

Research Statement

I began studying Economics only as a graduate, having first studied Mathematics as an undergraduate. It came as a revelation to be reading contemporary research almost from the start and to be able to relate some of the tools I had acquired and the precise ideas I was learning with public policy issues and the concerns of my peers in banking, consulting, and management. In many ways, I think that coming late to economics provides me with a somewhat fresh perspective and the zeal of a convert. After completing my masters, I spent a couple of years in consulting (in a niche firm specialized in regulatory and anti-trust work) before returning for my doctoral studies. This experience was useful to me, not only in developing some professional skills, but as a jumping off point for my thesis.

When contracting for consulting and other professional services, customers cannot know ahead of time the quality of the service they will receive, and so it is difficult to contract fully on the outcome of a transaction. As a result, reputation plays a critical role in terms of both consumers' choices to buy a service and the way that the consultant performs.

In addition, consulting firms prosper by recruiting, incentivizing and retaining talented consultants. Consulting firms are primarily composed of individual consultants whose behavior aggregates into the overall performance of the firms.

These general observations, internalized through my experience in the consulting industry before starting my doctoral work, led to a broad interest in reputation, information, and organizations—in particular, in the way that market structure and product-market competition interacts with decisions within organizations and individuals' careers. In my research, I have formalized and built on these ideas and the issues that they raise, using the tools of game theory and information economics to construct new models and generate insight.

Dynamics, information, and individual and collective incentives have application in fields including industrial organization, human resources and labor, finance, political science, and marketing. Though my initial interest was inspired by the example of professional services, the issues that arise and the theoretical tools and models that I have developed are widely applicable.

Reputation and Firm Behavior

A continuing theme and the work for which I am, perhaps, best known is the effect of reputational concerns on individuals and firms. I have sought to synthesize and extend existing models of reputation and use them to explore applied questions: First, does a firm end up with a reputation that reflects its true ability? Next, broadening the question to consider the interaction between different firms: How does competition affect reputation incentives? Then, inspired by the application to consulting and professional services, I have explored the interaction between individuals' reputations and the reputation of the firm. Finally, in a couple of different papers, I ask how decisions at the organizational level (such as recruitment policies, training, transparency, and organizational form) affect or manipulate individuals' reputation incentives?

In "Reputation and Survival," I consider the firm's decision to remain in the industry or to exit. The central question is: Can a run of bad luck drive an intrinsically good firm to exit the industry? The answer turns on the extent to which the firm knows its own quality since, if it does, the decision to remain in the industry despite continued failure can be informative in itself. Formally, I characterize (essentially unique) equilibria to show that a good firm will eventually leave the market when it does not know its own quality, but it never exits the industry when it does.

Whereas in “Reputation and Survival,” the firm cannot affect the outcome of a transaction (it is a pure adverse-selection model), my other work in this area supposes that by exerting effort, a firm or individual can improve the performance of a transaction (that is, there is some element of moral hazard).

In “Imperfect Competition and Reputational Commitment,” I ask whether the firm’s reputational incentives for such effort are stronger or more muted as an industry becomes more competitive, and I highlight several effects that, taken together, can lead to non-monotonicity. This might suggest that antitrust analysis in industries where reputation plays a key role requires considerable subtlety and judgment.

In “Something to Prove,” I argue that reputation incentives can be thought of as resolving uncertainty. Over time, customers gain a better idea of the seller’s ability, and so reputation incentives for effort diminish. However, a firm may want to commit to exerting effort when it is efficient to do so. I propose that a commonly observed organizational structure in the professional services—teams of seniors and juniors, with seniors selling partnership stakes to successful juniors—acts as such a commitment. In a team (which obscures the individual contributions of its members), the senior’s behavior affects the junior’s reputation. The senior, who would like to sell the firm to the junior, is motivated to work to improve the junior’s reputation. The paper formalizes the idea that when agents are young, they work to establish themselves and that when old, they work to maintain the reputation of their firms.

The paper also makes a somewhat more technical contribution, providing an endogenous mechanism (the choice to work in a team and the institution of a firm) for replenishing type uncertainty and sustaining reputation incentives.

