85(1), 2018) [2], I study a general class of games of *ex-ante* symmetric uncertainty. For instance, (i) a central bank and an economy learn about an inflation trend by observing prices, but the central bank can use monetary instruments to create surprise inflation and thus boost employment; (ii) a manager and a pool of analysts learn about a firm's fundamentals by observing earnings reports, but the manager can secretly inflate the reports to meet the analysts' forecasts; and (iii) a worker and a labor market learn about the worker's skills by observing output, but the worker can use her effort choices to boost production and thus increase her wages. In these settings, an agent benefits from taking actions that improve a market's perception of an unobserved variable, but the market participants understand those incentives. Thus, central banks, managers and workers end up being trapped into behaving as anticipated by the relevant markets, without being able to systematically fool them. What are the general forces that shape such equilibrium actions, and their implications for economic outcomes?

This paper extends Holmström's model [8] to allow for general payoffs. This is an important step forward for two reasons. First, the restriction to Hölmstrom's domain has precluded the analysis of many realistic phenomena such as trade-offs between inflation and employment, acute incentives to exceed earnings' thresholds, and labor market imperfections, all of which introduce *nonlinearities*. Indeed, beyond the linear case, attempting to use standard dynamic-programming tools for finding equilibria in which actions depend on beliefs is largely hopeless: the private beliefs issue materializes in a partial *functional* differential equation to be solved. My analysis bypasses this complexity by first uncovering a remarkable structure in the environment studied: equilibrium behavior must satisfy a simple ordinary differential equation (ODE). Subsequently, I develop a *verification theorem* for Markov equilibria—an equation-based method for finding equilibria of this kind—that consists of a system of two ODEs only. The class of applications and questions that can be examined has increased considerably as a result of this new toolkit.

Second, the restriction to linear settings has obscured our understanding of the forces at play: the general approach followed, via the ODE found, uncovers that a ratchet effect that is always at play. Consider a manager who secretly boosts a firm's earnings report. While this action generates a present gain, it leads the market to have unrealistic expectations for future performance. Conversely, the manager privately believes that the firm will systematically disappoint the market absent further manipulation that accounts for the market's incorrect optimism, and this underperformance will be more acute if the market expects firms with better performance to inflate their reports more. Attempting to exhibit abnormal performance then leads to more demanding expectations for future performance, and a ratchet effect occurs. This form of dynamic cost has important implications for outcomes, explaining (i) why market forces can be insufficient to incentivize effort by workers; (ii) why monetary authorities may naturally be committed to low inflation; and (iii) why firms that expect their true earnings to exceed key thresholds can in fact manipulate their reports more than their underperforming counterparts do.

Turning to games of asymmetric information, in **"Dynamic Oligopoly with Incomplete Information"** (with Alessandro Bonatti and Juuso Toikka, *Review of Economic Studies,* Vol. 83,

2017) [3], we ask how competition in a new market unfolds as firms learn about one another, with the corresponding implications for prices, quantities, and welfare. In fact, as firms proceed with their production decisions, the market price carries information about rivals' characteristics, such as their private costs. The flipside is that the market price offers a channel for influencing the behavior of competitors. Clearly, a low-cost firm may want to establish itself as such by inducing low prices via production expansion. However, a high-cost firm may want to overproduce at a loss today if doing so leads to a larger market share in the future by misleading its rivals into thinking that its costs are low. How does this jockeying for market leadership play out? Will it be completely wasteful due to unchanged market shares but lower profits?

In line with the intuition from the previous paragraph, the firms' incentives to manipulate the beliefs of their rivals lead to overproduction. These incentives, however, decrease as learning progresses, resulting in a market price that grows over time on average and ultimately settles around the price that arises when firms have complete information about the industry. Unlike the conventional wisdom from static analyses of oligopolies, however, this competition is not entirely wasteful. Specifically, learning and signaling generate a novel effect whereby each firm's output is most sensitive to its own cost at an intermediate point in the industry lifetime. Thus, a more efficient division of the market can offset overproduction losses, leading to profits that peak in the medium run.

