Research Statement

Daniel L. Chen

December 2017

1 Introduction

Economics and law is divided between the consequentialist view that optimal policy should be based on calculations of costs and benefits and a non-consequentialist view that policy should be determined deontologically: from duties we derive what is the correct law–what is right and just. Are there deontological motivations, and if there are, how might one formally model these motivations? What are the implications of things like deontological motivations for economics and policy, and what puzzles can be explained with deontological motivations that one cannot with standard models? And how do deontological motivations interact with consequentialist ones in the production of justice?

Answering questions like these has resulted in oTree—a programming language to study behavior in experiments (now used in 23 countries, 10 academic disciplines, private and public sectors, and high schools)—and the digital curation of 12 terabytes of archival and administrative data on judges and courts where normative commitments incubate. These tools bridge machine learning, causal inference, and normative theories of justice regarding equal treatment before the law and equality based on recognition of difference.

Ongoing research themes on consequences, formation, and measurement of normative commitments in law:

- **Law and Development**—tracing the incentives that led to what are now viewed as human rights violations
- **Markets and Morality**—how market forces interact with normative commitments
- **Behavioral Judging**—social and psychological, economic and political influences on legal ideas and production of justice
- **Law and Legitimacy**—role of legitimacy in legal compliance
- **Demography of Ideas**—economics of interpretation as a source of normative commitments

The research has also anchored successful grant applications totaling € 3 200 000 since 2013 from the European Research Council (Consolidator Grant for researchers 8-14 years from PhD), Swiss National Science Foundation, and Agence Nationale de la Recherche, and € 800 000 prior to 2013 from MacArthur Foundation, Ewing Marion Kauffman Foundation, Social Science Research Council, Templeton Foundation, Earhart Foundation, Institute for Humane Studies, National Institute of Child Health and Human Development, National Institutes of Health, and National Science Foundation.

The research papers have been accepted in leading economics journals (*American Economic Review, Econometrica, Quarterly Journal of Economics,* and *Journal of Political Economy*), leading computer science journals (*Journal of Machine Learning Research*), and peer-review law outlets (*Journal of Law and Economics*, Law and STEM Junior Faculty Forum, and Stanford-Yale Junior Faculty Forum).

2 Background

I completed a MIT Economics PhD in 2004 under Esther Duflo—a trailblazer in using experimental evidence for economic policymaking—and a Harvard Law JD in 2009. After a tenure-track assistant professorship in law, economics, and public policy at Duke University (USA), I received tenure at ETH Zurich (Switzerland) in 2012 and co-founded the ETH Center for Law and Economics, where I held the Chair of Law and Economics with endowed funding of € 660 000 per year and start-up of € 580 000. In 2015, I came to Toulouse School of Economics (France) and its Institute for Advanced Study, which prioritizes five themes related to my work—law and economics, political economy, behavioral economics, economic history, and economics and biology.

Much of my work uses economics methodology to study legal concepts. Most of my publications are in economics journals, with recent invitations to resubmit at *American Economic Review, Review of Economic Studies, American Economic Journal: Microeconomics, Review of Economics and Statistics,* and *Economic Journal*. Two articles are under review at *Journal of Political Economy*, for over 1 and 2 years, respectively. Other invitations to resubmit include leading outlets in machine learning, *Journal of Artificial Intelligence and the Law*, and in law, *Journal of Legal Studies*.

In addition to my appointments at Toulouse, I am a regular visitor at NBER (for access to U.S. Census data), director of oTree Open Source Research Association (an open-science CERN for social science experiments), and founder of Data Science Justice Collaboratory (faculty associate at Harvard Law School Berkman Klein Center for Internet and Society and project advisor at the NYU Courant Institute of Mathematical Sciences Center for Data Science). I have also served on the program committees of the European Economic Association, American Law and Economics Association, and Society for Institutional and Organizational Economics, made 500 presentations at universities and peer-reviewed conferences, and instructed spring/summer schools in law and economics, organizational, and behavioral economics.

In what follows, I present abstracts of my past work organized by theme—I like to think about the themes in terms of book projects. Submissions have been requested from *Oxford (2), Yale, MIT,* and *Virginia University Press.* Due to space constraints, the abstracts can be clicked to link to the paper or book-length manuscripts online.

3 Outline

My work in law and development traces the incentives that lead to human rights violations, such as religious, ideological, and gender-based violence. I examine how acts of violence derived their legitimacy from proto-legal sources and ultimately undermined that legitimacy. In **theorizing cultural differences: the economics of fundamentalism**, I document economic forces underlying the religious provision of social insurance, social sanctions, and social conservatism and, turning to one dimension of social conservatism, the economic incentives that give rise to gender violence, sexual harassment, and the regulation of the private domain. In doing so, I offer a theory for why and where church-state separation arose and develop high-dimensional methodologies for **measuring the consequences of normative commitments** using the random assignment of judges.

My work on markets and morality draws on experience with randomized control trials in developing countries to conduct a series of experiments in institutional and mechanism design that blur lab and field with the use of disaggregated labor markets. Observational data offers one lens on human behavior with which to evaluate legal rules. Experimental data offers another, but manipulating legal rules in field settings is practically infeasible. **Do free markets corrode moral values**? I test behavioral assumptions underlying the perceived dichotomy between law and economics approaches to law—morality vs. markets, deontological vs. consequentialist, expressive vs. incentives, subjective vs. objective, and duty vs. damages. On a theoretical level, **distinguishing sacred values from social preferences: theory, evidence, and relevance of deontological motivations** shows how the

presence of these commitments can problematize the random lottery incentive and strategy method widely used in economics and fundamentally challenge theoretical models that assume a convex—rather than concave—cost of deviating from moral and ideological bliss points. On a practical level, modularizing the code used for these studies led to the development of oTree—an open-source programming language for others to program their own lab, online, or field experiments—enhancing the ability to study normative economics, empirical moral psychology, experimental philosophy, and interactive epistemology in contextualized settings, and enabling researchers to explore the transmission and persuasion process of normative commitments whose incommensurability can lead to conflict and violence. oTree is presently used in an EU Horizon 2020 project to study behavior in large-scale networks, voting, macroeconomics, and mixed agent-based experiments.

My work on behavioral judging interrogates the **impossibility of objective judgments: priming, gambler's fallacy, mood, voice, and peer effects in courts**. I have digitally curated 12 terabytes of archival and administrative data on judges and courts where normative ideas incubate, textually and orally. I use these data to study the social and psychological, economic and political influences in **judicial analytics and behavioral foundations of polarization**. Two singular initiatives are: 1) the digital universe for a century-and-a-half of roughly a half-million cases in the U.S. Courts of Appeals, engineered 2 billion N-grams up to eight words in length, and constructed the networks of citations across cases and seating among judges, who are randomly assigned to panels of three. These cases are linked to the universe of publicly available datasets on the U.S. Supreme and District Courts, Administrative Office of the U.S. Courts, and hand-collected features in a 5% random sample and 2) digitized speech formants in 2500 hours of U.S. Supreme Court oral arguments for over a half-century—longitudinal data on speech intonation (linguistic turns) are rare—and manually clipped oral advocates' identical introductory sentences—speech variation holding words fixed is rare in high-stakes settings. I also curated the universe of 1 million decisions and 15 million hearing sessions in U.S. immigration courts, the universe of 1 million criminal sentencing decisions in U.S. District Courts, and the universe of judges' biographies that are the potential source of their normative commitments. An additional 1000 legal databases at all levels of federal and state government are tagged and linked.

My work on law and legitimacy formalizes recognition-respect theory and what it means for legal institutions, lawmakers, judges—to be indifferent, such that it violates our notion of justice. First, I examine whether legal sanctions can **deter or spur? british executions during world war i**. I digitize and link eleven distinct World War I British archival datasets comprising the universe of deserters reported in military diaries, police gazettes, and handwritten military trials, commuted and executed capital sentences, geocoded casualties, maps, officer lists, and the military units' order of battle. I document differential effects of executing British and Irish soldiers on their subsequent desertions and the practice of endogenous justice. Second, I introduce the idea of **revealed preference indifference**. I distinguish sympathy from empathy and link the latter to recognition, knowledge and respect of reference points. Information acquisition can be endogenous to preferences, which blurs the boundary between Becker's statistical and taste-based discrimination. Psychologists find many effects of moderate sizes in the lab, so settings where people are closer to indifference among options are more likely to lead to detectable effects outside of it. One way to know if a judge is indifferent, is by observing where behavioral biases arise. Their decisions become less predictable, holding fixed the predictions. Third, I illustrate these concepts using asylum judges and criminal prosecutors, who handle fifteen times the number of cases that judges do, but whose role is largely with little accountability. I do so with a unique dataset recording a half million defendants for a decade, vertically linked from the police arrest to the potential final sentence (vertical linkages from the time of arrest, including those sent home without a trial, otherwise do not exist).

My work on the demography of ideas traces the **geneology of ideology** in the corpus of judicial decisions and micro-level data on physician decision-making. I investigate the impact of corporate sponsored economics training on moral reasoning and criminal justice (sentencing harshness and disparate impacts become more justifiable under

theories of deterrence and statistical discrimination). I investigate the effects of pharmaceutical company payments to physicians on their prescriptions, patient outcomes, and patient adherence. I link administrative, unencrypted (through special permission) Medicare data to industry-physician relationships cleaned from litigation settlements prior to the Affordable Care Act and from the universe of payments reported after the Affordable Care Act and linked to NDC drugcode—a combined total of roughly 30 million payments—and the p-curve of underlying scientific papers linked by NDC drug code to price the lives cost by the scientific reproducibility crisis. The development of institutions is critical to preventing tragedies that stem from lack of choice, be the tragedies ones of violence, poverty, infectious diseases, or lack of health and education. But incentives to help members of shared identities frequently debilitate these institutions. On a thick vein, I study the macro and demographic forces and laws that aid or hinder discrimination in legal institutions, markets, and public policy.

I have broadly been motivated by an idea that can be termed, **hermemetrics: the measurement of meaning**, which is hermeneutics and econometrics or economics of interpretation, as a source of normative commitments. Hermemetrics examples include: (a) modeling the economic incentives behind the shift from pro-welfare religious interpretations a century ago to the anti-welfare posture of today's religious right and why church-state separation arose when and where it did, (b) measuring the effects of sexual harassment law—an interpretation of anti-discrimination law—on gender inequality by exploiting the random assignment of appellate judges interpreting the facts and the law differently but in a statistically predictable manner, (c) using a particular instance of interpretive injustice where British capital cases during World War I were randomly executed or commuted to estimate the deterrent and delegitimizing effects of the death penalty, (d) studying how individuals misinterpret the law by manipulating tax schedules and contracts in labor market settings, (e) understanding why and when lay interpretation of legal decisions have expressive or backlash effects, and (f) revelation of deontic motivations through revealed preference and in interpretations of legal and social texts. If economics is the science of how individuals optimize to budget constraints, hermemetrics studies how societies optimize in response to textual constraints.