An ongoing project, “Transparency, Career Concerns, and Incentives for Acquiring Expertise,” analyzes the effect of transparency on an individual’s incentives when he exerts effort to acquire knowledge to aid a decision, which he then goes on to make. The paper highlights that policymakers and interested observers should take care in advocating “transparency” and determines circumstances in which it may or may not be damaging. Moreover, the paper emphasizes that there are many different notions of transparency that might be relevant and might have differing effects.

The paper also plays a somewhat more technical role in incorporating approaches from two different literatures on reputation: a more established approach in which reputational concerns lead agents to take costly (but unobserved) actions; and the literature on “career concerns for experts,” in which decisions are typically costly and the underlying information is not observed. I show that these two different kinds of reputational concern can interact in a subtle and, perhaps, surprising way.

These papers and my work on reputation, motivated by the application to professional services, have led to some broader insights and knowledge in this area, which are, to a large extent, summarized in a wide-ranging review paper “Seller Reputation” (with Steven Tadelis). This review has already found its way onto a number of graduate reading lists.

Organizations and Personnel Economics

As mentioned above, in “Something to Prove,” I highlight a reputational rationale for teamwork as a dominant form of production and the partnership form as a dominant organizational form in professional services. In this and a number of other papers, I unpack the “black-box” of the firm and consider organizational decisions—in particular, and building on my interest in information and reputation, decisions that affect information of consumers or rival employers in ways that lead to higher profits.

In “Recruitment, Training and Career Concerns,” Juanjo Ganuza and I step back from manipulating the method of production directly and explore the role of recruitment and training policies in shaping employees’ reputation incentives. In particular, whereas numerous books on airport bookstore shelves exhort managers to recruit the best and train them, we maintain that different recruitment strategies can lead to the same

average quality of new recruits, and exactly opposite implications for how much effort these recruits would exert on the job.

In “Information and Human Capital Management,” Ian Jewitt, Clare Leaver and I highlight that a current employer can affect the information—and, thus, the perception—that rival employers have of a particular current employee. We abstract from the ways that current employers affect information release, though we have in mind organizational and production design choices, such as whether a consultant works alone or in a team and whether juniors have direct access to clients. We suppose that employers can commit to “disclosure policies” within a broad class of possibilities, and so firms compete for workers not only with salary, but also with potential career prospects (which depend on the firms’ disclosure policies). We characterize market outcomes and highlight consequences of such endogenous disclosure policies for profits, turnover, and wage dispersion.

In addition to the applied insight, the analysis raises an interesting theoretical observation: Information structures affect not only the level of adverse selection in a market, but also its distribution; there may be information structures that lead to the same level of adverse selection but with markedly different efficiency consequences. We make some progress in characterizing how.

In a paper that grew out of thinking through spatial models for my Ph.D. class in Industrial Organization and their application to professional services, I look at a different kind of decision for an organization (or an individual). “Breadth, Depth and Competition” explores how the extent of competition affects the decision to focus and specialize in a particular area or, instead, to choose breadth (at the cost of less knowledge in a particular area). In particular, I highlight that competition can have non-monotonic effects.

In ongoing work, “Reputation and Specialization,” Johannes Horner and I explore a different notion of specialization and show how reputational concerns and difficulties in commitment can lead individuals to specialize—loosely, an advantage of putting all your eggs in one basket is that it is easier to convince others that you will hold the basket carefully.

Illegal Organizations

My interest in organizations and reputation models, which are closely related to formalizations of the notion of trust, led to a collaboration and project on a somewhat different aspect: organization forms for illegal organizations. As an application, this presents an extreme case where formal contracts are ineffective (requiring the need for trust) and where external considerations play an important role and have to be traded off against the firm’s internal efficiency.

In “How to Organize Crime,” Mariagiovanna Baccara and I explore the diffusion of information within an organization. The key tradeoff is that wide diffusion of information (in particular, through fostering trust) can help the organization function more efficiently, but makes the organization more vulnerable (inasmuch as if any member leaves the organization—or is poached or captured—he takes more information with him). We provide a characterization of the optimal organizational structure. Our characterization rationalizes both hierarchical and cell-based organizations and allows to us compare authorities’ different detection policies.