A successful treatment of oligopolistic dynamics has been elusive due to the complex statistical inference that such firms must perform and the way in which this inference interacts with their production decisions. In particular, upon observing a low price, how much of this price is attributable to the signaling motive of a low-cost rival versus the manipulation motive of a high-cost one? The problem is that not only are the firms' costs private but so are the firms' beliefs about their rivals: to learn about its competitors, each firm will subtract its contribution to the price, but such output carries its cost. Because such beliefs matter for production, firms will have to make inferences about others' beliefs from the price signal too, with such inferences further influencing production, and so forth. This paper shows that (i) this inference problem can be handled in a *linear-quadratic-Gaussian* setting and that (ii) continuous time is key for analyzing the joint filtering-optimization problem of each firm and its consequences on market outcomes.

In **"Consumer Scores and Price Discrimination"** (joint with Alessandro Bonatti, *Review of Economic Studies*, Vol. 87(2), 2020) [4], we examine the welfare consequences of *consumer scores* used to price discriminate. Consumer scores combine all forms of personal data into statistics that proxy for consumers' unobserved characteristics, and there are growing concerns about their use in forms of market segmentation that can hurt consumers. Unlike credit scores, these scores are highly opaque. With an ultimate goal of informing policy on consumer privacy, we ask two questions in the areas of score *composition and transparency*. First, are there ways of aggregating information that benefit consumers despite the data being used against them? Conversely, what properties would a firm prefer such scores to have in order to better price discriminate? Second, do consumers benefit from directly knowing their scores, and if so, why?

We develop a linear-quadratic-Gaussian model of reputation where a consumer has a private willingness to pay and where the score is a historical average of past purchases—since past purchases correlate with current willingness to pay, firms can then use this score to engage in a form of price discrimination. In such a context, a ratchet effect is at play. Specifically, larger

purchases lead to higher scores and hence to firms ratcheting their prices. A sophisticated consumer then reduces her demand in an attempt to be perceived as a lower type and therefore drive future prices downwards—thus, data collection can benefit the consumer by generating discounts. Our focus is on how this ratchet-based mechanism operates as we vary the score's composition (measured by the rate at which it discounts past signals) and its transparency.

We first show that firms never want to have access to disaggregated data; rather, they always prefer data to be compressed into highly persistent scores. The reason firms favor such suboptimal scores *given* the underlying data is because of the better informational content of such data: an increase in persistence, by reducing the correlation between the score and the willingness to pay, lowers the sensitivity of prices to the score, so the consumer signals more of her information. Regarding transparency, we show that for data collection to ultimately benefit consumers, scores must be made observable. Indeed, it turns out that demand is less price-sensitive when consumer scores are hidden, an effect that magnifies the ratchet effect in that the consumer is more fearful and buys less, but that generates *higher* prices. At the core of this reduced sensitivity is a signaling effect of prices: when a score is hidden, a high price is a signal of a high score and hence a signal of high future prices too. Expecting low overall consumption, the value of manipulating the price falls; i.e., the consumer becomes less price sensitive.

My last paper in this area is entitled **"Signaling with Private Monitoring"** (with Aaron Kolb, mimeo) [5], which examines a new class of two-player dynamic signaling games with the key property that the receiver privately observes noisy signals of a sender's actions. This is an unexplored area within noisy signaling, but it is important for applications. Consider the example from the introduction: a leader implements changes intended to signal her knowledge of what the organization should do to adapt to new economic conditions, but she does not know with certainty how her actions are being interpreted. If this adaptation requires coordinating such changes and the organization's response, the leader must forecast the organization's inferences. As we show, such a *second-order belief* is also her private information, and so the organization will have to forecast the leader's belief and so forth. Is there a resolution to this "beliefs about beliefs" problem? How does the leader adapt the organization to the new environment?

