

# RESEARCH SUMMARY

**Alessandro Pavan**

Department of Economics  
Northwestern University

November 2012

The central theme of my research is the analysis of the role of information in strategic situations that are of interest to both micro and macro economists. In particular, my research focuses on global games, the social value of information and coordination, mechanism design, and two-sided markets. Below I explain the contribution of my work to each of these areas.

## 1 Global games

Global games are games of incomplete information in which the uncertainty about the underlying payoff structure is “large” in the sense that players do not exclude ex-ante any possible payoff profile (although certain profiles may well be considered more likely than others). Global games have been used to examine a variety of phenomena, including currency crises, bank runs, debt crises, investment spillovers, liquidity crashes, adoption of technology standards, and political change.

Building on results from Carlsson and van Damme (1992), Morris and Shin (2002), and Frankel, Morris and Pauzner (2003), the approach followed in most applications of global games is to assume certain exogenous information structures as a selection device—as a tool to achieve the convenience of unique-equilibrium comparative statics—without investigating what determines information in the first place. For many questions, however, understanding the endogeneity of information is the key to understanding the phenomenon under examination. The distinctive feature of my research in this area (joint with Marios Angeletos and Christian Hellwig) is in investigating the consequences of endogenizing the sources of information both on the determinacy of equilibria and on the predictions that these games deliver in applications.<sup>1</sup>

The first paper that endogenizes information in global games is [6] “*Signaling in a Global Game: Coordination and Policy Traps.*” Consider a central bank trying to prevent a speculative run against the domestic currency by raising interest rates or taking other defense measures that increase the cost of speculation. Such interventions convey information about the bank’s perception of the underlying fundamentals, as well as its willingness and ability to defend the currency. Importantly,

---

<sup>1</sup>Related to global games is also my work with Bruno Jullien [18] “*Platform Competition under Dispersed Information*”.

this information is endogenous, for it depends on the strategy of the policy maker. What the paper shows is that this endogeneity leads to multiple equilibria, despite the fact that market participants have heterogeneous information about the type of the policy maker.

The multiplicity in this paper originates from the combination of signaling and coordination; it vanishes when the policy maker faces a single receiver or when interventions do not convey information. Furthermore, contrary to the multiplicity in standard global games, this multiplicity does not rely on the policy being a public signal; it also obtains in environments in which the action of the policy maker is observed with idiosyncratic noise. Finally, contrary to the multiplicity in standard signaling games, the multiplicity documented here does not rely on the freedom to choose out-of-equilibrium beliefs; it is robust to perturbations in which the support of the policy signal is invariant to the policy choice, so that beliefs are always pinned down by Bayes' rule.

This paper has attracted attention both for its theoretical insights (warning against the use of incomplete information just as a selection device to achieve the convenience of unique-equilibrium comparative statics) and for its applied contribution (warning the policy maker that he may be trapped into a position in which the best he can do is to confirm market expectations). These predictions seem relevant in many environments, including IMF interventions during debt crises and central bank interventions during credit crunches.

Another natural source of endogenous information is learning from past outcomes in a dynamic setting. We consider such a possibility in [2] "*Dynamic Global Games of Regime Change: Learning, Multiplicity, and Timing of Attacks.*" Games of regime change are coordination games in which a status quo is abandoned, causing a discrete change in payoffs, once a sufficiently large number of agents "attacks" it. These games have been used to model a variety of crises phenomena, i.e. situations where a significant change in outcomes can be triggered without any significant change in fundamentals. Virtually all applications of these games have been confined to static frameworks: they abstract from the possibility that agents take multiple shots against the status quo and that agents' beliefs about the underlying fundamentals evolve over time. Yet, these two possibilities are important, both from an applied and from a theoretical perspective. First, crises are intrinsically dynamic phenomena. For example, in currency crises, speculators can attack again and again until they induce devaluation, and their expectations about the possibility to induce devaluation in the present naturally depend on the central bank's interventions in the past. Second, learning in a dynamic setting may critically affect the level of strategic uncertainty (i.e. uncertainty about other players' actions) and thereby the dynamics of coordination and the determinacy of equilibria.

A difficulty with extending global games to dynamic settings comes from the fact that one has to keep track of the endogenous evolution of beliefs. The paper provides a recursive algorithm that permits one to do so in a tractable way. This result is then used to show how dynamics can lead to multiple equilibria under the same conditions on the precision of exogenous private and public information that guarantee uniqueness in the static benchmark. We also show that fundamentals may predict eventual outcomes, e.g. whether a currency is devalued, but not the timing and number of attacks and that equilibrium dynamics alternate between phases of "tranquility," in which no

attack is possible, and phases of “distress”, in which an attack can be triggered by a shift in “sentiments,” without any change in fundamentals.