We extend this work in “Interrogation Methods and Terror Networks,” allowing the authorities a broader range of behavior. This was an invited chapter for the edited interdisciplinary volume *Mathematical Methods in Counter-terrorism*.

We continue to consider developments in this area and, in particular, in communication in organizations that face an external threat and in competition between different illegal organizations.

Information-gathering consumers

In the work on reputation, I (and the literature in general) treat consumers as largely passive. In applications, including professional services, consumers exert considerable effort prior to purchase to find about the firms that they are considering.

A more recent stream of work with Guillermo Caruana and Vicente Cuñat is based around consumers' active decisions to acquire information (or, perhaps, to just spend time thinking about a product). In contrast to much of the previous literature on information release and information gathering, we have in mind that this information gathering is related primarily to horizontal characteristics. That is, a consumer gathers information about whether a product is well suited to her, but this does not say much about how well suited it is to others.

When customers acquire information to help them assess the quality of a good or service, an important theme emerges: Even though they do not pass on information to others, there are externalities in consumers' (decentralized) decisions as to what information to acquire and use for making purchase decisions. The channel for this externality is that consumers' information-gathering decisions affect firms' decisions (in pricing and in investment).

Two consequences are brought out in "Information-gathering Externalities in Product Markets." First, a reduction in the cost of information gathering can reduce welfare. Second, diversity can be beneficial (that is, for a fixed size of population, every consumer may prefer a diversity of tastes among the population of consumers). This work can help formalize and evaluate popular criticisms of school rankings and information provision on medical practitioners, thereby contributing to ongoing debates over these measures.

Whereas in "Information-gathering Externalities in Product Markets," the focus is squarely on consumers' behavior, in "Information Gathering and Marketing," we consider a firm's integrated pricing, investment, and marketing strategies. We maintain the assumption that consumers are active in information gathering. We highlight, in particular, that with ex-ante heterogeneity among consumers, marketing can act as a non-price means of discrimination, and that, again, externalities can play a role in consumer information gathering. Here, the role is to allow the firm to commit to investment.

In "Costly Search and Design," we take a simpler model of consumer information gathering and put greater focus on product design. Simplifying the consumer problem enables us to consider competition. Each time a consumer pays a search cost, he receives a new price quote and draws a match from a firm-specific distribution of possible matches. Firms choose prices and "designs." For example, a manufacturer could make computers in black or with pink polka dots; in the latter case, more people are likely to find it a bad match, but those who find it a good match are likely to do so with a very high valuation. We are able to fully characterize equilibrium in this environment and, in particular, explore the effect of reduced search costs. We highlight that profits can be non-monotonic in search costs and show that the model simultaneously generates both the "long-tail" and "superstar" effects in sales distributions that have been empirically documented as consequences of the internet.

Other work

I have also written papers on themes that do not fit neatly into the taxonomy above, though they are related as they explore information in dynamic models. In "Campaigns with Coarse Cognition," I suppose that consumers (or voters) pay more attention when they are uncertain and that they consider the implications for campaign timing when the campaign has a fixed end-date (for example, Christmas for toys, or an election in a political economy application). In particular, I highlight that such coarse cognition on the part of consumers leads both to a pre-emption effect that moves rivals to campaign early and to a "last-word" effect, which leads them to move late. In simple examples, I can perform comparative statics exercises that show how asymmetries between rivals affect campaign timing.

Finally, Vicente Cuñat and I explore a fairly different application of information economics in “Long-term Debt and Hidden Borrowing.” In this paper, we characterize the optimal form of a long-term debt contract when an entrepreneur can borrow from a source that the primary lender cannot observe (for example, from the black market or through trade credit). Borrowers simultaneously access both sources of funds, even though the hidden borrower is more expensive. In addition, the presence of a hidden borrower leads to a relatively simple form for the primary debt contract (that is, one where repayments do not depend on the performance of the project) and to inefficiently late liquidation of poor projects.