We develop a linear-quadratic-Gaussian model of one-sided incomplete information and one-sided private monitoring where the main object of study is linear Markov equilibria using the players' private beliefs as state variables—and the leader's second-order belief plays a central role. Indeed, behind the private nature of the second-order belief is a natural, albeit largely unexplored, phenomenon: not knowing what the organization understood, the leader relies on her *past* behavior to forecast the organization's *current* belief. Critically, this implies that there must be separation via the second-order belief channel due to the leader's past actions carrying her type. We show via applications how this novel effect can significantly alter signaling and manipulation effects that arise in the traditional knife-edge case in which all signals are public.

At a technical level, this paper introduces a novel representation of a Gaussian second-order belief under linear strategies in terms of lower-order counterparts—this result circumvents the problem of the state space growing, and it encodes the signaling effect just described. Second, the paper introduces techniques for finding equilibria in asymmetric dynamic signaling games of the linear-quadratic-Gaussian type, as the receiver is allowed to signal his evolving private information back via a public channel—to the best of our knowledge, dynamic multisided

signaling has been addressed in symmetric settings only, such as in [3]. These methods are based on solutions to *boundary value problems* (systems of ODEs with initial and terminal conditions), and one of them is general enough to be exportable to other asymmetric environments.

Schematic representation. A common theme across the papers just described is that the combination of incomplete information with unobserved actions and/or private signals leads to private information in the form of private beliefs. The progression starts from settings where changes in such form of private information do not affect behavior, and it finalizes in settings where private second-order inferences are needed for characterizing behavior.