4 Law and Development

4.1 Theorizing Cultural Differences: The Economics of Fundamentalism

One of the most influential views of our time attributes a large part of the failure of development in the post-war period to group conflicts. Recent research in development economics has identified a large collection of policy innovations that would help the poor. But these policies often do not get adopted because of conflicts between groups. Researchers have traditionally focused on the number of groups that are in conflict with each other. Moreover, why social conflict occurs along ethnic-religious lines instead of class lines is a subject of much debate.

1. Club Goods and Group Identity: Evidence from Islamic Resurgence During the Indonesian Financial Crisis (*Journal of Political Economy*, 118(2), 2010) tests a model in which group identity in the form of religious intensity functions as ex post insurance. I exploit relative price shocks induced by the Indonesian financial crisis to demonstrate a causal relationship between economic distress and religious intensity (Koran study and Islamic school attendance) that is weaker for other forms of group identity. Consistent with ex post insurance, credit availability reduces the effect of economic distress on religious intensity, religious intensity alleviates credit constraints, and religious institutions smooth consumption shocks across households and within households, particularly for those who were less religious before the crisis.

2. Islamic Resurgence and Social Violence During the Indonesian Financial Crisis (*Institutions and Norms in Economic Development*, MIT Press, ed. M. Gradstein and K. Konrad, 179-200, 2007) also focuses on the intensity with which people identify with their groups. Violence is a negative externality with enormous social costs, so to the extent group identity and social violence (physical acts of destruction, killing, looting, attacks,

burning, clashes, taking hostages, etc., by one group against another) are related, policies taking into account intensity of group identity need to be considered. This paper documents a correlation between group identity and group conflict in the specific context of Islamic resurgence during the Indonesian financial crisis.

3. Economic Distress Stimulates Religious Fundamentalism presents a causal analysis. It exploits relative price shocks induced by the 1997 Indonesian financial crisis and variation in religious institutions across Indonesia before the crisis to identify the effect of economic distress on the relationship between religious institutions and social violence. In the cross-section, high religious intensity areas before the crisis have more social violence after the crisis. In the panel, social violence increases fastest where participation in Koran study also increases the fastest. In the two-stage least squares, instrumenting for economic distress using relative price shocks shows a causal relationship between economic distress and the relationship between religious intensity and social violence. These results are unlikely to be driven by omitted environmental variables. Credit availability mitigates this effect. Economic distress alone did not stimulate social violence but stimulates it in the presence of religious institutions. I explain these findings in a model where high marginal utilities during economic distress increase incentives for group conflict where group conflict increases the budget of insurance groups. With volatility, religions with stronger sanctions or more out-group bias are more stable and successful. As volatility declines, benign groups and religions become relatively successful.

4. Gender Violence and the Price of Virginity: Theory and Evidence of Incomplete Marriage Contracts builds and tests a model of marriage as an incomplete contract that arises from asymmetric virginity premiums and examines whether this can lead to social inefficiencies. Contrary to the efficient households hypothesis, women cannot prevent being appropriated by men once they enter marriage if they command lower marriage market opportunities upon divorce. Because men cannot or do not commit to compensating women for their lower ex post marriage market opportunities, marriage is an incomplete contract. Men may seek to lower women's ex ante "market wages" in order to induce entry into joint production. Inefficient or abusive marriages are less likely to separate. Equalizing virginity premiums may reduce domestic and non-domestic violence. Female circumcision and prices women pay doctors to appear virgin before marriage in many countries suggest asymmetric virginity premiums continue to exist. Evidence from China and the U.S. suggest asymmetric virginity premiums persist over economic development. Asymmetric virginity premiums are strongly positively correlated with female but not male virginity premiums. I use variation in religious upbringing to help estimate the effect of virginity premiums on gender violence in the U.S. The OLS relationship between virginity premiums and female reports of forced sex may be biased downwards if shame is associated with abuse and this shame is greater for women with higher virginity premiums. But the OLS relationship for males might not be biased downwards. Asymmetric virginity premiums are positively correlated with men forcing sex on women and paying women for sex. The model complements a growing empirical literature on inefficient households and human rights abuses, visible manifestations of female appropriability across time and space.

5. Why are religious groups with greater within-group charitable giving more socially conservative and opposed to the welfare state? The Political Economy of Beliefs: Why Fiscal and Social Conservatives/Liberals (Sometimes) Come Hand-in-Hand (TSE Working Paper No. 16-722; J. Lind) proposes and tests a theory where religious provision of social insurance explains why fiscal and social conservatism align. Religious groups that provide social insurance act as competitors to state-run insurance, and therefore resist expansion of the state into the provision of social insurance. The alignment between fiscal and social conservatism disappears when there is a state church and reverses for members of a state church. Moreover, plausibly exogenous increases in church-state separation precede increases in the alignment between fiscal and social conservatism, which mitigates concerns of the reversal being due to omitted environmental variables. The theory provides a novel explanation for religious history: as elites gain access to alternative social insurance, they judiciate increas-

ing church-state separation to create a constituency for lower taxes. This holds if religious voters exceed non-religious voters, otherwise, elites prefer less church-state separation in order to curb the secular left, generating multiple steady states where some countries sustain high church-state separation, high religiosity, and low welfare state, and vice versa. Religion, Welfare Politics, and Church-State Separation (*Journal of Ecumenical Studies*, 42(1), 42-52, 2007; J. Lind) uses this framework to explain the changing nature of religious movements, from Social Gospel to the Religious Right, and why church-state separation arose in the U.S. but not in many European countries.

4.2 Measuring the Consequences of Normative Commitments

My earlier work in law and development traced the incentives that led to what have now come to be viewed as human rights violations. Subsequent work explored the consequences of normative commitments exploiting the random assignment of normative interpreters. Measuring the Moral and Economic Consequences of Judicial Discretion is a book-length treatment, which proposes rapid impact analyses to aid judges deliberating hard cases. The current method of judicial decision-making relies on informal policy analysis and such interpretive discretion comes with bias. Recent econometric advances make possible methodologies for empirically evaluating the effects of law using the random assignment of judges interpreting the facts and the law in different ways. The first paper developing this method examined the empirical consequences of feminist legal theory introducing the concept of sexual harassment into American jurisprudence in the 1970s; over the next few decades, sex discrimination laws were reinterpreted to include sexual harassment. The second paper further developed the methodology–with theoretical econometricians–into a framework with broad applicability. The third paper extends the legal analysis. The next four papers ask if laws affect others' normative commitments. The final four papers illustrate the need for random assignment, present evidence for the exclusion restriction, document lay attention to legal decisions, and develop causal impacts of legal precedent on judges' decisions.

1. Sexual harassment is perceived to be a major impediment to female labor force participation. Insiders, Outsiders, and Involuntary Unemployment: Sexual Harassment Exacerbates Gender Inequality (*Review of Economics and Statistics*, resubmitted; TSE Working Paper No. 16-687; J. Sethi) uses the random assignment of U.S. federal judges setting geographically-local precedent, and the fact that judges' biographies predict decisions in sexual harassment cases, to document the causal impact of forbidding sexual harassment. Consistent with an insider-outsider theory of involuntary unemployment, but in contrast to a mandated benefits theory of employment protections, pro-plaintiff sexual harassment precedent spurred the adoption of sexual harassment human resources policies, encouraged entry of outsiders, and reduced gender inequality in labor supply and wages among the population. These effects were comparable to the effects of the Equal Employment Opportunity Act and greatest in the construction industry, which was heavily affected by sexual harassment litigation.

2. Sparse Models and Methods for Optimal Instruments with an Application to Eminent Domain (*Econometrica*, 80(6), 2369-2429, 2012; A. Belloni, V. Chernozhukov, C. Hansen) develops results for the use of Lasso and post-Lasso methods to form first-stage predictions and estimate optimal instruments in linear instrumental variables (IV) models with many instruments, p. Our results apply even when p is much larger than the sample size, n. We show that the IV estimator based on using Lasso or post-Lasso in the first stage is root-n consistent and asymptotically normal when the first stage is approximately sparse, that is, when the conditional expectation of the endogenous variables given the instruments can be well-approximated by a relatively small set of variables whose identities may be unknown. We also show that the estimator is semi-parametrically efficient when the structural error is homoscedastic. Notably, our results allow for imperfect model selection, and do not rely upon the unrealistic "beta-min" conditions that are widely used to establish validity of inference following

model selection. In simulation experiments, the Lasso-based IV estimator with a data-driven penalty performs well compared to recently advocated many-instrument robust procedures. In an empirical example dealing with the effect of judicial eminent domain decisions on economic outcomes, the Lasso-based IV estimator outperforms an intuitive benchmark. Optimal instruments are conditional expectations. In developing the IV results, we establish a series of new results for Lasso and post-Lasso estimators of nonparametric conditional expectation functions which are of independent theoretical and practical interest. We construct a modification of Lasso designed to deal with non-Gaussian, heteroscedastic disturbances that uses a data-weighted l1-penalty function. By innovatively using moderate deviation theory for self-normalized sums, we provide convergence rates for the resulting Lasso and post-Lasso estimators that are as sharp as the corresponding rates in the homoscedastic Gaussian case under the condition that $\log p = o(n1/3)$. We also provide a data-driven method for choosing the penalty level that must be specified in obtaining Lasso and post-Lasso estimates and establish its asymptotic validity under non-Gaussian, heteroscedastic disturbances.

3. Is it justified for states to appropriate private property rights? If so, should governments expropriate or regulate? Government Expropriation Increases Economic Growth and Racial Inequality: Evidence from Eminent Domain (*Economic Journal*, invited to resubmit; TSE Working Paper No. 16-693; S. Yeh) tests three conventional views: insecure property rights cause underinvestment, moral hazard cause overinvestment, or public use cause economic growth. We embed these mechanisms in a model and measure them using the random assignment of U.S. federal court judges setting geographically-local precedent. For a half-century, racial minority Democrats were more likely to strike down government appropriations while Republican former federal prosecutors were more likely to uphold them. We find that pro-government physical takings precedent stimulated subsequent takings, expropriation of larger parcels, highway construction, and growth in construction, transportation, and government sectors as well as agriculture, retail, and financial sectors, overall economic growth, and property values. However, racial minorities were increasingly displaced, unemployed, and living in public housing, and the service sector declined. Pro-government regulatory takings precedent also spurred economic growth and property values, but did not increase displacement or racial inequality.

