Summary. I am an economic theorist that is primarily interested in dynamic models. In particular, I am interested in contracting problems in which random events determine the actions available to the players at each point in time, and in matching markets where participants learn gradually about the quality of their match. In addition, I am also interested in studying optimal privacy regulation in various market settings. In what follows, I describe my research.

1. Dynamic Contracting with Random Opportunities. Interactions between a principal and an agent often take place in complex and dynamic environments: seasonality and random shocks affect demand, workers accumulate skills, and business opportunities arrive and disappear at random. I believe that the randomness inherent in such environments is a feature that should be explored rather than a bug that should be abstracted away by various smoothness assumptions.

My first paper in this research agenda, *Optimal Contracts with Randomly Arriving Tasks*,[1] explores the implications of the need to adapt a worker's tasks to the fluctuating needs of his employer. To this end, my coauthor and I analyze a model where different production opportunities arrive randomly in a manner that is i.i.d. across periods. Each period, the agent can exert publicly observable effort whose productivity depends on the realized production opportunity and the principal can provide a periodic wage. We find that, in general, as time goes by, the worker's expected effort within a period decreases while his periodic wage increases.

Our analysis offers new insights into wage ladders and seniority in workplaces by drawing a clear connection between the intertemporal variability of on-the-job assignments and the dynamics of effort and compensation. Specifically, sophisticated promotion dynamics may arise efficiently even in the absence of traditional frictions such as imperfect or asymmetric information about the worker's ability or actions, search frictions, accumulation of expertise, etc. Moreover, our results show that wage stickiness arises due to fluctuations in the worker's tasks.

My second paper in this research agenda, *Monotone Contracts*,[2] aims to uncover the qualitative impact of random opportunities in a general class of contracting problems. In essence, we consider contracting problems that can be represented as stochastic games—in each period the players play a randomly drawn stage game, observe its outcome, and collect payoffs—in which the principal commits to a long-term strategy and the agent reoptimizes his play at every history. For such contracting problems, my coauthor and I define a class of components that we refer to as activities. Our main result identifies a (tight) property of activities that guarantees

that, in optimum, and irrespective of the exact details of the contracting problem, the level of the activity changes over time only in the direction that favors the agent.

Our notion of an activity encompasses many prominent aspects of contacting problems. Examples include worker's daily effort, monthly wage, production volume, level of authority of a bureaucrat or a unit in an organization, financing to an entrepreneur, quality of supplied products, and more. Thus, in addition to offering a generalization and unification of several seemingly unrelated monotonicity results in the existing literature, our framework paves the way for deriving monotonicity results for new, more intricate, contractual components in richer settings. Moreover, the "detail-free" nature of our approach circumvents the need to fully characterize optimal contracts, and so it can be used to derive qualitative properties even in settings where it is challenging to write a tractable model that exhibits all the essential features of the setting of interest.

My third paper, *Dynamic Nonmonetary Incentives*,[3] focuses on the impact of asymmetric information about the availability of opportunities. In particular, my coauthor and I study applications in which instantaneous monetary transfers are infeasible or inferior to other methods of compensation and, more importantly, there is uncertainty regarding the availability of future compensation opportunities. We show that when compensation is provided by randomly arriving rewards, the unique optimal way to provide incentives to an agent is by a "time-constrained carte blanche," i.e., allowing him to enjoy all rewards that arrive in a fixed time interval that begins at the present moment. When there are multiple types of rewards, then the agent will have a different "time-allowance" for each reward type. We find that, in optimum, the principal often permits reward types for which there is a high cost of providing a util to the agent even when it is not necessary to do so. Thus, our finding can rationalize many seemingly inefficient forms of compensation.

My fourth paper in this line of research, *A Theory of Front-line Management*,[4] has a more applied motivation. Large organizations generally adopt a hierarchical structure and use incomplete contracts. Thus, mid- and low-level managers can fulfill a fundamental role within organizations beyond their traditional supervisory role. Namely, such managers, if given sufficient discretion, can effectively adapt the assignments of their subordinates to provide timely responses to unexpected opportunities and threats. That is, they can become an effective means to *complement* workers' incomplete labor contracts. The managerial problem faced by such managers is unique in that they need to incentivize a worker to exert effort on rare occasions (when an event that is not covered by the baseline contract occurs) and they typi-

cally lack control over wages. Moreover, there is no compelling reason to assume that, within their specific interaction, the worker and the manager discount the future in the same way.