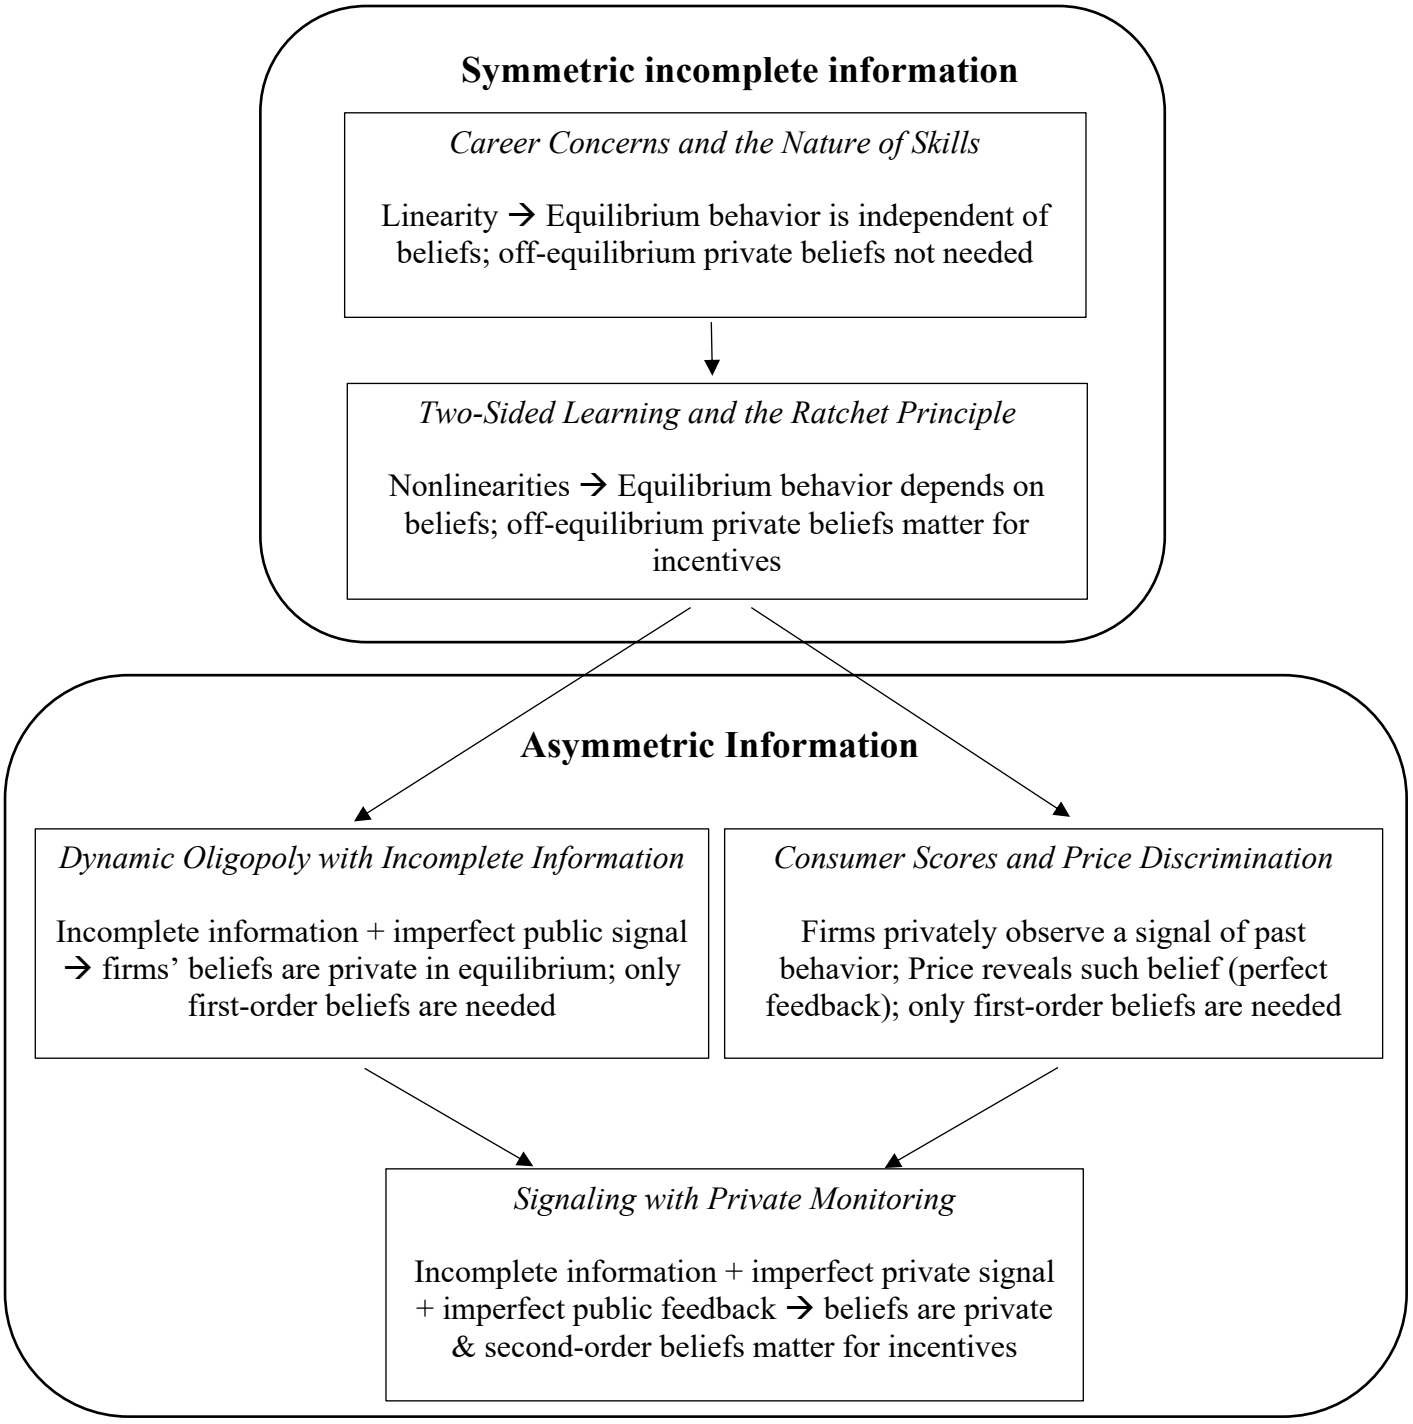

Other work

I am also interested in examining the functioning of specific markets in detail. In **"Sequential Procurement Auctions and Their Effect on Investment Decisions"** (joint with Nicolas Figueroa, *Rand Journal of Economics* Vol. 46, 2015) [6], we examine a question that is central to the performance of buyer-supplier relationships: how can buyers encourage relationship-specific investments on behalf of suppliers, while at the same time promoting competition among them in a way that minimizes the total cost of procurement? While promoting competition in every period can result in low procurement costs in the short run, it can generate weak incentives to undertake cost-reducing investments, which can be undesirable over long horizons.