The multiplicity of equilibria that emerges in global games with endogenous information is however very different from the multiplicity in standard coordination environments with complete information and need not preclude concrete and testable predictions. We demonstrate this point in [12] “*Selection-Free Predictions in Global Games with Endogenous Information and Multiple Equilibria.*” While [6] focuses on documenting multiplicity and showing its robustness to alternative information assumptions, the questions this paper addresses are the following: What predictions, if any, one can deliver regarding equilibrium outcomes that are robust across all equilibria? How different are these predictions from those that obtain under complete information? We address these questions in the context of a flexible family of games of regime change, which have been used to model, inter alia, speculative currency attacks, debt crises, political change, and party leadership. We show that global games retain significant selection power and yield concrete testable predictions even when the endogeneity of information (in this paper originating from the signaling role of policy interventions) sustains multiple equilibria. From a theoretical viewpoint, the contribution here is in constructing a novel procedure of iterated deletion of non-equilibrium strategies that delivers probabilistic predictions that an outside observer—an econometrician—can form under arbitrary equilibrium selections.<sup>2</sup> Interestingly, none of the predictions identified by this procedure could have been made on the basis of the complete-information variant of the model. What is more, even though uniqueness does not generically obtain when information is endogenous, the equilibrium set continues to exhibit a sharp discontinuity reminiscent of that in standard global games: as the noise in the agents’ information vanishes, the set of equilibrium outcomes becomes a measure-zero subset of its complete-information counterpart. In other words, while the sharpness of the predictions improves as the noise gets smaller, it disappears in the complete-information version of the model. These points underscore how global-game techniques can retain a strong and useful selection bite despite the endogeneity of information and the ensuing equilibrium multiplicity.

Summarizing, the contribution of my research in global games is in showing that information is an integral part of the analysis, not just a selection device. Studying the sources of information permits a better understanding of the phenomena under examination and can give guidance on what information structures are more likely to emerge in different environments. Endogenizing information can bring back multiple equilibria; this is because the same information can be interpreted differently depending on which strategy has led to it. This multiplicity is however very different from the one that emerges under common knowledge and need not preclude useful predictions. In this respect, I find the recent debate about uniqueness versus multiplicity in global games potentially distracting from what, to me, is the main contribution of this literature: having highlighted the role that information plays in sharpening the mapping from model assumptions to model predictions. The insightfulness of this mapping, and not the determinacy of equilibria, is

---

<sup>2</sup>This procedure is different from the one in standard global games because of the endogeneity of information.

what matters in applications.

## 2 Social value of information and coordination

What are the welfare effects of the information disseminated by prices, market experts, or the media? Should central banks disclose the information they collect and the forecasts they make in a transparent and timely manner? Also, is the equilibrium acquisition and use of information socially efficient? If not, how does this inefficiency depend on the value agents assign to aligning their decisions to those of others? And what policies can restore efficiency or at least improve upon equilibrium welfare?

While the positive properties of the interaction between information and coordination have been largely examined, the *welfare* implications of such interaction are far less understood. Understanding what primitive forces create inefficiency in the acquisition and use of information is the distinctive feature of my ongoing research in this area.

In [1] “*Efficient Use of Information and Social Value of Information*,” we examine a tractable class of economies (games) with externalities, strategic complementarity or substitutability, and incomplete information on commonly relevant fundamentals. Contrary to global games, the distinctive feature of this class is that the equilibrium is unique, regardless of the information structure. Examples of economies that fit into this class are the large Bertrand and Cournot games examined in Vives (1990), the beauty contest model of Morris and Shin (2002), the investment spillover model in Angeletos and Pavan [8], and the business cycle models of Woodford (2002), Hellwig (2005), and Lorenzoni (2010).

We first identify possible inefficiencies in the equilibrium response to different sources of information; this is done by characterizing an efficiency benchmark that captures the best society can do under the sole constraint that information must remain decentralized. This efficiency benchmark, more than standard second-best concepts that allow the planner to transfer information from one agent to another, is what in these games permits one to answer the type of questions raised above.

We then show how the discrepancy, if any, between the equilibrium and the efficient use of information determines the social value of information, i.e. the comparative statics of equilibrium welfare with respect to the information structure. Examples of the results that the paper delivers are the following. In economies in which the equilibrium is efficient, welfare necessarily increases with the precision of any type of information. In contrast, in economies in which inefficiency emerges only under incomplete information, more accuracy (i.e. less noise in the agents’ forecasts of underlying fundamentals) necessarily boosts welfare, while more commonality (i.e. more correlation in the noise across agents) decreases welfare if and only if the equilibrium degree of coordination is excessively high (i.e. the private value of aligning decisions is higher than the socially optimal one). Finally, in economies that are inefficient even under complete information, ignorance can be a bliss, i.e. welfare may decrease with both the commonality and the accuracy of information (this happens when there is a certain correlation between the complete-information equilibrium and the first best).