4. When do policies generate expressive or backlash effects? Whether policies shift preferences is relevant to policy design. Recent economic models suggest that where a proscribed activity is prevalent, permissive laws liberalize attitudes toward partakers while increasing utility. The opposite occurs in communities where the proscribed activity is rare. How Do Rights Revolutions Occur? Free Speech and the First Amendment (TSE Working Paper No. 16-705; S. Yeh) tests a model of law and norms using an area of law where economic incentives are arguably not the prime drivers of social change. From 1958–2008, Democratic judges were more likely than Republicans to favor progressive free speech standards. Using the random assignment of U.S. federal court judges setting geographically-local precedent, we estimate that progressive free speech standards liberalized sexual attitudes and behaviors and increased both crime rates and the spread of sexually transmitted diseases. We then randomly allocated data entry workers to enter newsarticles of court decisions. Progressive decisions liberalized sexual attitudes and shifted norm perceptions even for data entry subjects, but not self-reported behavior. These results present evidence of law's expressive power–with fundamental implications for decision making in social and political settings and for the empirical predictions of theoretical models in these domains.

5. To build on the previous study and the field data using random assignment of judges, The Construction of Morals (*Journal of Economic Behavior and Organization*, 104, 84-105, 2014; S. Yeh) reports additional contextualized field experiments. We randomize data entry workers to transcribe newspaper summaries of liberal or conservative court decisions about obscenity. As before, we find that liberal obscenity decisions liberalize individual and perceived community standards and increase utility. Yet we also find that religious

workers become more conservative in their values, identify as more Republican, view community standards as becoming more liberal, and report lower utility. Workers update beliefs about the prevalence of sexual activities differently in response to liberal or conservative decisions. These results provide causal evidence for the law having indirect social effects that may amplify or attenuate deterrence effects and suggest that legitimacy of law can affect utility and self-identification.

6. The previous framework relies on a model of law having backlash *or* expressive effects, but the next paper develops a theoretical model where law has backlash, *then* expressive effects. Policies Affect Preferences: Evidence from Random Variation in Abortion Jurisprudence (*Journal of Political Economy*, under review, 12/15; TSE Working Paper No. 16-723; V. Levonyan, S. Yeh) exploits the random assignment of U.S. federal judges creating geographically local precedent and the fact that judges' politics, religion, and race predict decision-making in abortion jurisprudence. Instrumenting for abortion jurisprudence with exogenous judicial characteristics, we estimate the impact of abortion jurisprudence on state laws, campaign donations, and abortion attitudes. We verify information transmission in that pro-life abortion jurisprudence caused restrictive state laws and increased campaign donations to pro-choice causes. Pro-choice abortion decisions shifted preferences against legalized abortion in the short-run, but in the longer-run, abortion views followed court decisions. Pro-choice decisions affected Republicans while pro-life decisions affected Democrats. Counterfactual exercises suggest that had abortion cases in the last half-century been decided the opposite way, the increase in pro-life attitudes among Republicans would have been steeper and Democrats would have been more pro-choice. Our estimates complement a historical narrative that turning to the courts to vindicate rights often led to resistance and subsequent acceptance and we present a model consistent with these facts.

7. Law and Norms: A Machine Learning Approach to Predicting Attitudes Towards Abortion (K. Kwan, M. Maass, L. Ortiz) uses U.S. Courts of Appeals cases, sociological and attitudinal indicators, criminal statistics, and a variety of survey data to predict societal attitudes towards abortion. Understanding the predictors of societal attitudes has been widely investigated using individuals' responses to surveys and polls. We create two classification models: proabortion attitudes for health reasons related to the mother, fetus, or rape and proabortion attitudes for any other personal nonhealth related reason. To address high dimensionality, we employ factor analysis to group indicators. Logistic regression and random forests performed best among three types of classifiers evaluated on AUC and accuracy. For pro abortion attitudes related to health related abortions, the most important factor contains sociological and attitudinal indicators about the frequency of contact with family and friends. The most important U.S. Courts of Appeals indicators include the religion, political affiliation, and ABA ratings of the judges in the Circuit pool. The most important crime indicators are the rates of violent and property crimes. For pro abortion attitudes towards nonhealth related abortions, the most important factor contains sociological and attitudinal indicators about job satisfaction, religious preference, religion raised in, and beliefs about the Bible. The most important U.S. Courts of Appeals indicators include the religion and political affiliation of the judges in the Circuit panel. These legal indicators are grouped in the same factor as the most important crime indicators, which are rates of crimes against society, crimes against property, and violent crimes.

8. Distinguishing Between Custom and Law: Empirical Examples of Endogeneity in Property and First Amendment Precedents (*William & Mary Bill of Rights Journal*, 21(1081), 2013, with S. Yeh) discusses the relationship between custom and law to highlight the phenomenon of endogeneity when empirically evaluating the effects of laws. An important literature evaluates the roles of laws in motivating behaviors, including investigations of whether or how laws influence customs and social norms. Traditional economic analysis, for example, posits that official or codified laws are expected to influence behaviors by formally incentivizing a particular action or communicating values. Enhancing this strand of thought, there is a growing contribution

from historical and empirical analyses that link laws to broader societal changes over time. At the same time, a valuable discourse examines how customs may determine both de facto laws and formally enacted laws, including the court precedents that are rendered. Whether they are directly codified into a legal test or informally referenced, customs can influence the formal laws that are adopted in a community and beyond. Indeed, some scholars have argued that evolving customs and norms have influenced the Court in its decisions. The subsequent effects of these formal laws and court decisions are of tremendous interest to policymakers and judges, and with policy concerns in mind, we argue that one must not ignore the endogenous feedback between aggregate behaviors, customs, and laws. That is, while customs may shape or influence laws, laws can also shape customs through their effects on behaviors or norms in the aggregate. The endogeneity that custom produces suggests that simply by observing a correlation between law and behavior is not enough to assert that a law in itself is effective or to assert that social trends and evolving customs are driving legal change. We propose a greater role for experimental or quasi-experimental evidence among scholars interested in the empirical study of law and sketch out an original empirical strategy that could overcome the endogeneity between custom and law,

9. Do judges directly affect economic outcomes at the time they are revealed? The Shareholder Wealth Effects of Delaware Litigation (*American Law and Economics Review*, forthcoming; TSE Working Paper No. 16-683; A. Badawi) collects data on the record of every action in hundreds of derivative cases and merger class actions involving public companies filed in the Delaware Court of Chancery from 2004 to 2011. We use these data to estimate how markets respond to litigation in the most important court for corporate disputes in the United States. The detail in the dataset allows us to explore how case characteristics such as the timing of case filing, the presence of certain procedural motions, litigation intensity, and judge identity relate to firm value. Unlike previous studies, we document that that negative abnormal returns are associated with the filing of derivative cases, and that association is particularly strong for cases that are first filed in Delaware and are not related to a previously disclosed government investigation. We also develop some evidence that market participants can anticipate litigation intensity and respond by valuing the firm equity less. Finally, we find little evidence of abnormal returns associated with judicial assignment at the time of filing of derivative cases, but we do observe an association between judicial assignment and case filing for merger cases.

10. Do people pay attention to court outcomes? Does Appellate Precedent Matter? Stock Price Responses to Appellate Court Decisions of FCC Actions (in Empirical Legal Analysis: Assessing the Performance of Legal Institutions, 2013, with S. Yeh and A. Araiza) tests the effects of federal appellate court decisions of Federal Communications Commission (FCC) actions on stock prices using differences-in-differences (DID) and an instrumental variables approach. This study exploits the random assignment of appellate judges to three-judge panels and the fact that a judge's (1) party affiliation and (2) race predict outcomes in appellate court decisions of FCC actions to instrument for anti-industry decisions, which favor the public interest. This instrumental variable approach demonstrates a causal relationship between appellate court decisions of FCC actions and changes in stock prices of media firms relative to the stock prices of non-media firms. The differences-in-differences (DID) analysis shows that federal appellate court decisions against media businesses decreased media stock prices. The instrumental variables analysis shows that these appellate court decisions decreased media stock prices relative to non-media stock prices, one and two years after the court decisions. Recent studies indicate that stock prices serve as a proxy for competition and that decreased media competition may correspond to an increase in variety of programming. These findings suggest that when deciding against media businesses, the courts effectively reinforced the purpose of the FCC to serve the public interest by promoting a diversity of viewpoints.

11. Do judges follow legal precedent? Public enforcement of law relies on the use of public agents, such as judges,

to follow the law. Are judges motivated only by strategic interests and ideology, as many models posit, rather than a duty to follow the law? Judicial Compliance in District Courts (TSE Working Paper No. 16-715; J. Frankenreiter, S. Yeh) uses the random assignment of U.S. Federal judges setting geographically-local precedent to document the causal impact of court decisions in a hierarchical legal system. We examine lower court cases filed before and resolved after higher court decisions and find that lower courts are 29-37% points more likely to rule in the manner of the higher court. The results obtain when the higher court case was decided in the same doctrinal area as the pending case and when the higher court case was decided on the merits. Reversals by the higher court have no significant effects. These results provide clean evidence that judges are motivated to follow the law and are not solely motivated by policy preferences.

12. The final paper returns to the motivating question in this section, the economics of fundamentalism. Religious Freedoms, Church-State Separation, and Religiosity: Evidence from Randomly Assigned Judges (E. Ash) provides evidence on how laws related to the establishment or exercise of religion affect the level of religious activity. We exploit the random assignment of judges to regional appellate courts, along with the fact that judge politics, religion, and other characteristics predict decision-making in church-state jurisprudence, to obtain exogenous variation in religion caselaw. Our first-stage effect is driven by the fact that Democratic, minority-religion judges tend to favor greater separation of church and state than their Republican, majority-religion counterparts. In the second-stage estimates for the effect of law on outcomes, we report two findings. First, judicial decisions reducing government subsidies to mainstream religions reduce measures of supply-side religious services. Second, judicial decisions that strengthen the free exercise of fringe religions increase self-reported religiosity.

5 Markets and Morality

5.1 Do Free Markets Corrode Moral Values?

The previous section studied how market forces *interact* with normative commitments. This section studies the impact of market forces *on* normative commitments. At least since Adam Smith and David Hume, scholars have offered hypotheses about the effect of a citizen's economic experience on his or her moral life as an individual. It has been asserted that competition may bring a winner-take-all mentality and a lack of concern for others or that exposure to market values will lead us to abandon non-utilitarian forms of moral thought, treating every moral issue instead in terms of costs and benefits. (Whether this is viewed as a bad thing, of course, depends on one's attitude towards utilitarianism.) On the positive side, the proponents of the so-called *doux commerce* thesis (a theory popularized by 18th century political philosophers) have proposed that a competitive market, with its disruptive effect on geographical and tribal isolation, will actually have morally improving effects, increasing our care for and understanding of others. Of particular interest today are questions like: Could the structure of employment affect moral attitudes? What is the effect of pharmaceutical company payments to physicians, and the potential role for regulation, transparency, and accountability in medicine? What is the impact of third-party litigation funding and can the law be influenced by market forces? Returning to the mechanism of the *doux commerce* thesis, do markets exacerbate—or disrupt—tribal conflict?

