



INSTITUT
D'ÉCONOMIE
INDUSTRIELLE



Toulouse Network for Information Technology

Annual Meeting Sept 30-Oct 1, 2016 Redmond

Lists of abstracts

Daron Acemoglu (MIT) Demographics and Automation

Abstract

There is increasing concern about the future of jobs and wages as robots and other new technologies such as artificial intelligence take over tasks previously performed by labor. Much of our current knowledge about the likely consequences of these changes comes from feasibility studies that estimate which occupations and tasks can be automated in the near future and then extrapolate from these estimates into potential employment losses. However, such approaches ignore important equilibrium considerations--- in particular, both the possibility that technologically feasible automation may not be economically attractive for firms, and more importantly, that the replacement of workers in some tasks may in turn lead to the expansion of other tasks and industries. In this paper, we estimate the local equilibrium effects of the introduction of robots on employment and wages at the commuting zone level. We first show that, theoretically, the equilibrium effects of robots on employment and wages can be estimated by regressing these variables on exposure to robots---the counterfactual number of robots in a local labor market computed from the national penetration of robots into each industry and the local distribution of employment across industries. Using this approach, we show large and robust negative effects of robots on employment and wages. According to our estimates, each additional robot reduces employment by about seven workers and average wages by about 0.0015 percent.

Susan Athey (Stanford) Using Machine Learning for Personalized policy Recommendations: Welfare Benefits of Big Data

Abstract

This paper introduces new methods for constructing personalized policies, such as online recommendations, personalized medical treatment decisions, or allocation to government programs; the paper also provides an approach to evaluate the benefits of such personalized policies and establishes statistical properties of the estimates, so that valid confidence intervals can be created. New statistical efficiency properties are established for the method. Several applications are provided, and the benefits of “big data” are quantified.

Nicholas Bloom (Stanford)
Management as a Technology

Abstract

Are some management practices akin to a technology that can explain company and national productivity, or do they simply reflect contingent management styles? We collect data on core management practices from over 11,000 firms in 34 countries. We find large cross-country differences in the adoption of basic management practices, with the US having the highest size-weighted average management score. We present a formal model of “Management as a Technology”, and structurally estimate it using panel data to recover parameters including the depreciation rate and adjustment costs of managerial capital (both found to be larger than for tangible non-managerial capital). Our model also predicts (i) a positive effect of management on firm performance; (ii) a positive relationship between product market competition and average management quality (part of which stems from the larger covariance between management with firm size as competition strengthens); and (iii) a rise (fall) in the level (dispersion) of management with firm age. We find strong empirical support for all of these predictions in our data. Finally, building on our model, we find that differences in management practices account for about 30% of cross-country total factor productivity differences.

Glenn Ellison
Diversity in IT sector

Abstract

The underrepresentation of minorities in the IT workforce follows underrepresentation in high school and college coursework. At the high school level AP computer science courses are disproportionately found at high-achieving schools and only 6% of students who achieve passing grades are black or Hispanic. Increasing this number will require expanding the set of schools offering such classes and/or increasing the number of minority students who take them at high-achieving schools. A challenge related to the latter is that several public K-12 and university systems have recently shifted from race-based affirmative action plans to race-neutral alternatives. This paper explores the degree to which race-neutral alternatives are effective substitutes for racial quotas using data from the Chicago Public Schools (CPS). We develop a theoretical framework that motivates quantifying the efficiency cost of race-neutral policies by the extent admissions decisions are distorted more than needed to achieve a given level of diversity. According to this metric, CPS's race-neutral system is only 20-25% efficient as a tool for increasing minority representation at CPS's top two exam schools. Even though CPS's system is based on socioeconomic disadvantage, it also turns out to also be less effective than race-based policies at increasing the number of low-income students. We examine several alternative race-neutral policies and find some that are more efficient than the CPS policy, but none that avoid substantial efficiency costs.

Matthew Gentzkow (Stanford)
Partisan language

Abstract

We study trends in the partisanship of Congressional speech from 1873 to 2009. We define partisanship to be the ease with which an observer could infer a congressperson's party from a fixed amount of speech, and we estimate it using a structural choice model and methods from machine learning. The estimates reveal that partisanship is far greater today than at any point in the past. Partisanship was low and roughly constant from 1873 to the early 1990s, then increased dramatically in subsequent years. Evidence suggests innovation in political persuasion beginning with the Contract with America, possibly reinforced by changes in the media environment, as a likely cause. Naive estimates of partisanship are subject to a severe finite-sample bias and imply substantially different conclusions.

