Research Statement

Summary

I am an applied microeconomic theorist. At a broad level, my research has focused on how economic agents strategically exploit or manage counterparties' incomplete information in settings where (i) interactions are *dynamic* and (ii) *learning* about key underlying *fundamentals* takes place. To explain the interactions that I have studied, 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. However, only the worker observes her effort choices. Effort can be then used as a substitute for ability to generate high 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 firms' quantities are 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 in order to 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-maximizing practices, but such knowledge is hard 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 a mutual uncertainty about each party's knowledge.

In all these settings, there is a fundamental variable—a worker's skills, a firm's costs, a management practice—that is payoff-relevant. Moreover, the true values of such fundamentals are hidden from some, or sometimes even all, economic agents. Employers, firms and organizations then form beliefs about such variables, and these beliefs become key in dictating their behavior. Crucially, signals originating from output measures, market prices, or even subjective interpretations, are used as inputs for such learning purposes, but their imperfect nature creates two phenomena: certain types of agents may engage in opportunistic behavior (e.g., a high-cost firm mimicking an efficient one), or may find it impossible to reveal true information (e.g., a leader not fully transmitting her embodied knowledge). Ultimately, in my research, I ask the following question: what are the implications of such forms of incomplete information for outcomes such as firms' outputs, profits, and organizational learning?

A satisfactory treatment of this question requires incorporating features that match the informational environments that economic agents face: namely, allowing them to learn about unobserved variables, to take actions that are hidden to 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. To address this complexity, my work blends game theory with statistical filtering theory while exploiting the power of *continuous-time methods*, i.e., of the connections between stochastic calculus, dynamic programming and differential equations. Critically, the advantage of the continuous-time approach is not necessarily its impact at the specific level (i.e., enabling a particular breakthrough that is elusive in discrete time). Rather, the value of the approach is its impact on the whole: it provides a way to manage statistical learning, optimization, identification of forces, existence of equilibria, and computation of behavior all at the same time and with a transparent economic message.

The rest of this statement summarizes five papers on the topic of the *learning and signaling games* described above—a schematic representation of their progression, with an emphasis on how the private beliefs issue is manifested, is presented at the end of their exposition. I conclude with a section on other published work.

Learning and Signaling Games

Economic agents seldom have complete knowledge of the variables that affect their interactions. The first two papers that I describe next explore games of *symmetric incomplete information*: in the beginning (and in equilibrium), agents are equally uninformed about a common source of uncertainty. The last 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 at the same time that they gradually learn about their environments.

The first paper in this area is titled **"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 outputs in an attempt to be perceived as more skilled. In this paper, I ask the following question: how does this mechanism operate when skills accumulate over time via past effort choices?

Understanding how such career concerns affect skill acquisition is currently of particular relevance because (i) information about performance flows easily among employers, and because (ii) human capital is a key driver of economic growth. Most of the literature on reputation in labor markets—which originates from Holmström's seminal article [7]—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 the traditional literature, therefore, I show 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. Interestingly, 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 [7] to allow for effort to affect on-the-job skill acquisition. Critically, it is this generality that enables me to discover that a statistical identification problem is at the core of the aforementioned discrepancy: 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 that it was due to a change in skills; conversely, if the change in output is due to skills, learning attaches a positive probability that it was 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 resulting from a skills shock. Since skills have persistence, so does the market's belief; hence, effort can generate a strong response of wages despite having a short-term effect on output only. 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, so effort has an effect on wages that is unequivocally *weaker* than its impact on output. The paper also discusses policy responses to this inefficiency.

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 aggregate 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 all these settings, an agent benefits from taking actions that improve a market's perception of an unobserved variable. However, the market participants understand those incentives, and hence will take such motives into account when learning from signals such as prices, earnings reports, or output. Thus, central banks, managers and workers end up being trapped into behaving as conjectured by their relevant markets, without being able to systematically fool the public. Nonetheless, the actions that they take do have economic consequences. What are the general forces that shape such actions, and what are their implications for outcomes such as inflation, earnings reports, and output?

This paper extends Holmström's model [7] to allow for general payoffs. This is an important step forward in the literature 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 naturally introduce *nonlinearities*. Indeed, beyond the linear case, attempting to use traditional dynamic-programming tools for finding equilibria in which actions depend on current beliefs is largely hopeless: the private beliefs issue materializes in a type of partial *functional* differential equation to be solved. To address this challenge, I first uncover a remarkable structure of the class of games studied: equilibrium behavior must in fact 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 economic 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 constructed, uncovers that a ratchet effect that is at play in all these settings. 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 of future

performance. From the manager's perspective, therefore, she acquired a private belief that indicates that the firm will systematically disappoint the market absent further manipulation that accounts for the market's incorrect optimism. Moreover, 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 a more demanding expectation of future performance, and a ratchet effect occurs. Crucially, this form of dynamic cost has important implications for outcomes: it has helped us (i) explain why market forces can be insufficient to incentivize effort provision by workers; (ii) understand why monetary authorities may naturally be committed to keeping inflation low; and (iii) uncover novel surprising results, such as firms that expect their true earnings to exceed key thresholds can in fact manipulate them more actively than their underperforming counterparts do.