We study the interaction between a front-line manager and a worker, and characterize the "managerial style" as a function of the players' relative patience and information. The broader contribution of this paper is twofold. First, we propose a tractable model of dynamic contracting with asymmetric discounting. More importantly, our analysis has novel implications for organizational design. Our analysis shows that evaluating individual contracts in isolation may be insufficient as the relative properties of workers' and managers' respective contracts affect the dynamics of their interaction. For example, if these contracts lead to different intertemporal incentives, then a manager who is relatively more patient than the worker would increase the worker's long-run effort, whereas a relatively impatient manager would improve the worker's retention rate through the adoption of a seniority system. Moreover, we show that, in some cases, the organization benefits from inducing informational frictions between workers and front-line managers.

2. Search and Matching.

My recent work has focused on exploring two-sided matching markets in which agents learn gradually about the quality of a proposed match. In the paper *Search, Dating, and Segregation in Marriage*,[5] my coauthors and I study a marriage market in which, before committing to a match, a couple who have just met can start dating to learn if they are a good fit for one another or not. The contribution of this paper is twofold. At the methodological level, we develop a novel tractable framework that incorporates pre-match information acquisition into the matching-with-search-frictions paradigm.

We then apply this framework to examine if premarital learning can explain observed patterns in the marriage market. We find that not only does our model predict observed matching patterns — e.g., people are more likely to marry a partner with a similar level of education but, occasionally, they will also marry a partner with a very different level of education – but it can also explain why interracial segregation remains prevalent despite the overwhelming approval rate of such marriages (94% of Americans approve of interracial marriages). We further show that frictions in cross-racial learning (á la Phelps's (1972) notion of statistical discrimination) can lead to segregation patterns that are consistent with those that are observed in the data: e.g., that segregation becomes less common as education increases. Fi-

nally, we show that due to pre-match information acquisition a social planner may want to induce negative assortative matching even under conditions that ensure that in equilibrium there will be positive assortative matching.

In the paper *Search, Matching, and Online Platforms*,[6] we consider matching markets in which matches are facilitated by an online platform. We consider a model with vertical heterogeneity among agents in which (post-match) agents learn gradually about the quality of their match. We examine how the platform's profit-maximizing incentives shape the effect of technological advances on pricing, profits, and consumer welfare.

The main tradeoff faced by the platform lies in managing its repeated clientele. The decision of whether to terminate a match and return to the platform is endogenous, as it depends on the quality of the service provided by the platform, as well as its pricing policy. This creates a dynamic repeated clientele tradeoff in which, on the one hand, the platform wants to provide high-quality service to attract repeated clientele, but, on the other hand, it does not want to induce matches that are so good that agents never terminate them.

We find that due to this tradeoff the platform has a strong incentive to improve the speed of search on the platform but a disincentive to provide users with information that will facilitate their learning about the quality of a match. Moreover, we show that technological advances (i.e., faster search of better-quality information) lead the platform to reduce the fee that it charges from its users; A result the contradicts the basic intuition that the price of a good sold by a monopolist is increasing in its quality. The main contribution of this paper is to shed light on the incentives of online platforms in adopting new technologies. In addition, at the methodological level, our paper is one of the first to study matching markets with horizontal differentiation and derive clear comparative statics for the equilibrium of such a matching market.

3. Privacy Design. I am also interested in the economics of privacy. Motivated by advances in personalized medicine and the proliferation of privacy laws, in the paper, *What Should a Firm Know? Protecting Consumers' Privacy Rents*,[7] my coauthor and I consider markets in which the consumer's type impacts the payoff of both the consumer and the firm. We ask how a designer should structure the firm's information about the consumer's so that the firm has enough information to facilitate mutually beneficial transactions, but is prevented from extracting too much surplus through exploitation of its informational advantage. Our main result is to derive

a geometric characterization of the consumer-optimal information structure. If the firm can make only two offers to the consumer, then the payoff space is divided into two half-spaces by a (straight) downward-sloping line that passes through the origin. Each half-space corresponds to an offer, and the firm is only allowed to learn to which half-space the state of the world belongs. If the firm can make more offers, then the analysis is more involved, but the basic result extends in the sense that the payoff space is partitioned into polyhedral cones (i.e., sets that are bounded by a finite number of hyperplanes). This result provides a guideline for the evaluation of privacy policies by emphasizing the more relevant dimensions of the state and policy spaces. In particular, it suggests a measure of "closeness of states" that specifies which states the firm should be allowed to distinguish between. We apply this result to the design of efficient health insurance markets.