1. Markets, Morality, and Economic Growth: Competition Affects Utilitarian Judgment (TSE Working Paper No. 16-692) studies the impact of employment structure on three normative issues: utilitarian versus deontological values, other-regarding preferences, and charitable donations. Through a labor market intermediary, I randomly assigned workers to competitive or piece-rate work conditions. The groups were given a moral question posing a conflict between utilitarian and deontological values, and offered a choice to make a charitable donation. The moral question was accompanied by an illustration that made salient outgroup

considerations. Several results emerge: Competitively structured work experiences increased deontological value choices, deontological commitments towards outgroup members, and donations by productive workers relative to non-productive workers; and the effects on deontological value choices differ by a country's level of economic development. I reconcile these results with a theoretical model based on experimental findings in affective moral psychology. When competition is perceived as unfair or unfamiliar, negative affect triggers deontological value choices, but when it is perceived as familiar or even fun, positive affect increases utilitarian attitudes. If utilitarian attitudes lead to market-oriented policies, multiple steady states arise where some countries sustain high utilitarianism, market-orientation, and economic growth, and vice versa. I use this perspective to help explain the intellectual history of the *doux commerce* thesis.

2. Mandatory Disclosure: Theory and Evidence from Industry-Physician Relationships (*Journal of Legal Studies*, 2nd revise and resubmit; TSE Working Paper No. 16-716; V. Levonyan, E. Reinhart, G. Taksler) presents and test a model of mandatory disclosure. The effects of disclosure laws on what is being disclosed are typically unknown since data on disclosed activity rarely exist in the absence of disclosure laws. We exploit data from legal settlements disclosing $316 million in pharmaceutical company payments to 316,622 physicians across the U.S. from 2009-2011. States were classified as having strong, weak, or no disclosure based on whether the data was reported only to state authorities (weak) or were publicly available (strong). Strong disclosure law was associated with reduced payments among doctors accepting less than $100 and increased payments among doctors accepting greater than $100. Weak disclosure states, despite imposing administrative compliance costs to industry, were indistinguishable from no disclosure states. This result suggests that the mechanism for fewer small payments in strong disclosure states was physicians' reduced willingness to accept payments rather than the imposition of significant administrative costs on industry. We conduct additional analysis holding fixed the cost for pharmaceutical companies of disclosing data, which was possible because Massachusetts began releasing payment data online during our sample period. Differences-in-differences analyses and multiple regression yield similar estimates for each payment category: Mandatory disclosure reduced payments for speaking and for meals but increased payments for consulting activities. Significant disclosure aversion reducing conflicts of interest is consistent with the policy goals of mandatory disclosure, though the increased payments among those receiving large payments may have been unintended.

3. The alienability of legal claims holds the promise of increasing access to justice and fostering development of law. While much theoretical work points to this possibility, no empirical work has investigated the claims, largely due to the rarity of trading in legal claims in modern systems of law. Can Markets Stimulate Rights? On the Alienability of Legal Claims (*RAND Journal of Economics*, 46(1), 23-65, 2015) develops a principal-agent framework where litigation funders provide expertise in reducing uncertainty in agents' disutility of production. The model leads to the counterintuitive prediction that litigation funders prefer cases with novel issues, and social surplus is positively correlated with legal uncertainty. Consistent with the model, court backlog, court expenditures, and a slowing in average time to completion are associated with third-party funding; cases with third-party funding receive more citations and are reversed less often than comparable cases without such arrangements. A Market for Justice: A First Empirical Look at Third Party Litigation Funding (*Journal of Business Law*, 15(3), 2013; D. Abrams) takes the first step toward empirically testing some of these theoretical claims using data from Australia. We find some evidence that third party funding corresponds to an increase in litigation and court caseloads. While third party funding appears to have effects on both the cases funded and the courts in jurisdictions where it is most heavily used, the overall welfare effects are still ambiguous.

4. Do Markets Overcome Repugnance? Muslim Trade Response to Anti-Muhammad Cartoons tests a theory of trade, trust, and conflict. After Danish newspapers published Anti-Muhammad cartoons—on the intensive

margin—exports from Danish countries to Muslim countries decreased by 23% for 20 months but rebounded by 24% thereafter. More secular Muslim countries were less affected by the controversy. Final goods (food, transport, and manufactured goods) were more affected, while intermediate goods (crude materials and commodities exports) were unaffected. The rebound virtually eliminated Muslim backlash to the obscenity deemed blasphemous from an Islamic perspective. Interestingly, Islamic exports to Denmark were unaffected—Danish money was acceptable—throughout this time period. On the extensive margin—the number of product categories trading—did not rebound. These results present a mixed picture on the costs of free speech. They are consistent with trade models with fixed costs preventing importers from easily shifting between countries, and consistent with markets sometimes eliminating the effect of moral attitudes against opposing ideologies.

5. What are the consequences of intermediating moral responsibility through complex organizations or transactions? Intermediated Social Preferences: Altruism in an Algorithmic Era (in Advances in the Economics of Religion, forthcoming, Palgrave, ed. J. P. Carvalho, S. Iyer, J. Rubin) examines individual decision-making when choices are known to be obfuscated under randomization. It reports the results of a data entry experiment in an online labor market. Individuals enter data, grade another individual's work, and decide to split a bonus. However, before they report their decision, they are randomized into settings with different degrees of intermediation. The key finding is that less generosity results when graders are told the split might implemented by a new procurement algorithm. Those whose decisions are averaged or randomly selected among a set of graders are more generous. These findings relate to "the great transformation" whereby moral mentalities are shaped by modes of (a)social interaction.

5.2 Distinguishing Sacred Values from Social Preferences: Theory, Evidence, and Relevance of Deontological Motivations

In the last few decades, there has been a gradual expansion of the domain of preferences considered by theory. There is the homo-oeconomics view – that people are only motivated by material consequences of their decisions on their own payoffs. Then, models were expanded to incorporate fairness into economics, for example, caring about the consequences of their decisions for others and what people think of my intentions. We consider another slice, that is, can people be motivated simply out of duty–purely internal consequences? In general, consequentialism and deontological motivations are very hard to distinguish. What we propose is a method using revealed preference to identify non-consequentialist motivations and a model of deontological motivations as lexicographic ("first, do your duty"). An approximation of lexicographic preferences is a concave cost of deviating from duty. In economics it is standard to model costs as being convex. This is often plausible, reflecting the idea that small deviations from a bliss point are rather negligible, while larger deviations may have big consequences. However, there is one domain in which experimental data seem to suggest that costs are in fact concave–the domain of ideology and morality. Popular terms like the "What the Hell Effect" and "Slippery Slope" capture exactly the perception that individuals are perfectionist in this domain, and they do not distinguish much between small and large deviations from their bliss points. The cognitive cost of wrongdoing in a large scale is not much larger than in a small scale, hence the need for perfectionists to be careful when considering even a small deviation from the "right" principle. Whether the costs of deviating from ideological or moral bliss points is convex or concave has implications for how to integrate different cultures and religion, and how to address extremism (rightwing, Islamists, etc.) in particular. Different sanctioning structures (be it interpersonal, informal, or coming from an authority) affect the choice of individuals with non-majoritarian ideologies to integrate. In particular, being harsh on small cultural deviations may backfire and create subcultures. This framework forms the basis for studying more detailed questions of collision between cultures and it provides a large number of implications that can be tested both empirically and experimentally.

1. Social Preferences or Sacred Values? Theory and Evidence of Deontological Motivations (TSE Working Paper No. 16-714; M. Schonger) proposes an experimental method that can detect an individual's deontological motivations by varying the probability of the decision-maker's decision having consequences. It uses two states of the world, one where the decision has consequences and one where it has none. We show that a purely consequentialist decision-maker whose preferences satisfy first-order stochastic dominance will choose the decision that leads to the best consequences regardless of the probability of the consequential state. A purely deontological decision-maker is also invariant to the probability. However, a mixed consequentialist-deontological decision-maker's choice changes with the probability. The direction of change gives insight into the location of the optimand for one's duty. We provide a formal interpretation of major moral philosophies and a revealed preference method to detect deontological motivations and discuss the relevance of the theory and method for economics and law.

2. A Theory of Experiments: Invariance of Equilibrium to the Strategy Method of Elicitation and Implications for Social Preferences (TSE Working Paper No. 16-724; M. Schonger) develops implications of deontological motivations for economics methodology. Most papers that employ the strategy method (SM) use many observations per subject to study responses to rare or off-equilibrium behavior that cannot be observed using direct elicitation (DE), but ignore that the strategic equivalence between SM and DE holds for the monetary payoff game but not the game participants actually play, which is in terms of utilities. To illustrate the severity of this issue, we formalize the mapping from the monetary payoff game to this actual game. A theorem provides necessary and sufficient conditions for strategic equivalence to apply. When the domain of preferences includes commonly-modeled motivations, such as intentions or disappointment aversion, or less-common ones, such as self-image or duty, strategic equivalence fails and thus the invariance to the method of elicitation does not apply. We use results from the past literature and our own experiments to investigate how well this theorem explains when results with SM and DE differ. We manipulate the salience of off-equilibrium considerations in our own experiments to demonstrate that SM and DE are not strategically equivalent, contrary to conventional wisdom. Three results emerge. First, not accounting for the bias in the estimation when decisions at one information set can influence the utility at another information set can render significant differences in decision-making. Second, the bias can be large and equivalent to some of the other causal effects being measured. Third, subtle interventions on salience can magnify these differences by a similar amount.

3. Who Cares? Measuring Attitude Strength in a Polarized Environment (C. Cavaille and K. Van der Straeten) builds on research in social psychology, we propose a model of survey response in which individuals' policy preferences are characterized by two parameters: their attitude on an issue, and their attitude strength. Strong attitudes are behaviorally relevant and stable over time while weak attitudes are easily manipulated with only limited behavioral consequences. We assume that the psychological cost to individuals of not reporting their attitude depends positively on the issue strength. We derive predictions about how respondents will answer survey questions under two different survey techniques, Likert scales and Quadratic Voting (QV). The QV method gives respondents a fixed budget to "buy" votes in favor or against a set of issues. Because the price for each vote is quadratic, it becomes increasingly costly to acquire additional votes to express support or opposition to the same issue. We formally show that QV better measures preference strength. This, we argue, is especially true in a polarized two-party system where individuals with weak preferences are more likely to mechanically default to the party policy position. Results from an online pilot survey provide preliminary support for our argument.