Although the framework developed in [1] is fairly abstract, its insights are relevant for many applications. For example, Morris and Shin (2002) show that in economies that resemble Keynes’ beauty-contest metaphor for financial markets, more precise public information can reduce welfare. In contrast, public information is always welfare improving in the investment game examined in Angeletos and Pavan [8] “*Transparency of Information and Coordination in Economies with Investment Complementarities*” as well as in the business cycle models of Hellwig (2005) and Roca (2006) and in the large Bertrand games of Vives (1990). The equilibrium is essentially the same in all these economies, so why are the welfare effects of public information so different? As paper [1] shows, this is because the social value of information depends not only on the form of strategic interaction, but also on external effects that determine the gap between equilibrium and efficient use of information.

The work initiated in [1] has now been extended in a few promising directions. In [14] “*Policy with Dispersed Information,*” we develop a novel theory of taxation that aims at correcting inefficiencies in the equilibrium response to different sources of information. We generalize the framework in [1] to allow for arbitrary information structures and consider economies in which the decentralized use of information can fail to be efficient not only because of payoff externalities but also because of informational externalities. With payoff externalities, inefficiency manifests itself in excessive non-fundamental volatility (overreaction to common noise) or excessive cross-sectional dispersion (overreaction to idiosyncratic noise). With informational externalities, inefficiency manifests itself in suboptimal social learning (lower quality of information contained in macroeconomic data, financial prices, and other indicators of economic activity). In either case, a novel role for policy is identified: even if the government cannot centralize and communicate information to society, it can improve welfare by manipulating the incentives agents face in using their available sources of information. The key insight is that this can be done—and efficiency can be restored—by appropriately conditioning marginal taxes on aggregate activity.

Related to [1] and [14] is also [3] “*Socially Optimal Coordination: Characterization and Policy Implications.*” In this short paper—written for the 2006 meetings of the European Economic Association—we ask whether the heightened inertia (slow response to changes in fundamentals) and volatility (variation in aggregate activity generated by correlated noise in information) documented in many macro models with complementarities and heterogenous information are undesirable from a social viewpoint. We argue that the key to answering this question is the relation between the equilibrium and the socially optimal degree of coordination. The former summarizes the private value of aligning individual decisions, whereas the latter summarizes the value that society assigns to such alignment once all externalities are internalized. Anticipating results from [14], the paper shows how the socially optimal degree of coordination can be induced by appropriately designing the contingencies in the tax system.

All the papers cited above assume that the private information upon which agents base their decisions is exogenous. In [21] “*Information Acquisition and Welfare*” (joint with Luca Colombo and Gianluca Femminis) we relax this assumption and push the research agenda to a different

level: We endogenize the agents' private information. The analysis is carried out within the same framework introduced in [1] which offers the convenience of equilibrium uniqueness while allowing for a rich class of payoff interdependencies that may create possible wedges between the equilibrium and the efficient acquisition of information. The distinctive feature of this paper relative to the rest of the literature is the characterization of the efficient acquisition of private information. We start by showing why efficiency in the use does not guarantee efficiency in the acquisition of information. We then proceed by relating the discrepancy between the equilibrium and the efficient acquisition of private information to the primitives of the environment, as well as to the way information is used in equilibrium.

Importantly, the paper also shows how the acquisition of private information may change the social value of public information, i.e., the comparative statics of equilibrium welfare with respect to the quality of public information. We illustrate the implications of these results in a variety of applications that include beauty contests, monetary economies with price-setting complementarities, and economies with negative production externalities.

The interaction between information and coordination is also the focus of [13] "*Beauty Contests and Irrational Exuberance: a Neoclassical Approach*" (joint with Marios Angeletos and Guido Lorenzoni). In contrast to previous work, in this paper, strategic complementarities emerge endogenously as a result of an information spillover from the real sector to financial markets.

It is well known that financial markets look at data on aggregate investment for signals about underlying economic fundamentals (e.g. the profitability of a new technology or a new sector). At the same time, firms' incentives to invest increase with expected financial prices (this is because high financial prices raise the value of installed capital). In this paper we study the positive and normative implications of this two-way feedback during periods of intense technological change (e.g. the internet bubble at the end of the 90's, the more recent investment boom in China, or the latest financial crisis), when information about the profitability of new technologies/markets/securities is widely dispersed.

Because high aggregate investment is "good news" for profitability, asset prices increase with aggregate investment. Because an entrepreneur's incentives to invest in turn increase with the financial market assessment of his capital, an entrepreneur is willing to invest more when he expects others to invest more. This endogenous complementarity induces entrepreneurs to rely more on common sources of information regarding profitability and less on idiosyncratic sources of information.