I now turn to my research on 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 over the lifetime of the industry. In fact, as firms proceed with their production decisions, the market price can carry 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 future market share via misleading its rivals into thinking that its costs are low. How does this jockeying for market leadership play out? Will it result in a wasteful competition due to unchanged market shares but lower profits due to depressed prices?

The question that we pose is an important one within the theory of industrial organization. In line with the previous intuitions, overproduction does arise from firms' incentives to manipulate the beliefs of their rivals. In addition, prices increase in expectation partly due to such manipulation incentives falling as learning progresses, ultimately settling around the price that arises when firms have complete information about the industry—critically, similar features have been documented empirically in some markets. Unlike the conventional wisdom from static analyses of oligopoly, 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, the overproduction losses can be offset by a more efficient division of the market, leading to profits that peak in the medium run.

A successful treatment of oligopolistic dynamics has been an elusive topic—this is due to the complex statistical inference that firms must perform and how this inference interacts with their production decisions. In particular, upon observing a low price, how much of it is attributable to a signaling motive by a low-cost rival versus a manipulation motive by a high-cost one? If production costs were the only source of private information, the analysis would not be difficult. The problem is that firms' beliefs about their rivals' costs are also private: to learn about firm A, firm B will subtract its contribution to the price, but B's output depends on its private cost. However, this implies that firms have to make inferences about others' beliefs, with such inferences in turn influencing production, and so forth. A key contribution of this paper is to show that (i) this inference problem can be handled in a *linear-quadratic-Gaussian* setting and

that (ii) continuous time is critical for being able to analyze the joint filtering-optimization problem faced by the firms, as well as for teasing out the properties of economic outcomes.

A second paper in this area is **"Consumer Scores and Price Discrimination"** (joint with Alessandro Bonatti, *Review of Economic Studies*, forthcoming) [4]. In this paper, we examine the welfare consequences of using consumer scores to price discriminate. Consumer scores are metrics that combine all forms of personal data into summary statistics that proxy for consumers' unobserved characteristics. The growing concerns about such scores are due to potential adverse uses in various forms of market segmentation. Moreover, unlike credit scores, these scores are highly opaque. With an ultimate goal of informing policy on consumer privacy, we seek to answer 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 like such scores to have to better price discriminate? Second, do consumers benefit from having direct access to their scores, and if so, why?

We examine these questions in a linear-quadratic-Gaussian setting where a consumer has a privately observed 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 the score to engage in a form of price discrimination. In such a context, the ratchet effect is again at play. Specifically, larger purchases lead to higher scores and hence to firms ratcheting their prices. A sophisticated consumer, understanding this possibility, reduces her demand to manipulate future prices downwards—thus, data collection can benefit the consumer by inducing lower prices. Our focus is on studying how this ratchet-based mechanism operates as we vary (i) the score's persistence (our measure of score composition) and (ii) the score's transparency.

Our first contribution pertains to uncovering practical ways of aggregating information that can be used by firms to better price discriminate—indeed, this is a major concern in any market where personalized dynamic prices are used. Specifically, to maximize their price-discrimination power, firms never want to have access to disaggregated data; instead, they prefer data to be compressed into highly persistent scores. The reason why firms favor such suboptimal scores *given* the underlying data is because of the better informational content of such data: by reducing the correlation between the score and the consumer's willingness to pay, an increased persistence reduces the sensitivity of prices to the score, so the consumer signals more of her information.

Our second contribution is key for informing the debate on consumer privacy: we show that the possibility of data collection benefiting consumers via low prices relies heavily on making scores transparent. In fact, we show that demand is less price-sensitive when consumer scores are not directly observable: this effect then magnifies the ratchet effect, in that the consumer is more fearful and consumes less, and instead leads to *higher* prices, thus hurting the consumer in both dimensions. Furthermore, we show that a sophisticated consumer facing hidden scores can be worse off than a naïve one—i.e., awareness and transparency are complementary policies. Interestingly, the reduced sensitivity is due to a signaling effect of prices: when a score is hidden, a high price is a signal of a high score, and hence of high future prices; expecting low overall consumption, the value of manipulating the price falls, and demand is less price sensitive.