In *Privacy, Market Structure, and Prices in Competitive Search Markets*,[8] we analyze a different aspect of privacy that is relevant even when personalized pricing is not practiced by firms. Namely, we study how privacy regulation affects firms' ability to reach consumers through interest-based advertising in markets with search frictions. To this end, we consider a two-sided search market in which consumers receive targeted advertisements from firms and also actively search for firms themselves. We assume that the viability of targeted advertising depends on privacy regulation, and ask how consumer welfare is impacted by privacy laws due to the induced differences in the distribution of the number of price quotes that the consumer observes before purchasing the good. We show that privacy protection leads to a reduction in consumer surplus if firms price discriminate between the search and advertising markets. If firms do not price discriminate, then privacy protection can either increase or decrease consumers' surplus and firms' profits, and we characterize conditions for both. At a broader level, our analysis raises doubts about the efficacy of modern privacy protection laws (e.g., the EU's "General Data Protection Regulation") that give consumers control over their personal information. In particular, we show that even though each consumer benefits individually from sharing her information, consumer welfare may be higher when privacy is maintained. That is, we argue that there is a social aspect to privacy protection that may justify a more paternalistic intervention to protect consumers' privacy.

4. Other Projects. I have worked on dynamic learning problems also in a principal–agent settings. In *Sequential Learning*,[9] my coauthors and I study a moral hazard problem that typically arises in R&D and standardization processes (e.g., drug ap-

proval) that involve learning and approval by multiple heterogeneous players who operate in a sequential manner. In such ventures, the player who moves first (e.g., a pharmaceutical company) may have an incentive to exaggerate and even fabricate positive findings to convince the player who moves second (e.g., the FDA) to approve the venture. Taking this into account, the second mover may be less inclined to trust the first mover's findings and may want to examine the venture or the findings himself.

We develop a flexible framework that enables us to study the dynamics of efficient sequential collaboration in the realm of this moral hazard problem. We find that, virtually regardless of the learning technology that is available to the player who moves second (i.e., whether he looks for positive findings, faults, or fraud), in the Pareto-dominant equilibrium of our sequential learning game, the player who moves first fabricates positive findings on the equilibrium path and the frequency with which she does so goes down gradually as time progresses until, at some point in time, she becomes honest. As a result, the player who moves second treats the first mover's submissions more favorably the later he receives the project. We conclude that occasional manipulation is inherent in the early stages of an efficient sequential collaboration and that trust is built monotonically and gradually.

Finally, I have also worked on Monetary Economics. Historically, modern economies have looked to central banks to carry out two main functions: to act as lenders of last resort during credit crunches and to maintain price stability. An important question is whether these two goals conflict with one another: central banks that act as lenders of last resort create money to lend to distressed banks, and such liquidity injections can be inflationary.

In the paper *Money Under the Mattress: Inflation and Lending of Last Resort*,[10] my coauthors and I examine whether the two key functions of central banks—ensuring price stability and lending during crises—necessarily conflict. We develop a nominal model of bank runs à la Diamond and Dybvig (1983) in which individuals can store the money they withdraw "under the mattress." In this setting, lending of last resort need not be inflationary. Whether it is so depends on the interest rates the central bank charges on its loans. Our results suggest that the central bank must not charge a rate that is too low if it wants to ensure price stability, and must charge a high rate if it wants to robustly attain the ex-ante efficient outcome. At a broader level, the main contribution of our work is to highlight the implications of allowing agents to store money outside of the financial system.

References

1. Bird, D. and Frug, A. (2021): Optimal Contracts with Randomly Arriving Tasks, *Economic Journal*, 131(637):1905–1918.

2. Bird, D. and Frug, A. (2022): Monotone Contracts, *Theoretical Economics*, 17(3):1041–73.

3. Bird, D. and Frug, A. (2019): Dynamic Nonmonetary Incentives, *AEJ Microeconomics*, 11(4):111–150.

4. Bird, D. and Frug, A. (2024): A Theory of Front-line Management, Working paper.

5. See Antler, Y., Bird, D., and Fershtman, D. (2022): Search, Dating, and Segregation in Marriage, Working paper.

6. Antler, Y., Bird, D., and Fershtman, D. (2024): Search, Matching, and Online Platforms, CEPR DP18719.

7. Bird, D., and Neeman, Z. (2022): What Should a Firm Know? Protecting Consumers' Privacy Rents, *AEJ Microeconomics*, 14(4):257–95.

8. Bird, D., and Neeman, Z. (2024): Privacy, Market Structure, and Prices in Competitive Search Markets, *Journal of Law, Economics, & Organization*, Forthcoming.

9. Antler, Y., Bird, D., and Oliveros, S. (2023): Sequential Learning, *AEJ Microeconomics*, 15(1):399–433.

10. Barlevy G., Bird, D., Fershtman D., and Weiss, D. (2024): Money under the Mattress: Inflation and Lending of Last Resort, *Journal of Economic Theory*, 217:105804.