The positive contribution of the paper is in showing how this endogenous complementarity reduces the impact of fundamental shocks (shifts in underlying profitability) and amplifies the impact of expectational shocks (correlated errors in the entrepreneurs' assessment of profitability). The normative contribution is in showing that these effects are also symptoms of inefficiency: investment reacts too little to fundamental shocks and too much to expectational shocks. This inefficiency originates in the dispersion of information, not in the fact that entrepreneurs care about financial prices. In this respect, the paper is also the first one to provide a complete micro-

foundation for Keynes' beauty-contest-like inefficiencies in the interaction between real and financial activity.

### 3 Mechanism design

My research in mechanism design can be divided in five related and interdependent areas. The first one is competing principals (also known as common agency). The second is privacy in sequential contracting. The third is dynamic mechanism design. The fourth is managerial compensation. The fifth is the design of auctions for divisible goods (e.g., Treasury auctions).

#### 3.1 Common agency

Common agency refers to a situation in which multiple principals contract non-cooperatively with the same agent. Depending on the application of interest, the principals can be manufacturers selling to a common retailer, lobbyists exerting influence on a politician, firms hiring a common consultant, or federal and state authorities regulating a firm.

There are many ways common agency can be modelled depending on the restrictions one imposes on the mechanisms that the principals can use to select their contracts. One possibility is to assume that principals simply make take-it-or-leave-it offers to the agent (these offers can be price-quantity pairs in a trading relationship or an employment contract in a labor relationship). Alternatively, one can think of the principals engaging in more sophisticated negotiation procedures during which the parties exchange information (for example about competing offers) before selecting the final contract.

While in games with a single mechanism designer, the entire set of sustainable outcomes can be conveniently characterized restricting the principal to offering simple direct revelation mechanisms in which the agent reports his "type," i.e. his exogenous private information, unfortunately, this is not the case in games with competing principals. The reason is that the agent may possess relevant private information not only about his exogenous type but also about the endogenous offers made by (and the decisions taken with) other principals.

My research in this area (joint with Giacomo Calzolari) aims at identifying mechanisms that are flexible enough to support all outcomes of interest in applications, while at the same time retaining tractability.

In [9] "*Truthful Revelation Mechanisms for Simultaneous Common Agency Games*," we consider games in which multiple principals contract simultaneously with the same agent. We introduce a new class of revelation mechanisms that, although it does not always permit a complete equilibrium characterization, it facilitates the characterization of the equilibrium outcomes that are typically of interest in applications (those sustained by pure-strategy profiles in which the agent's behavior in each relationship depends only on payoff-relevant information such as the agent's type and the decisions he is inducing with the other principals). We then illustrate how these mechanisms can be put to work in environments such as menu auctions, competition in nonlinear tariffs, and moral

hazard settings. Lastly, we show how one can enrich the revelation mechanisms, albeit at a cost of an increase in complexity, to characterize also equilibrium outcomes sustained by non-Markov strategies and/or mixed-strategy profiles.

Most of the literature on common agency assumes that the principals contract simultaneously with the agent. While this is a natural benchmark, there are many interesting environments in which contracting is sequential. Despite a fast growing number of applications, no general characterization results have been established for these games. This is what we do in [11] “*Sequential Contracting with Multiple Principals.*”

One of the difficulties with sequential contracting is that there is no single extensive form that fits all applications. Our baseline model assumes contracting is private (in the sense that downstream principals do not observe upstream mechanisms nor the decisions taken in these mechanisms) and that the sequence of bilateral relationships is exogenous (in the sense that the agent cannot choose with whom to contract at each date). We first show that all PBE outcomes can be characterized through pure-strategy profiles in which the principals offer menus of contracts and delegate to the agent the choice of the contractual terms. We then show that, in most cases of interest for applications, the characterization of the equilibrium outcomes is further facilitated by the fact that the principals can be restricted to offering incentive-compatible extended direct revelation mechanisms in which the agent reports the endogenous payoff-relevant decisions contracted upstream in addition to his exogenous private information. This is the case e.g. when the agent’s strategy is Markov, i.e. when it depends on upstream histories only through their payoff-relevant component (we also show that restricting the agent’s strategy to be Markov is without loss when information is complete and decisions are deterministic). Finally we show how the aforementioned results must be adjusted to accommodate alternative assumptions about the observability of upstream histories and/or the timing of contracting examined in the literature.