4. Non-Confrontational Extremists (TSE Working Paper No. 16-694; M. Michaeli, D. Spiro) develops another implication of deontological motivations for economics methodology. We present a critique of a standard assumption in economics—that the cost of deviating from one's bliss point is convex—with fundamental im-

plications for decision making in social and political settings and for the empirical predictions of theoretical models in these domains. We show that the concavity of ideological and moral costs holds in (1) naturally generated data, (2) where stakes are high, (3) when decisions are made by experts, and (4) when the consequences have the effect of silencing people with opinions that are far from the consensus, because they are the ones for whom it is most difficult to stand on their ground. In a high-stakes field setting (U.S. Courts of Appeals), we examine which individuals, on an ideological scale, conform most to the opinions of others. Legal precedents are set by ideologically diverse and randomly composed panels of judges. Using exogenous predictors of ideology and rich voting data we show that ideological disagreements drive dissents against the panel's decision, but ideologically extreme judges are caving in: they are the least likely to dissent and their voting records are the least correlated with their predicted ideology. Meanwhile, moderately ideological judges are dissenting the most despite evidence that they are more often determining the opinion. Our theoretical analysis shows that these findings are most consistent with a model of decision making in the presence of peer pressure with a concave cost of deviating from one's ideological convictions—perfectionism.

5. Modularizing the code across experiments led to oTree: An Open Source Platform for Online, Lab, and Field Experiments (*Journal of Behavioral and Experimental Finance*, 9(1), 88-97, 2016; M. Schonger, C. Wickens). oTree is an open-source and online software for implementing interactive experiments in the laboratory, online, the field or combinations thereof. oTree does not require installation of software on subjects' devices; it can run on any device that has a web browser, be that a desktop computer, a tablet or a smartphone. Deployment can be internet-based without a shared local network, or local network-based even without internet access. The programming language is Python, a popular, open-source programming language. www.oTree.org provides the source code, a library of standard game templates and demo games which can be played by anyone. oTree has been translated into French, German, Italian, Japanese, Spanish, and Russian. It is used in Computer Science, Economics, Engineering, Mathematics, Physics, Political Science, Psychology, and Sociology for classroom, lab, MTurk, and field experiments without internet. It is used in Australia (Melbourne, Newcastle, Perth, Queensland), Austria (Innsbruck, Vienna), Belgium (Leige), Canada (Guelph, Toronto), China (Beijing), Czech (Prague), Finland (Aalto), France (CReA Defense, Lille, Montpellier, Nice, Toulouse), Germany (GfK Marketing Research, Mannheim, Munich), Hungary (Academy of Sciences), Italy (Bologna, European University Institute), Japan (Tokyo), Kenya (Nairobi), Korea (Seoul), Netherlands (Amsterdam, Maastricht, Nijmegen, Tilburg, United Nations University, Utrecht), Norway (Norwegian School of Economics), Russia (RANEPA), South Africa (Pretoria), Spain (Madrid, Malaga, Valencia, Zaragoza), Sweden (Gothenburg, Stockholm), Switzerland (ETH, Geneva, Lausanne, Zurich), U.K. (Cambridge, Lancaster, Oxford), and U.S. (Boston College, Colby, Columbia, Iowa State, Michigan State, Northwestern, NYU, Ohio State, Princeton, Treasury Department, UCSC, UCSD, University of Chicago, Vanderbilt, Stanford, Yale).

6. Ambiguity aversion has recently been proposed as a behavioral foundation for polarization. Ambiguity aversion has also been used to explain a wide range of phenomena in law and policy: incomplete contracts, stock market volatility, abstention from voting, why prosecutors offer and defendants accept harsh plea bargains—and why individuals update priors after receiving identical information in an opposing manner. Is Ambiguity Aversion a Preference? (TSE Working Paper No. 16-703; M. Schonger) presents evidence problematizing the experimental basis for ambiguity aversion. Ambiguity aversion is the interpretation of the experimental finding known as the Ellsberg paradox that most subjects violate probabilistic sophistication: They prefer betting on events whose probabilities are known (objective) to betting on events whose probabilities are unknown to them (subjective). However in typical experiments these unknown probabilities are known and often determined by the experimenter. Thus the typical Ellsberg experiment is a situation of asymmetric information. People may try to avoid situations where they are the less informed party in an asymmetric situation setting. Indeed doing

so is often normatively appropriate. Thus avoidance of situations of informational asymmetry is a potential confound in typical Ellsberg experiments. Paying to avoid information asymmetry in an Ellsberg experiment would constitute the misapplication of a heuristic to the unfamiliar experimental situation. To eliminate this confound, this paper proposes a new source of ambiguity: participant generated ambiguity. Instead of the experimenter filling an Ellsberg urn, the opaque Ellsberg urn is filled by the other subjects in a laboratory session. We find that eliminating asymmetric information while leaving ambiguity in place, makes subjects more than willing to choose the ambiguous bet rather than the objective one even when choosing the objective bet is costless.

7. Testing Axiomatizations of Ambiguity Aversion (TSE Working Paper No. 16-717; M. Schonger) empirically interrogates existing ambiguity aversion theories. The study of the normative and positive theory of choice under uncertainty has made major advances through thought experiments often referred to as paradoxes: the St. Petersburg paradox, the Allais paradox, the Ellsberg paradox, and the Rabin paradox. Machina proposes a new thought experiment which posits a choice between two acts that have three outcomes. As in the Ellsberg paradox there are three events, but while the Ellsberg paradox has two (monetary) outcomes in Machina there are three. Machina shows that four prominent theories of ambiguity aversion predict indifference between the acts. Introspection, however, suggests that many people might very well strictly prefer one act over the other. This paper makes four contributions: first, to our knowledge, it is the first to experimentally implement the Machina thought experiment. Second, we employ a novel method to simultaneously elicit the certainty equivalent of an embedded lottery. Third, our results—across three experiments—indicate non-indifference, which rejects earlier theories of ambiguity aversion, but is consistent with a newer one, which we apply to explain our results. Fourth, we show that independence is a sufficient condition for indifference in the Machina paradox, and thereby explains why so many models predict indifference.

6 Behavioral Judging

By justice, we think of equal treatment before the law and equality based on recognition of difference. We can imagine a set of covariates X that should lead to the same prediction or predictability of outcomes $Y = f(X) + \varepsilon$. X should improve the prediction. And there's a set of W's that should not ($y \perp W, var(\varepsilon) \perp W$). We generally think of the W's as immutable and the X's as mutable as products of choices ($a \to X, a \nrightarrow W$); though, this view is also changing – W's can also be mutable if they are expressions of one's identity. Judges who have self-concept of being just, if they misperceive what is a fair or random sequence may actively negatively autocorrelate trying to do what they think is right. Judges who think of themselves as moral and have lexicographical preferences–first, do your duty–and approximate these lexicographical preferences with a concave cost of deviating from your duty may intentionally cave in when the benefits of deviation is too high. Judges who see people like themselves, randomly, excessively, may start to distinguish themselves; dissimilation in the management of self-identity. Judges may unintentionally bias, though economists may still opine that judges are intentionally assigning longer sentences to defend their ego. Judges can also indirectly bias when they say they're impartial but they are actually partial. When the W's influence their decisions, but they shouldn't; or when the X's should affect their decisions, but they do not. And to be sure self-image can positively affect justice, when they see another set of X's identical to a previous situation, they follow legal precedent (and not simply instrumentally for fear of reversal); or when they see new moral theories and start to evaluate their decisions in terms of cost and benefits. The first two papers examine when the W's (previous decision and defendant's name) inappropriately influence judicial decisions. The next three papers examine how W's can be mutable if they are expressions of one's identity (in the U.S., less so in Europe), and their influence on judicial decisions and labor market outcomes. The next six papers examine when the W's (politics, weather, race, shared judicial features) influence decisions, but they shouldn't.

6.1 Impossibility of Objective Judgments? Priming, Gambler's Fallacy, Mood, Voice, and Peer Effects in U.S. Courts

1. Decision-Making Under the Gambler's Fallacy: Evidence From Asylum Courts, Loan Officers, and Baseball Umpires (*Quarterly Journal of Economics*, 2016, 131(3): 1181-1241; NBER Working Paper No. 22026; TSE Working Paper No. 16-674; T. Moskowitz, K. Shue) finds consistent evidence of negative autocorrelation in decision-making that is unrelated to the merits of the cases considered in three separate high-stakes field settings: refugee asylum court decisions, loan application reviews, and major league baseball umpire pitch calls. The evidence is most consistent with the law of small numbers and the gambler's fallacy – people underestimating the likelihood of sequential streaks occurring by chance – leading to negatively autocorrelated decisions that result in errors. The negative autocorrelation is stronger among more moderate and less experienced decision-makers, following longer streaks of decisions in one direction, when the current and previous cases share similar characteristics or occur close in time, and when decision-makers face weaker incentives for accuracy. Other explanations for negatively autocorrelated decisions such as quotas, learning, or preferences to treat all parties fairly, are less consistent with the evidence, though we cannot completely rule out sequential contrast effects as an alternative explanation.

2. Implicit egotism—in particular, unconscious associations that individuals have with others who share their names or first initials—is a mainstay of modern psychology textbooks. Using unique data on 48,988 randomly assigned defendants from 1988-1999, Implicit Egoism in Sentencing Decisions: First Initial Name Effects with Randomly Assigned Defendants (TSE Working Paper No. 16-726) finds that judges assign 8% longer sentences when they match on first initials with the defendant. The effects are larger for Negroes, which could be due to behavioral biases playing a stronger role in evaluations when decision makers are nearly indifferent. The effect is consistent with threatened egotism, in which individuals motivated to manage self-image create social distance from negatively-valenced targets perceived to be associated with the self.

3. The emphasis on "fit" as a hiring criterion has raised the spectrum of a new form of subtle discrimination. Under complete markets, correlations between malleable characteristics and outcomes should not persist. Yet using data on 1,901 U.S. Supreme Court oral arguments between 1998 and 2012, Covering: Mutable Characteristics and Perceptions of Voice in the U.S. Supreme Court (*Review of Economic Studies*, invited to resubmit; TSE Working Paper No. 16-680; Y. Halberstam, A. Yu) documents that voice-based snap judgments based on lawyers' identical introductory sentences, "Mr. Chief Justice, (and) may it please the Court?", predict court outcomes. The connection between vocal characteristics and court outcomes is specific only to perceptions of masculinity and not other characteristics, even when judgment is based on less than three seconds of exposure to a lawyer's speech sample. Consistent with employers mistakenly favoring lawyers with masculine voices, perceived masculinity is negatively correlated with winning and the negative correlation is larger in more masculine-sounding industries. The first lawyer to speak is the main driver. Among these petitioners, males below median in masculinity are 7 percentage points more likely to win in the Supreme Court. Female lawyers are also coached to be more masculine and women's perceived femininity predict court outcomes. Republicans, more than Democrats, vote for more feminine-sounding females, while Democrats, but not Republicans, vote for less masculine-sounding men. A de-biasing experiment finds that information reduces 40% of the correlation between perceived masculinity and perceived win, and incentives reduce another 20%. A model shows how the information treatment identifies statistical discrimination and the incentives treatment identifies taste. Perceived masculinity explains additional variance relative to and is orthogonal to the best random forest prediction model of Supreme Court votes. In a separate paper written for a general scientific audience, Perceived Masculinity Predicts U.S. Supreme Court Outcomes (*PLoS-ONE*, 11(10), e0164324; TSE Working Paper No. 16-682; Y. Halberstam, A. Yu) documents the correlation between vocal characteristics

and court outcomes. While this study does not aim to establish any causal connections, our findings suggest that vocal characteristics may be relevant in even as solemn a setting as the Supreme Court of the United States. Previous studies suggest a significant role of language in the court room, yet none has identified a definitive correlation between vocal characteristics and court outcomes.