My last paper in this area is titled **"Signaling with Private Monitoring"** (with Aaron Kolb) [5]. In this paper, we examine information transmission via actions between an informed sender and

a receiver when the former does not observe the signals that her actions generate. This is an unexplored area within signaling games, but it is an important one for applications. In the context of my previous work, for instance: there is ample evidence of the use of subjective metrics for evaluating workers' performances, and the latter can carry information about workers' unobserved abilities via their effort choices; likewise, firms do secretly observe consumer scores, which are linked to consumers' unobserved willingness to pay via their online behavior.

Allowing for private signals of behavior opens the possibility for examining a wide range of new economic questions, but it introduces substantial challenges. 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 successful adaptation requires coordinating such changes and the organization's response, the leader must forecast the organization's inferences in order to reduce coordination costs. As we show, such *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 in this case? How are behavior, the extent of knowledge transfer, and performance affected by the organization's information system?

This paper develops a class of linear-quadratic-Gaussian games as a framework for answering questions like these, where we construct linear Markov perfect equilibria using the players' private beliefs as state variables. Critically, behind the private nature of the second-order belief is a natural phenomenon: the leader must rely on her *past* behavior to forecast the organization's *current* belief. This effect alters the informational content of the leader's actions relative to the knife-edge case in which all signals (and hence, beliefs) are public: different leader types separate not only due to their different adaptation motives, but also because they expect their organizations to hold different beliefs—i.e., there is separation along the second-order belief channel. Our ability to characterize equilibria featuring this dependence on past play allows us to give a deeper look to the role of information structures within organizations: in particular, we show that worsening the information fed to a leader can result in more knowledge transmission.

At a technical level, the paper introduces a novel representation of a Gaussian second-order belief under linear strategies in terms of lower-order counterparts—this result is key for "closing" the state space, 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—to the best of our knowledge, dynamic multisided signaling has been addressed in symmetric settings only, such as in [3]. These methods are tailored to private- and common-value environments, and are based on solutions to *boundary value problems* (systems of ODEs with initial and terminal conditions). Crucially, one of these approaches is novel in the literature of dynamic games, and is applicable to games beyond the class examined.

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.