Related to [11] is [10] “*On the Use of Menus in Sequential Common Agency.*” In this paper, we construct two examples to show why, in certain sequential contracting environments, simple menus may fail to sustain all possible outcomes. The first example features an environment in which downstream principals observe upstream contracts, but not the mechanisms used to select them. In this environment, restricting the principals to offering menus means restricting the extent to which different principals can have different out-of-equilibrium beliefs about the mechanisms used upstream to select an off-equilibrium contract. When the agent’s strategy is non-Markov, this means imposing restrictions on the principals’ expectations about the agent’s behavior downstream. Such restrictions may preclude the possibility of sustaining certain outcomes.

The second example features an environment in which downstream principals observe both the mechanisms and the contracts selected upstream. The problem with simple menus is that they do not permit the principals to use payoff-irrelevant information as a correlation device. In the absence of alternative instruments (e.g. sunspots or cheap talk messages), this means restricting the extent to which a principal can correlate her decisions with those of other principals.



## 3.2 Privacy

Consider a buyer-seller relationship. Suppose the seller expects the buyer to contract downstream with another seller. When is it in her interest to grant the buyer full privacy, i.e. to commit not to disclose any information to the downstream seller? We address this question in [4] “*On the Optimality of Privacy in Sequential Contracting*” (joint with Giacomo Calzolari) which builds on my job market paper “Optimal Design of Privacy Policies.”

The reason why a downstream seller may be interested in receiving information from an upstream seller is twofold. First, the buyer’s willingness to pay for the downstream product/service may depend on the products/services purchased upstream. Second, even in the absence of complementarities, knowing what products the buyer has purchased upstream may permit the downstream seller to better tailor her contract offers and price discriminate.

To shed light on what determines the optimality of privacy, we model the exchange of information between the two sellers as a common agency game in which two principals contract sequentially with the same agent. The main theorem then shows that when (a) the upstream principal is not personally interested in the downstream level of trade, (b) the agent’s valuations are positively correlated (i.e., the sign of the single crossing condition is the same for upstream and downstream decisions), and (c) preferences in the downstream relationship are separable, then it is always optimal for the upstream principal to grant the agent full privacy. On the contrary, when any of these conditions is violated, there exist preferences for which (partial) disclosure is strictly optimal, even if the downstream principal does not pay for the information she receives. The paper also examines the welfare effects of disclosure and shows that, contrary to what is often believed, disclosure does not necessarily reduce consumer surplus and in some cases may even yield a Pareto improvement.

From a methodological viewpoint, my work on privacy illustrates how to design mechanisms that optimally screen the agent’s types and signal information to a third party. In [4], the third party is a downstream principal. In [5] “*Monopoly with Resale*” (joint with Giacomo Calzolari), the third party is another buyer in a resale market. This paper examines the intricacies associated with the design of revenue-maximizing mechanisms for a monopolist who expects her buyers to resell. Two cases are considered: resale to a third party who does not participate in the primary market and inter-bidder resale, where the winner resells to the losers. We show the following results. First, the monopolist may find it optimal to use a stochastic selling procedure, for example, using lotteries and/or inducing the buyers to randomize over different bids. Second, the monopolist may need to adopt a disclosure policy richer than the simple announcement of the decision to sell to a particular buyer (for example, she may need to disclose information about the bids). Stochastic selling procedures and richer disclosure policies permit the monopolist to better control the beliefs of the participants in the secondary market so as to extract more surplus from those bidders who participate in the primary market. The paper also shows that, when the distribution of bargaining power in the resale game depends not only on the allocation of the good in the primary market, but also on personal bargaining abilities, then it is generically impossible to obtain the same revenue as when resale can be prohibited. This result warns against what is predicted by the literature that

assumes sellers have full bargaining power in the resale game (e.g. Zheng, 2002).

### 3.3 Dynamic Mechanism Design

In ongoing work with Ilya Segal and Juuso Toikka [15] “*Dynamic Mechanism Design*” we study the problem of how to design incentive-compatible mechanisms in dynamic environments in which the agents’ private information arrives gradually over time and in which decisions are made over multiple periods. The environments we consider are fairly general in that the agents’ types are drawn from continuous (and possibly non-Markov) processes, decisions are allowed to affect the type distributions and payoffs are not restricted to be separable over time. The model covers as special cases such problems as the allocation of private and public goods to agents whose valuations follow a stochastic process, the procedures for selling new experience goods to consumers who refine their valuations upon consumption, the design of multi-period procurement auctions for bidders whose costs evolve stochastically over time and may exhibit learning-by-doing effects, and the design of optimal dynamic taxes for workers whose productivity evolves over time.

We deliver general necessary and sufficient conditions for incentive compatibility, and then show how to use these conditions to characterize optimal (profit-maximizing) mechanisms in applications.