4. Using data from 1946–2014, Is Justice Really Blind? And Is It Also Deaf? (in Computational Analysis of Law, Santa Fe Institute Press, submitted, ed. M. Livermore and D. Rockmore) shows that audio features of lawyers' introductory statements and lawyers' facial attributes improve the performance of the best prediction models of Supreme Court outcomes. We infer face attributes using the MIT-CBCL human-labeled face database and infer voice attributes using a 15-year sample of human-labeled Supreme Court advocate voices. We find that image features improved prediction of case outcomes from 63.8% to 65.6%, audio features improved prediction of case outcomes from 66.8% to 68.8%, image and audio features together improved prediction of case outcomes from 66.9% to 67.7%, and the weights on lawyer traits are approximately half the weight of the most important feature from the models without image or audio features. Predictions of Justice votes with image and audio features however remained more similar relative to their baselines. We interpret this difference to suggest that human biases are more relevant in close cases.

5. Electoral Cycles Among U.S. Courts of Appeals Judges (*Journal of Law and Economics*, forthcoming; TSE Working Paper No. 16-704; C. Berdejo) finds evidence consistent with experimental studies that document the contexts and characteristics making individuals more susceptible to priming. Just before U.S. Presidential elections, judges on the U.S. Courts of Appeals double the rate at which they dissent and vote along partisan lines. Increases are accentuated for judges with less experience and in ideologically polarized environments. During periods of national reconciliation—wartime, for example—judges suppress dissents, again, especially by judges with less experience and in ideologically polarized environments. We show the dissent rate increases gradually from 6% to nearly 12% in the quarter before an election and returns immediately to 6% after the election. If highly experienced professionals making common law precedent can be politically primed, it raises questions about the perceived impartiality of the judiciary. We cannot rule out the possibility that judges—who profess to be unbiased—are intentionally biased, which raises different questions regarding intentional bias of professionals who claim to be unbiased.

6. To investigate the mechanism for why U.S. Courts of Appeals judges elevate their dissents for ten months prior to Presidential elections, Priming Ideology: Why Presidential Elections Affect U.S. Judges (TSE Working Paper No. 16-681) develops a theoretical model showing that the salience of partisan identities can explain this pattern. I link judges to their states of residence, and exploit temporal variation in the importance of a state during the electoral season. Dissents are elevated in swing states and in states that count heavily to winning the election, when these states are competitive. An active U.S. Senate election in the jurisdiction further elevates dissents. I link administrative data on case progression and frequency of campaign advertisements in judges' states of residence. Dissents occur shortly before publication, increase with monthly increases in campaign ads, and appear for cases whose legal topic, economic activity, is most heavily covered by campaign ads. Finally, I link the cases to their potential resolution in the Supreme Court. Dissents before elections appear on more marginal cases that cite discretionary miscellaneous issues and procedural (rather than substantive) arguments, which the Supreme Court appears to recognize and only partly remedy.

7. The behavioral response to presidential elections can also be intentional. Less than 1% of U.S. Federal judges report political motivations for retirement and resignation. However using two centuries of data, Judicial Exits from the U.S. Courts of Appeals are Politically Motivated (TSE Working Paper No. 16-721) shows that 13% of retirements and 36% of resignations follow political cycles. When the President comes from a

different political party as judge's party of appointment, U.S. Courts of Appeals judges are *less likely to retire* in each of the three quarters *before* a Presidential election. In contrast, judges are *more likely to resign* in each of the four quarters *after* a Presidential election when the President comes from the judge's party of appointment. Politically motivated exits have increased significantly in recent years to constitute one-fifth of retirements since 1975. I am able to uncover these patterns by analyzing the data at the quarter-to-election level, while prior research has relied on self-reports or yearly analysis. That highly experienced professionals making common law precedent can be self-deceiving raises questions about judicial impartiality. I cannot rule out the possibility that judges are outright deceptive, which also raises the question of outright deception of professionals who claim to be fair.

8. Recent work in cognitive science provides overwhelming evidence for a link between emotion and moral judgment. Mood and the Malleability of Moral Reasoning (TSE Working Paper No. 16-707) detects intra-judge variation in judicial decisions driven by factors completely unrelated to the merits of the case, or to any case characteristic for that matter. I show that asylum grant rates in U.S. immigration courts and sentencing decisions in U.S. district courts differ by the success of the court city's NFL team on the night before, and by the city's weather on the day of, the decision. My data including one-and-a-half million decisions spanning three decades allows me to exclude confounding factors, such as scheduling and seasonal effects. Most importantly, my design holds the identity of the judge constant. On average, U.S. immigration judges grant an additional 1.4% of asylum petitions–and U.S. district judges assign 0.6% fewer prison sentences and 5% longer probation sentences—on the day after their city's NFL team won, relative to days after the team lost. Bad weather on the day of the decision has approximately the opposite effect of a team win. Notably, the effects of NFL games on asylum decisions are driven entirely by unrepresented parties. This would be suggestive that the effects are due to judge decision-making as opposed to the game outcomes affecting other court participants such as lawyer behavior.

9. In United States District Courts for federal criminal cases, prison sentence length guidelines are established by the severity of the crime and the criminal history of the defendant. Events Unrelated to Crime Predict Criminal Sentence Length investigates the sentence length determined by the trial judge, relative to this sentencing guideline. Our goal is to create a prediction model of sentencing length and include events unrelated to crime, namely weather and sports outcomes, to determine if these unrelated events are predictive of sentencing decisions and evaluate the importance weights of these unrelated events in explaining rulings. We find that while several appropriate features predict sentence length, such as details of the crime committed, other features seemingly unrelated, including daily temperature, baseball game scores, and location of trial, are predictive as well. Unrelated events were, surprisingly, more predictive than race, which did not predict sentencing length relative to the guidelines. This is consistent with recent research on racial disparities in sentencing that highlights the role of prosecutors in making charges that influence the maximum and minimum recommended sentence. Finally, we attribute the predictive importance of date to the 2005 U.S. Supreme Court case, United States v. Booker, after which sentence length more frequently fell near the guideline minimum and the range of minimum and maximum sentences became more extreme.

10. Federal courts are a mainstay of the justice system in the United States. In What Matters: Agreement Between U.S. Courts of Appeals Judges (*Journal of Machine Learning Research*, forthcoming; TSE Working Paper No. 16-747; X. Cui, L. Shang, J. Zheng), we analyze 387,898 cases from U.S. Courts of Appeals, where judges are randomly assigned to panels of three. We predict which judge dissents against copanelists and analyze the dominant features that predict such dissent with a particular attention to the biographical features that judges share. Random forest achieves the best classification. Dissent is roughly half-driven by case features and half-driven by personal features.

6.2 Judicial Analytics and Behavioral Foundations of Polarization

Judicial Analytics and the Great Transformation of American Law (Journal of Artificial Intelligence and the Law, revise and resubmit) presents an overview of ongoing machine learning analysis of 12 terabytes of judicial data. Three broad themes emerge. First, legal factors broadly predict decisions more than political factors alone. Second, implicit associations (bias) in judicial texts or oral arguments can be surprisingly predictive. Third, machine learning holds the promise for automated impact analyses of judicial decisions.

1. The (Non-)Polarization of U.S. Circuit Court Judges, 1930-2013 (E. Ash, W. Lu) employs a supervised learning approach to measure the polarization of U.S. Circuit Court judges using the text of court opinions for the years 1930-2013. Our results show persistent but low partisanship of court opinions in the past century, with little evidence of an increase or decrease in partisanship. We visualize voting networks of Circuit Court judges, and we nd that judges are not polarized, although the graph exhibits clustering within the courts. Third, we study the behavior of Circuit Court judges during Supreme Court vacancies; judges who were candidates of nomination write fewer discretionary opinions when the Senate is controlled by the opposing political party.

2. In "Precedent vs. Politics? Case Similarity Predicts Supreme Court Decisions Better Than Ideology," using the universe of U.S. Circuit Court cases appealed to the Supreme Court since 1946, we show that case similarity among Circuit Court opinions achieves better prediction accuracy of Supreme Court decisions relative to the current best prediction model, which is based on ideology of judges and trends of how they vote. Relative to the benchmark prediction accuracy of 59%, textual measures of case similarity achieve prediction accuracy of 64%. We interpret this improvement to suggest that precedent matters more than politics alone. Combining case similarity with ideological features further improves accuracy to 72%, suggesting that ideology affects interpretation of precedent. We also offer our model available as a web app.

3. In "Affirm or Reverse? Using Machine Learning To Help Judges Write Opinions," we predict higher court reversals of lower court decisions. Every year more than 300,000 civil and criminal cases are heard in the district courts all over the U.S. Less than 5% of these cases are appealed and heard in circuit courts. For most of the cases, the circuit court either affirms the decision of the district court or reverses it. Out of the cases heard in circuit courts, only about 2% are heard in the Supreme Court. The Supreme Court again can again either affirm or reverse the circuit court decision. We build a model to predict the higher court decision using the lower court's opinion using a digital corpora of circuit court cases to 1891 and district court cases to 1923. Comparing a wide variety of classification, dimensionality reduction, and oversampling techniques, we are able to achieve an accuracy of 79% in circuit courts and 68% accuracy in the Supreme Court.

4. Recent work in natural language processing represents language objects (words and documents) as dense vectors that encode the relations between those objects (Blei, 2012; Mikolov et al., 2013). These methods have recently been adapted to the analysis of human social behavior (e.g. Caliskan et al., 2017). In "Judge Embeddings: Vector Representations of Legal Belief," (E. Ash). This paper explores the vectorization of legal beliefs, with the goal of understanding judicial reasoning and the causal impacts of law. We illustrate the usefulness of these vectors in three ways. First, we show that they recover intuitive institutional connections between judges. Second, we show the vectors can be used as features in a decision prediction task. Third, we show that they can be used to measure implicit bias by judges toward women and racial minorities

5. In "Law vs. Fact: Motivated Reasoning and Judicial Fact Discretion," using the universe of U.S. Circuit Court cases since 1891, we divide opinions into law vs. fact. We map sentiment towards policy (e.g., culture, immigration, nationalization, democracy) at the paragraph level of each opinion. We show that a judge's decision to affirm or reverse can be predicted by these sentiment features with AUC scores of 0.95 using AdaBoost, which performs the best classification.