To address this question, we employ a mechanism-design approach and characterize the cost-minimizing mechanism and investment level in an environment in which two projects are pursued sequentially and in which the winner of the first project can invest in a cost-reducing technology between auctions. We obtain two results. First, when the buyer is able to commit to a two-period mechanism and induces participation in both periods, the optimal mechanism gives an advantage to the first-period winner in the second auction: this is because the promise of a future advantage (and hence of high future profits) intensifies the strength of competition among suppliers today. Second, we show that the advantage granted to the first-period winner induces an investment level that is higher than the socially efficient level—to the best of our knowledge, this paper is the first to show that *overinvestment* can arise in an auction context. The relevance of this finding is twofold: it mitigates the short-term inefficiency of allocating the second project to an inefficient incumbent, and it can be a powerful tool for reducing long-term costs.

Finally, motivated by the recent misinformation attempts taking place in social media, in **"Fake News in Social Media: A Supply and Demand Approach"** (with Jorge Vásquez, mimeo) [7] we develop a model of fake news creation and diffusion to answer the following questions. First, how effective at reducing fake news are the introduction of third-party fact checkers and algorithms that filter news for users? Second, to what extent does concentration in fake news production affect market outcomes? These questions are of key importance because the response of social media platforms has focused on improving the detection and certification of misinformation, as well as on weakening the incentives of fake news producers. An advantage of our model is that we can provide answers in a framework akin to supply and demand.

Our work emphasizes a recurrent theme in the public debate: successfully addressing the fake news problem requires users to actively verify the information that they see. Thus, in our model, users can learn their authenticity of the news encountered at a cost (e.g., reading fact checks), and such verification incentives turn out to interact in non-trivial ways with policy instruments. First, the introduction of third-party fact checkers that lower such costs can be ineffective: the demand for misinformation—the fraction of users who share a news without verifying it—is reduced only for intermediate levels of fake news prevalence, as users' incentives vanish if fake news are rare or abundant. Second, these reductions create convexities in the demand for misinformation that can be exploited by trivial segmentation techniques, such as splitting users into subpopulations identical to the original one. Finally, using imperfect filters can backfire: encountering a news item makes any user more optimistic about the item's veracity, thus leading to more unverified sharing and hence to more production and diffusion of misinformation. Our model highlights the importance that natural measures of elasticity have for policy evaluation.

List of Publications and Working Papers

[1] "Career Concerns and the Nature of Skills," *American Economic Journal: Microeconomics,* 2018, **10**, 152-189

[2] "Two-Sided Learning and the Ratchet Principle," *Review of Economic Studies,* 2018, **85**, 397-351

[3] "Dynamic Oligopoly with Incomplete Information" with Alessandro Bonatti and Juuso Toikka, *Review of Economic Studies*, 2017, **83**, 503-546

[4] "Consumer Scores and Price Discrimination," with Alessandro Bonatti, *Review of Economic Studies,* 2020, **87**, 750-791

[5] "Signaling with Private Monitoring" with Aaron Kolb, mimeo.

[6] "Sequential Procurement Auctions and Their Effect on Investment Decisions" with Nicolas Figueroa, *RAND Journal of Economics,* 2015, **46**, 824-843

[7] "Fake News in Social Media: A Supply and Demand Approach," with Jorge Vásquez, mimeo.

Other cited work

[8] "Managerial Incentive Problems: A Dynamic Perspective," by Bengt Holmström, *Review of Economic Studies,* 1999, **66**, 169-182