The cornerstone of our analysis is the derivation (and validation) of an envelope formula for the derivative of an agent’s equilibrium expected payoff with respect to his private information in any Bayesian incentive-compatible mechanism. The formula represents the impact of an (infinitesimal) change in the agent’s current type on his equilibrium expected payoff. It accounts both for the familiar direct effect of the current type on the agent’s utility, as well as for the impact that all future types have on the utility, weighted by the effect that the current type has on the type distribution in each of the subsequent periods, which is both direct and indirect through its impact on the distribution of types in intermediate periods. All these stochastic effects are summarized in an *impulse response function* that describes the effect of the current type on all future ones by representing future types as a combination of the current type, of the decisions taken over time, and of independent shocks. Importantly, we identify assumptions on the primitive environment that guarantee that this formula is a necessary condition for incentive compatibility in any Bayesian incentive-compatible mechanism.

We then show how, in quasi-linear environments, this formula yields a dynamic “revenue-equivalence” result and an expression for the designer’s objective as “dynamic virtual surplus” which is instrumental to the characterization of optimal mechanisms and to the analysis of the dynamics of distortions under such mechanisms.

Lastly, we turn to the characterization of sufficient conditions for incentive compatibility and show how the latter can be verified by using the envelope formula to check that the allocation rule satisfies various monotonicity conditions, appropriately defined for the dynamic environment.

The results in this paper have various applications that go from the design of managerial compensation schemes (see [17] and [19] below) to the design of “bandit auctions” for the sale of experience goods to buyers who refine their valuations upon consumption. We show how the

profit-maximizing mechanism can be obtained as the solution to a multi-armed bandit problem and consists in allocating in each period the good to the buyer with the highest virtual Gittins index.

Related is also an older solo paper [16] “*Long-term Contracting in a Changing World.*” This paper studies the dynamics of distortions in an environment that stylizes a buyer-seller relationship. It first shows, through a simple example with finitely many types, that distortions need not vanish over time and need not be monotone in the shocks to the buyer’s valuation. The paper then discusses situations in which (a) the optimal mechanism is the same irrespective of whether the shocks are the buyer’s private information or are observed also by the seller, and (b) the distortions in the optimal quantities are independent of whether the shocks are transitory or permanent.

### 3.4 Managerial Compensation

Related to the aforementioned body of work are also two recent papers with Daniel Garrett. In [17] “*Dynamic Managerial Compensation: On the Optimality of Seniority-based schemes,*” we apply some of the techniques developed in [15] to study the properties of optimal incentive schemes for a manager who faces costly effort decisions and whose ability to generate profits for the firm changes stochastically over time. The optimal compensation scheme is obtained as the solution to a dynamic mechanism design problem with hidden actions and persistent shocks to the agent’s private information.

We show that, when the agent is risk-neutral, the optimal contract can often be implemented with a simple pay package where the manager’s compensation is linear in the firm’s profits. Furthermore, the power of the incentive scheme typically increases over time, thus providing a possible justification for the practice of putting more stocks and options in the package of managers with a longer tenure in the firm. Contrary to other explanations proposed in the literature (e.g. declining disutility of effort, or career concerns), the optimality of seniority-based schemes in our model is not driven by variations in the agent’s preferences or in his outside option. It results from an optimal allocation of the manager’s informational rents over time.

Building on the insights from the risk-neutral case, we then explore the properties of optimal incentive schemes for risk-averse managers. We find that risk-aversion reduces (and in some cases can even reverse) the profitability of seniority-based schemes whose power of incentives increases, on average, over time. The reason is that these schemes entail a high sensitivity of compensation to performance precisely in those periods in which the manager faces high uncertainty about his ability to generate cash flows for the firm. Increasing the sensitivity of compensation to performance over time thus means exposing the manager to a great deal of risk. Whether risk-averse managers with a longer tenure in the firm receive more or less high-powered incentives than younger ones then depends on the interaction between the degree of risk-aversion and the dynamics of the impulse responses of the process governing the evolution of managerial productivity.

In [19] “*Managerial Turnover in a Changing World,*” we develop a dynamic theory of managerial turnover in a world where the quality of the match between a firm and its top managers changes

stochastically over time. We characterize the joint dynamics of retention, compensation, and effort decisions under the profit-maximizing contract and compare them to their efficient counterparts.

Our key positive result shows that the firm's optimal retention decisions become more permissive with time. What in the eyes of an external observer may look like "entrenchment" is, in our theory, the result of a fully-optimal contract in a world where incumbent managers possess privileged information about the firm's prospects under their own control.

Our key normative result shows that, compared to what is efficient, the firm's optimal contract either induces excessive retention (i.e., inefficiently low turnover) at all tenure levels, or excessive firing at the early stages of the relationship followed by excessive retention after sufficiently long tenure.

These results are obtained by endogenizing the firm's separation payoff accounting for the fact that its performance under each new hire is going to be affected by the same information and incentives frictions as in the relationship with the incumbent.