6. In "Measuring Implicit Bias and its Consequences Using Semantics Derived Automatically from Judicial Corpora," we show that word embeddings in judicial corpora reflect human-like semantic biases measured by the Implicit Association Test. We use a purely statistical machine-learning model trained on the universe of U.S. Circuit Court opinions since 1891. We show that male judges write opinions whose semantics reflect attitudinal disparities in favor of males while female judges write opinions whose semantics are gender-neutral. We show that judges whose semantics reflect positive attitudes towards government are more likely to vote in favor of government regulation. We compare our empirical approach to orthogonalized machine learning to measure the causal impact of judge's implicit biases on outcomes.

7. In "Law and Sentiment: Word Embeddings in U.S. Courts Impact Population Attitudes," we ask the question, do judge's attitudes, reflected in court opinions, impact population attitudes? We present an application of machine learning and causal inference using the universe of U.S. Circuit Court opinions from 1891. We assess the sentiment of each paragraph on thermometer (sentiment) questions in the American National Election Survey. We calculate cosine similarity of each paragraph to each thermometer target (e.g., Republicans, Democrats, woman, feminists, etc.) using word2vec, and use a sentiment analyzer to compute the average sentiment (positive or negative) of each opinion towards each target. We then use LASSO to select biographical characteristics predictive of attitudes and use the selected characteristics as instruments in a two-stage least square regression. We show that sentiment reflected in court opinions impact population attitudes in an opposite direction. These results complement other quantitative and qualitative work suggesting that court actions often led to resistance.

8. In "Predicting Punitiveness and Sentencing Disparities from Judicial Corpora," using over 1 million sentencing decisions linked to judge identity and the digital corpora of U.S. district court opinions, we show that judges' writings can predict average harshness and racial and sex disparities in sentencing decisions. We document significant reductions in mean square error relative to a naive prediction (the mean of the training data) on the test dataset by approximately 24 percent in predicting punitiveness.

9. Investigating Variation in English Vowel-to-Vowel Coarticulation in a Longitudinal Phonetic Corpus (*Proceedings of the 18th International Congress of Phonetic Sciences*, 2015; C. Abrego-Collier, J. Phillips, B. Pillion, A. Yu) investigates the nature of individual variation in speech, particularly the mechanism underlying such variability, is increasingly important, especially for research on sound change, since such investigations might help explain why sound change happens at all and, conversely, why sound change is so rarely actuated even though the phonetic preconditions are always present in speech. The present study contributes to the literature on inter- and intraspeaker variation in coarticulation, a major precursor to sound change, by focusing on the degree of coarticulation stressed vowels have on neighboring unstressed vowels using recordings from a longitudinal phonetic corpus of oral arguments before the Supreme Court of the United States. Significant inter-speaker variation in height coarticulation, both anticipatory and carryover, is observed, while no evidence for systematic inter-speaker variability in backness coarticulation is found. There is also no evidence for intra-speaker variation in coarticulation over the course of 205 days. Mimicry: Phonetic Accommodation Predicts U.S. Supreme Court Votes (A. Yu) digitizes speech patterns in US Supreme Court oral arguments and shows that lawyers converge to justices are more likly to win their vote, and that justices who converge to each other during oral arguments are more likely to vote together.

1. In "Implicit Bias in Supreme Court Speech: Inferences of Gender Attitudes from Vocal Patterns Predict Judicial Decisions," using 15 years of Supreme Court oral arguments, we show that vocal intonation of gendered words (e.g., actor vs. actress) classify vocal intonations of neutral words into stereotypically male (e.g., logical, ability, think) and female (e.g., looking, cook, goodwill), surprisingly suggesting the relevance of people's

perceptions of gender being revealed in how people speak. Furthermore, the vocal intonations of judges' speaking these words are predictive of their decisions. Justice Stephen Breyer can be predicted with 73% accuracy while Justice Antonin Scalia can be predicted with 58% accuracy. These results complement other work indicating that perceptions of gender improve predictions of Supreme Court outcomes and continue to play a role in a manner more complex and nuanced than conventionally perceived.

2. In "Dialects of Ideology," we examine whether speech variation beyond word choice, that is, fluctuations in the way one speaks holding the words fixed is predictive of ideology in the U.S. Supreme Court. We use lawyers' campaign donations as a commonly-used measure of political ideology. We find that audio significantly improves prediction accuracy of ideology relative to using the text alone. AUC increases from 0.55 to 0.61, even in a setting as solemn as the Supreme Court.

3. In "Does Dicta Matter? Evidence from Environmental Law," we pose the question whether judicial decisions affect outcomes only through their verdicts alone, or whether the manner of an opinion's writing (dicta) also affects outcomes. We characterize U.S. federal judges as vectors of their opinions and analyze the causal impact of random assignment of judges with different opinion styles in environmental law. We document a very strong first stage in opinion style, for example, persistent citation habits. Instrumenting for the judge's past style, we show that dicta impacts emission levels of various air pollutants. This result complements other work showing that the words of tax statutes impact tax revenues more than tax rates do.

4. In "Identifying Policy Levers: Automatic Case Classification and Validation using Citations," to build inputs for end-to-end machine learning estimates of the causal impacts of law, we consider the problem of automatic case classification. We consider the quasi-supervised multi-class problem using as training set the Chicago Judges Project (a hand-coded dataset of thousands of cases in over 20 politically salient topics). Our model achieves 84% correct classification using only the opinion's text. We show that, among citations used at least three times, 78% fall within the same cluster/predicted category. Similarly, in "Predicting Medical Malpractice of Health Practitioners Using Microdata on Treatment Choices," Medicare data from 2006-2012 and Florida malpractice claims can help build a predictive model of malpractice risk at the physician level and be used as inputs in law and economics analysis of malpractice.

7 Law and Legitimacy

According to communitarian philosophers, the idea of ego or self-concept can be traced to Enlightment and Romantic ideals of the self—Enlightenment ideals of self-knowledge and self-mastery, presuming that each person has an interior space, and memory of oneself, and Romantic ideals of self-expression and authenticity, being true to oneself. Recognizing that there are different sources of the self, according to philosophers, constituted a moral revolution, in which projects of personal identity came to be important beyond economic self-interest. Questions of identity became personally significant; refusals of acceptance and respect, deeply challenging. We can see projects of personal identity being mobilized in progressive (the transformation of gender and sexual identities and claims to equal rights) and conservative politics (the claims to national and communal identities defended by populist movements). We can see projects of personal identity in self-esteem and authenticity being taught as values in school ("find your passion"). We can see projects of personal identity and self-esteem–violations of self-esteem–in contemporary discussions of microaggression, trigger warnings, and privilege disparities. According to these philosophers, recognizing that everyone has their own way of being human facilitates respect for individuals, but also for different cultures.

7.1 Deter or Spur: British Executions During World War I

Deter or Spur? British Executions During World War I is a book-length version of The Deterrent Effect of the Death Penalty? Evidence from British Commutations During World War I (*American Economic Review*, 2nd invited to resubmit; TSE Working Paper No. 16-706). During World War I, the British military condemned over 3,000 soldiers to death, but only executed 12% of them; the others received commuted sentences. Many historians believe that the military command confirmed or commuted sentences for reasons unrelated to the circumstances of a particular case and that the application of the death penalty was essentially a random, "pitiless lottery." Over two death sentences a day precluded careful consideration to execute or commute by the Commander-in-Chief, who was responsible for the final decision.

In addition, we have the names of all the soldiers—those who received the death sentence and those who deserted—and we can see, to a first approximation, if they were Irish or British. The Irish were a large subordinated minority that declared independence right after the war. Like minorities elsewhere, they were sentenced to death at a much higher rate than the general population. But conditional on the death sentence, they were equally likely to be executed. Holding punishment constant allows the potential identification for a channel beyond deterrence, namely legitimacy, to explain compliance to the law.

Many psychologists, sociologists, political scientists, and legal theorists have all emphasized the importance of legitimacy in courts, organizations, and nation states. But the existing quantitative evidence has been limited to surveys asking different communities whether they think the lawgiver is legitimate and correlating these perceptions with rates of crime, which can be correlated for unobserved reasons. In particular, they can be correlated if unequal punishment leads to higher incarceration rates and lower perceptions of legitimacy or if economic inequality drives crime rates and perceptions. This identification of the potential deterrent and delegitimizing effects of executions hinges on whether executions were random.

Using a dataset on all capital cases during World War I, I find that the data are consistent with an essentially random process. Using this result, I exploit variation in commutations and executions within military units to identify the deterrent effect of executions, with deterrence measured by the elapsed time within a unit between the resolution of a death sentence (i.e., a commutation or execution) and subsequent absences within that unit. Absences are measured via handwritten trial records and "wanted" lists prepared by British military police units searching for deserters and preserved in war diaries and police gazettes. I find some limited evidence that executing deserters deterred absences, while executing Irish soldiers, regardless of the crime, spurred absences, particularly Irish absences. I present a model where perceived legitimacy of authority affects why people obey the law.

The differences between the situation I study and contemporary criminal justice scenarios are vast, so a more nuanced understanding of the differences is required in order to draw policy lessons from the WWI experience. Because the British experience provides an extremely low-bar test for the death penalty, finding a deterrence effect in the context of WWI would certainly not be a strong argument, leaving aside moral issues, that the death penalty is good policy. However, a negative result showing no deterrent effect might have more policy salience since if we ever expected to find an effect, it would be in the WWI context: executions took place almost immediately—in a manner purposefully designed to maximize their deterrent effect—and death sentences were given out very frequently and quite arbitrarily for military desertion. We would still expect that on the margin more executions should deter absences and if we find this not to be the case, it would suggest that the threat of future death for crimes is not as strong a disincentive as we might imagine. Despite these differences, this study offers some insights potentially capable of greater generalization. The granularity and richness of the data begets questions that are sometimes ignored in standard crime rate studies.

7.2 Revealed Preference Indifference: Legitimacy, Law, and Recognition-Respect

Whether the law is legitimate can be related to how individuals are treated, the identity of the lawmaker, and one's reference points. It can be studied in contextualized field experiments and observed in failed policy changes. Most research on judicial decision-making examine when extraneous factors influence their decisions; inverting the question, evidence when legal factors should affect their decisions, but they do not, raise questions of snap or predetermined judgement. These ideas are summarized in Using Machine Learning to Detect Human Rights Abuses (in Computational Analysis of Law, Santa Fe Institute Press, submitted, ed. M. Livermore and D. Rockmore).