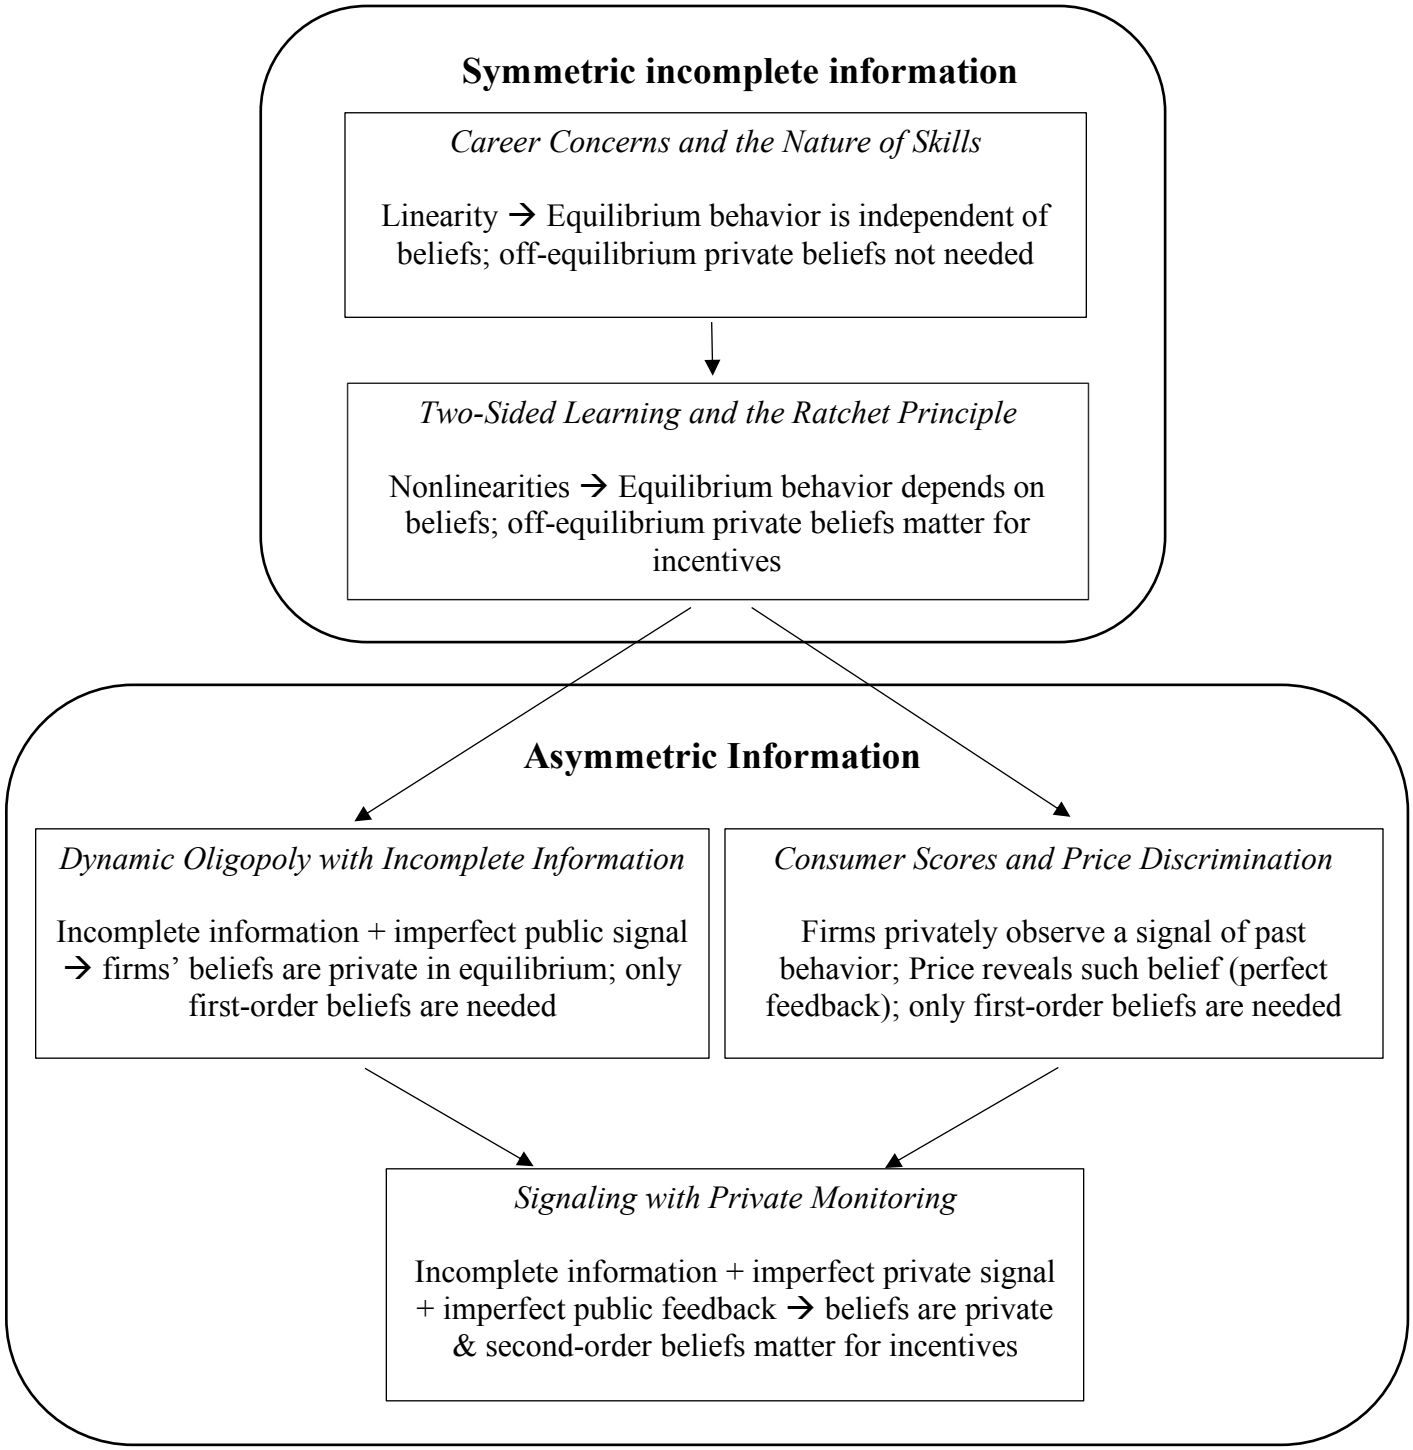

Other work

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 in 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 the 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 find 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.

List of Papers and Publications

[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,* 2019, forthcoming

[5] "Signaling with Private Monitoring" with Aaron Kolb, in preparation for submission.

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

Other cited work

[7] "Managerial Incentive Problems: A Dynamic Perspective," by Bengt Holmström, *Review of Economic Studies,* 1999, **66**(1), 169-182