On the methodological side, this paper applies techniques developed in [15] to a novel dynamic programming problem whose solution endogenizes the principal's outside option.

### 3.5 Auctions for divisible goods

My very first paper (coming out of my undergraduate thesis at Bocconi University) is a joint work with Marco LiCalzi on auctions for divisible goods: [7] "*Tilting the Supply Schedule to Enhance Competition in Uniform-Price Auctions.*" Uniform-price auctions were originally designed for Treasury securities; they have now become a popular mechanism also for on-line initial public offerings of unseasoned shares (Open IPOs), for electricity markets, and for markets for emission permits.

In a uniform-price auction, bidders submit an entire demand (supply) schedule for the divisible good being offered. The seller compares the aggregate demand with the aggregate supply and then computes a clearing (stop-out) price. Demand above the stop-out price is awarded in full, while marginal demand is prorated. Since all buyers pay the same price, these auctions are analogous to a Walrasian market. The important difference is that demand schedules are submitted strategically. This difference makes uniform-price auctions susceptible to substantial underpricing. In fact, as shown first in Wilson (1979), by submitting high inframarginal demands, bidders can prevent competition on prices and support equilibria in which the stop-out price is significantly below its Walrasian equivalent.

A feature of these auctions is that supply is inelastic and fixed in advance. This seemingly innocuous assumption implies a strategic asymmetry between the bidders and the seller: the former can use their demand schedules to inhibit price competition, but the latter cannot use her supply schedule to enhance it. Allowing the seller to use an elastic supply is a simple way to limit underpricing. Intuitively, while the steepness of the competitors' demand curves has a price effect which increases the marginal cost of a higher bid, an increasing supply schedule induces a quantity effect that raises a bidder's marginal revenue from inducing a higher stop-out price. Making the quantity effect greater than the price effect enhances competition on higher prices. What we show

in [7] is that precommitting to an increasing supply schedule can be optimal for the seller even when accounting for the costs of issuing a quantity different from a target supply.

## 4 Two-sided Markets

Two-sided markets are markets where a platform’s ability to price on one side depends on its pricing strategy on the other side. Equivalently, these are markets where the product sold by the platform is, to a large extent, “access to agents on the other side.”

In my work with Renato Gomes [20] “*Many-to-Many Matching Design*” we consider the problem of a platform that designs price-discrimination schemes to match users from one side of the market to users from the other side. Examples of the design of many-to-many matching services include the provision of Cable TV packages, the design of health insurance networks, but also the design of credit cards networks.

For concreteness, consider the problem of a Cable TV provider contracting with TV channels on one side of the market and with viewers on the other side. The Cable company’s problem can be seen from two perspectives. The more familiar one is that of designing a menu of packages of channels to offer to the viewers. The mirror image of this problem consists in designing a price schedule for the channels whereby prices are contingent on the packages in which the channel will be included (and hence on the number of viewers the channel will be able to reach—more viewers yields higher advertising revenue). By the very nature of the matching problem, the menu of channels offered to the viewers pins down the quantity schedule faced by the channels and the price schedule offered to the channels pins down the packages that the platform can offer to the viewers. As such, when designing its profit-maximizing menus, the Cable company has to internalize the cross-side effects of the schedules offered to the two sides of the market.

In order to investigate the optimality of a large variety of pricing strategies, we tackle the problem using a mechanism design approach. We allow the platform to offer *any* many-to-many matching rule that satisfies a weak reciprocity condition. We derive necessary and sufficient conditions for the welfare- and the profit-maximizing mechanisms to employ a *single network* or to offer a menu of non-exclusive networks (*multi-homing*). We characterize the matching schedules that arise under a wide range of preferences and deliver various testable predictions that relate the structure of the optimal matching rule to the distribution of match qualities. Our analysis sheds light on the distortions brought in by the private provision of broadcasting, health insurance and job matching services.

At the theoretical level, what distinguishes the problem in this paper from a standard screening problem (e.g., Mussa and Rosen (1978) and Maskin and Riley (1983)) is twofold. First, the platform faces feasibility constraints with no equivalent in the adverse selection and price discrimination literatures. Second, each agent is both a customer and an input in the matching production function. The “customer” role of an agent is summarized by his/her willingness to pay while the “input” role is captured by the agent’s attractiveness for the other side. This feature of matching markets implies that the cost of procuring an input is endogenous (it depends on the entire matching rule)

and incorporates nontrivial strategic considerations.

The first part of the paper considers a market where preferences on each side are vertically differentiated: Any two agents from the same side agree on the relative attractiveness of any two agents from the opposite side (but possibly differ in their willingness to pay to reach these agents). In the second part of the paper, we extend the theory to markets where preferences on each side exhibit both vertical and horizontal differentiation. In other words, we let the relative attractiveness of any pair of agents vary with the agent they are matched with. Allowing for horizontal differentiation gives flexibility to the model and uncovers properties for the optimal matching rule that appear to square well with what we see in applications.