Renato Gomes (TSE)
Drip Pricing

Abstract

Drip pricing is a technique used by online retailers of goods and services whereby a retail price is advertised at the beginning of the purchase process, following which additional fees, only avoidable at a cost to consumers, are then incrementally disclosed (or "dripped"). One important example is that of credit card surcharging, which has been the subject of intense regulatory debates in the last decade. In this paper we present a theory of drip pricing centered at the notion of missed sales. Our analysis identifies two kinds of missed sales: ex-ante (which occur when consumers who would buy the good at the checkout point refrain from engaging in the purchase process by fear of price drips) and ex-post (which occur when consumers at the checkout point refrain from buying the basic good due to the inconvenience of avoiding drips). Because of missed sales, we show that retail prices signal price drips: the higher the retail price, the less greedy the firm is on dripping extra fees. We explain how equilibrium retail and drip prices change with the level of consumer information, and derive optimal regulation. We then apply the positive and normative implications our model to the case of card surcharging.

Bruno Jullien (TSE)
Privacy Protection

Abstract

We study the incentives of a website to sell its customers' personal information. Third parties buying that information can either benefit or harm consumers, who learn about their vulnerability to unwanted intrusions through experience. The cost to the website of selling the information is the risk that bad experience may cause consumers to end their relationship with the website. The measures adopted by the website to mitigate that cost are neither contractible nor discernible by consumers. Nevertheless, in equilibrium, the website has incentives to be cautious about selling information or spend resources to screen third parties. We characterize the equilibrium privacy policy of the website and its welfare properties, and discuss the difficulty of welfare-improving regulations.

Josh Lerner (Harvard)
Observations about standard setting in practice

Abstract

For much of the past 150 years, firms and governments have struggled with the proliferation of intellectual property rights covering key technologies. Numerous commentators have suggested that such “patent thickets” can have socially detrimental consequences: overlapping intellectual property rights make it expensive for final good producers to commercialize innovative products and difficult for entrepreneurs to enter new markets. These worries have intensified with the explosion of litigation that has occurred globally over the past three decades.

Patent pools and standard-setting bodies—both of which encompass formal or informal organizations where owners of intellectual property share patent rights with each other and third parties--have been proposed as ways in which firms can address this “patent thicket” problem. Neither are new developments: each has a history that spans over one hundred years, and each has continually raised questions similar to the ones businesses, regulators, and others face currently. Patent pools, agreements in which patent holders license their collective intellectual property to each other or independent actors, in the mid-nineteenth century offered a chance to avoid damaging litigation, while affecting innovation, competition, and technological change in uncertain ways. Standard-setting organizations, or groups who focus primarily on producing or coordinating technical standards, stimulated decades of industrial coordination and manufacturing efficiency, and also highlighted the difficulties of formalizing methods of reaching consensus, disseminating standards, and representing interested parties. Both institutions have faced questions that have persisted through time, as they have grappled with the inherent challenges associated with these institutions.

Indeed, these institutions seem to be experiencing a dramatic revival: for example, an estimate suggests that sales in 2001 of devices based in whole or in part on pooled patents reached at least \$100 billion. Moreover, in the past few years, the interest in these approaches is spreading; for instance, the biomedical research community has expressed a keen interest in the development of patent pools for biomarkers for cancer, patents relative to HIV/AIDS and SARS, as well as for biotechnologies applied to agriculture and animal cloning. We may well be returning to a situation similar to the early days of 20th century, when many (if not most) important manufacturing industries had an arrangement with one or more pools or standard-setting organizations.

But despite their robust revival, these practices are not as well documented as one might expect. While certain court cases have attracted extensive discussion, much less work has codified the day-to-day working of these organizations. To better understand these issues, we have, in the course of our research into standards and patent pools spanning the past decade, held many dozens of extensive conversations with corporate executives, lawyers, pool and standards setting organization leaders, and regulators. This paper—drawn from the opening chapter of Jean Tirole’s and my new book on knowledge-sharing organizations-- will focus on how these institutions work today in practice.

Jonathan Levin (Stanford)
Platform pricing for ride sharing

Abstract

We study optimal pricing for ride-sharing and transportation platforms. We show that efficient allocation requires a subsidy, and that imposing budget balance leads to a mix of externality and elasticity-based pricing. We characterize the optimal platform response to market growth, to differing geographies and to varying consumer preferences over price versus wait time. We also analyze optimal dynamic pricing, and explain why competitive price determination inevitably leads to under-supply.

Heidi Williams (MIT)
Patents and follow on Innovation

Abstract

We investigate whether patents on human genes have affected follow-on scientific research and product development. Using administrative data on successful and unsuccessful patent applications submitted to the US Patent and Trademark Office, we link the exact gene sequences claimed in each application with data measuring follow-on scientific research and commercial investments. Using this data, we document novel evidence of selection into patenting: patented genes appear more valuable — prior to being patented — than non-patented genes. This evidence of selection motivates two quasi-experimental approaches, both of which suggest that on average gene patents have had no effect on follow-on innovation.