In [18] “*Platform Competition under Dispersed Information*” (joint with Bruno Jullien), we consider a model where platforms compete in prices on two sides and where each agent from each side possesses private information both about his/her valuations for the products/services offered by the two platforms as well as about the distribution of valuations in the cross-section of the population. Dispersed information is a natural feature of many two-sided markets (e.g., the market for video-games consoles and the market for new operating systems).

The paper uses techniques from the global-game literature to arrive to a formula for the equilibrium prices that relates the latter to the inverse semi-elasticity of the residual demands, accounting for the interaction between information and network effects. We show that, unlike the case of full information (e.g., Armstrong (2006)), the equilibrium price on each side depends not only on the intensity of the network externality on the opposite side but also on the intensity of the own-side network externality. In particular, when preferences are positively correlated between the two sides, an increase in the intensity of the network externality on side  $i$  raises the price on side  $i$  and reduces the price on side  $j$ .

We also show that the equilibrium prices depend on the underlying information structure only through an index of mutual forecastability that captures the two sides’ ability to forecast each other.

We then use the results to study the effects of various advertising campaigns that increase the agents’ ability to estimate their own valuations and/or the distribution of valuations on the other side of the market.

## References

- [1] “Efficient Use of Information and Social Value of Information,” *Econometrica*, Vol. 75(4), July 2007, 1103-1142 (with G-M. Angeletos)
- [2] “Dynamic Global Games of Regime Change: Learning, Multiplicity, and Timing of Attacks,” *Econometrica*, Vol. 75(3), May 2007, 711-756 (with G-M. Angeletos and C. Hellwig)

- [3] “Socially Optimal Coordination: Characterization and Policy Implications,” *Journal of the European Economic Association (PEP)* Vol. 5(3), May 2007, 585-593 (with G-M. Angeletos)
- [4] “On the Optimality of Privacy in Sequential Contracting,” *Journal of Economic Theory*, Vol. 130(1), September 2006, 168-204 (with G. Calzolari)
- [5] “Monopoly with Resale,” *Rand Journal of Economics* Vol. 37(2), Summer 2006, 362-375 (with G. Calzolari)
- [6] “Signaling in a Global Game: Coordination and Policy Traps,” *Journal of Political Economy*, Vol. 114(3), June 2006, 452-485 (with G-M. Angeletos and C. Hellwig)
- [7] “Tilting the Supply Schedule to Enhance Competition in Uniform-Price Auctions,” *European Economic Review*, Vol. 49, January 2005, 227-250 (with M. LiCalzi)
- [8] “Transparency of Information and Coordination in Economies with Investment Complementarities,” *American Economic Review (PEP)* Vol. 94(1), May 2004, 91-98 (with G-M. Angeletos)
- [9] “Truthful Revelation Mechanisms for Simultaneous Common Agency Games,” May 2011, *American Economic Journal: Microeconomics* 2(2), 132-190 (with G. Calzolari)
- [10] “On the Use of Menus in Sequential Common Agency,” April 2007 *Games and Economic Behavior*, September 2008, Vol. 64(1), 329-334 (with G. Calzolari)
- [11] “Sequential Contracting with Multiple Principals,” *Journal of Economic Theory*, February 2009, Vol. 144(2), 503-531 (with G. Calzolari)
- [12] “Selection-Free Predictions in Global Games with Endogenous Information and Multiple Equilibria” October 2012 (with G-M. Angeletos), forthcoming in *Theoretical Economics*
- [13] “Beauty Contests and Irrational Exuberance: A Neoclassical Approach” March 2010 (with G-M. Angeletos and G. Lorenzoni)
- [14] “Policy with Dispersed Information,” *Journal of the European Economic Association*, March 2009, Vol. 7(1), 1-50 (with G-M. Angeletos)
- [15] “Dynamic Mechanism Design” August 2012 (with Ilya Segal and Juuso Toikka)
- [16] “Long-Term Contracting in a Changing World,” December 2007.
- [17] “Dynamic Managerial Compensation: On the Optimality of Seniority-Based Schemes” October 2011 (with Daniel Garrett)
- [18] “Platform Competition under Dispersed Information” November 2012 (with Bruno Jullien)
- [19] “Managerial Turnover in a Changing World,” October 2012 (with Daniel Garrett), forthcoming in *Journal of Political Economy*.

[20] “Many-to-Many Matching Design” October 2012 (with Renato Gomes)

[21] “Information Acquisition and Welfare”, May 2012 (with Luca Colombo and Gianluca Ferraris)