1. White Privilege in Law: The Reproduction of Racial Hierarchy in Sentencing Disparities introduces individual-level prosecutorial (as well as individual-level judicial) data. I observe screening (including declination decisions), charging, trial, and sentencing outcomes in a unique dataset vertically linked from the time of arrest. I leverage the random assignment of cases to prosecutors and judges and ask whether judge-based or prosecutor-based disparities interact in a courtroom setting – i.e., does the race of the prosecutor influence the extent of judge-based disparities. Two results emerge. First, cases assigned to white judges and black prosecutors receive shorter sentences. Second, black-white sentencing disparities reverse for cases assigned to black judges and black prosecutors. These results are consistent with narratives of racial hierarchy in the law whereby black-white disparities are rendered and reproduced.

2. Clash of Norms: Judicial Leniency on Defendant Birthdays (with A. Philippe) examines, while judges are supposed to treat defendants equally, how judges react on days when there is a strong social norm for generosity and gift giving? Offenders' birthday provides a good natural experiment as the social norm is unambiguous and the event is exogenous. Using 6 million decisions in France, we show that judges are more lenient when decisions take place on the defendants' birthday. However, this is only true if the defendant shows up to court. The leniency is clearer for probation time than for prison time. In 600,000 decisions in U.S. federal courts, the day part of prison sentences is affected while the month part is not. These results suggest that mental accounting and reference points influence judicial sentencing

3. How Prosecutors Exacerbate Racial Disparities: Screening Gaps, Race Effects, and Courtrooms Interactions links the universe of individuals in a district attorney's office over a decade with many stages of random assignment. Three facts emerge. First, racial disparities in criminal sentencing magnify once discretion in prosecutorial screening is taken into account. 50% of arrestees are screened out by prosecutors, who disproportionately favor whites. Taking this screening gap into account greatly magnifies real sentencing disparities. Second, race of screener prosecutors matters significantly. White prosecutors screen in fewer arrestees and these defendants get assigned shorter sentence lengths. Black prosecutors screen in far more black arrestees. Third, racial interactions in the courtroom affect sentencing levels and disparities. When defendants are assigned to black trial prosecutors and white judges, they receive substantially lighter sentences. Black trial prosecutors and black judges eliminate or reverse the sentencing gap, such that longer sentences are rendered for white defendants. This eliminates the theoretical possibility that black-white sentencing disparities reflect unobserved heterogeneity beyond race.

4. The appointment of Sonia Sotomayor to the Supreme Court in 2009 was criticized as sacrificing merit on the altar of identity politics. According to critics, Sotomayor was simply "not that smart." For some conservative critics, her selection illustrated the costs of affirmative action policies, in that this particular choice was going to produce a lower quality Supreme Court. For liberal critics, many were concerned that the President, by selecting Sotomayor, was squandering an opportunity to appoint an intellectual counterweight to conservative Justices like Antonin Scalia, Samuel Alito, and John Roberts. Using a set of basic measures of judicial merit, such as publication and citation rates for the years 2004 to 2006, when Sotomayor was on the Court of

Appeals for the Second Circuit, 'Not that Smart': Sonia Sotomayor and the Construction of Merit (*Emory Law Journal*, 61(4), 2012, with G. Charles and M. Gulati) compares her performance to that of her colleagues on the federal appeals courts. Sotomayor matches up well.

5. Tastes for Desert and Placation: A Reference Point-Dependent Model of Social Preferences (*Research in Experimental Economics*, conditionally accepted; TSE Working Paper No. 16-725) proposes a model of behavior in social interactions where individuals maximize a three-term utility function: a conventional consumption utility term and two "social" terms that capture social preference. One social term is a taste for desert, which is maximized when the individual believes the other person is getting what they deserve. The second social term measures the target individuals' anger or gratitude from the interaction which is determined by a value function derived from prospect theory. After introducing the model and generating a series of comparative statics results and derived predictions, I report the results of a series of quasi-field experiments on social preferences. I discuss how the model explains several paradoxes of empirical moral philosophy that are less explicable by current economic models of social preference focusing on outcomes and intentions.

6. The emergence of online labor markets makes it far easier to use individual human raters to evaluate materials for data collection and analysis in the social sciences. Designing Incentives for Inexpert Human Raters (*Proceedings of the ACM Conference on Computer Supported Cooperative Work*, 2011, with A. Shaw and J. Horton) reports the results of an experiment—conducted in an online labor market—that measured the effectiveness of a collection of social and financial incentive schemes for motivating workers to conduct a qualitative, content analysis task. Overall, workers performed better than chance, but results varied considerably depending on task difficulty. We find that treatment conditions, which asked workers to prospectively think about the responses of their peers—when combined with financial incentives—produced more accurate performance. Other treatments generally had weak effects on quality. Workers in India performed significantly worse than U.S. workers, regardless of treatment group.

7. In some online labor markets, workers are paid by the task, choose what tasks to work on, and have little or no interaction with their (usually anonymous) buyer/employer. These markets look like true spot markets for tasks rather than markets for employment. Despite appearances, Are Online Labor Markets Spot Markets for Tasks? A Field Experiment on the Behavioral Response to Wages Cuts (*Information Systems Research*, 27(2), 403-423; J. Horton; TSE Working Paper No. 16-675) find via a field experiment that workers act more like parties to an employment contract: workers quickly form wage reference points and react negatively to proposed wage cuts by quitting. However, they can be mollified with "reasonable" justifications for why wages are being cut, highlighting the importance of fairness considerations in their decision making. We find some evidence that "unreasonable" justifications for wage cuts reduce subsequent work quality. We also find that not explicitly presenting the worker with a decision about continuing to work eliminates "quits," with no apparent reduction in work quality. One interpretation for this finding is that workers have a strong expectation that they are party to a quasi-employment relationship where terms are not changed, and the default behavior is to continue working.

8. Interim Report on a Preschool Intervention Program in Kenya (resting paper; P. Glewwe, M. Kremer, S. Moulin) evaluates an educational program that professionalized an informal educational system. Teacher training, classroom materials, and incentives for teacher attendance was provided to fifty preschools randomly selected from one hundred preschools in rural Kenya. Teachers were eligible for bonuses of up to 85% of pre-program salary depending on their attendance. Headmasters acted as monitors and distributed funds. In practice, headmasters typically paid the entire bonus to teachers regardless of attendance, which tended to crowd out parental contributions to teacher salary. Teacher training significantly reduced the number of

minutes spent on the blackboard. The point estimates suggest that the program improved teacher attitude, energy, effort, control, and organization. The program increased progression to grade one by the end of three years; however, it also significantly decreased written test scores after two years.

9. Early Predictability of Asylum Court Decisions (*Proceedings of the ACM Conference on AI and the Law*, forthcoming; M. Dunn, L. Sagun, H. Sirin) presents evidence of judges ignoring information. In the United States, foreign nationals who fear persecution in their home country can apply for asylum under the Refugee Act of 1980. Over the past decade, legal scholarship has uncovered significant disparities in asylum adjudication by judge, by region of the United States in which the application is filed, and by the applicant's nationality. These disparities raise concerns about whether applicants are receiving equal treatment under the law. Using machine learning to predict judges' decisions, we document another concern that may violate our notions of justice: we are able to predict the final outcome of a case with 80% accuracy at the time the case opens using only information on the identity of the judge handling the case and the applicant's nationality. Moreover, there is significant variation in the degree of predictability of judges at the time the case is assigned to a judge. We show that highly predictable judges tend to hold fewer hearing sessions before making their decision, which raises the possibility that early predictability is due to judges deciding based on snap or predetermined judgments rather than taking into account the specifics of each case. Early prediction of a case with 80% accuracy could assist asylum seekers in their applications.

10. In Can Machine Learning Help Predict the Outcome of Asylum Adjudications? (*Proceedings of the ACM Conference on AI and the Law*, forthcoming; J. Eagel), we analyzed 492,903 asylum hearings from 336 different hearing 2 locations, rendered by 441 unique judges over a thirty-two year period from 1981-2013. We define the problem of asylum adjudication prediction as a binary 4 classification task, and using the random forest method, we predict twenty-seven years of refugee decisions. Using only data available up to the decision date, our model correctly classifies 82 percent of all refugee cases by 2013. Our empirical analysis suggests that decision makers exhibit a fair degree of autocorrelation in their rulings, and extraneous factors such as, news and the local weather may be impacting the fate of an asylum seeker. Surprisingly, granting asylum is predominantly driven by trend features and judicial characteristics–features that may seem unfair–and roughly one third-driven by case information, news events, and court information. Ignoring information is related to motivated cognition (the interpretation of information in a polarized manner) and may be related to perfectionism. If one's actions are unlikely to change as a result of information acquisition, why update one's beliefs? In "Endogenous Information Acquisition: Wikileak State Department Cables Predict Asylum Decisions Almost Perfectly," relative to the previous best prediction model of asylum court decisions of 80%, we show that refugee asylum outcomes are nearly perfectly predicted (with 98% accuracy) after incorporating US diplomatic communications captured in Wikileak cables. We interpret this to suggest that international current events interact with domestic immigration decisions in a meaningful way. Successfully predicting these decisions may be used to counsel refugees on their application's chance of success.

11. Legal theorists have suggested that literature stimulates empathy and affects moral judgment and decision-making. Law and Literature: Theory and Evidence on Empathy and Guile (*Review of Law and Economics,* forthcoming; TSE Working Paper No. 16-684) presents a model to formalize the potential effects of empathy on third parties. Empathy is modeled as having two components–sympathy (the decision-maker's reference point about what the third party deserves) and emotional theory of mind (anticipating the emotions of another in reaction to certain actions). I study the causal effect with a data entry experiment. Workers enter text whose content is randomized to relate to empathy, guile, or a control. Workers then take the Reading the Mind in the Eyes Test (RMET) and participate in a simple economic game. On average, workers exposed to empathy become less deceptive towards third parties. The result is stronger when workers are nearly

indifferent. These results are robust to a variety of controls and model specifications.

12. "Algorithms as Prosecutors: Lowering Rearrest Rates Without Disparate Impacts and Identifying Defendant Characteristics 'Noisy' to Human Decision-Makers," (D. Amaranto, E. Ash, L. Ren, C. Roper) investigates how machine learning might bring clarity to human decisions made during the criminal justice process. Our data comes from all cases at the New Orleans District Attorney's oce for the years 1988-1999. We exploit random assignment of prosecutors, prosecutorial discretion, and heterogeneity across prosecutors in charge rates to compare prediction models to judicial decision makers. Our model of defendant rearrest, trained using defendant and oense characteristics, selects higher-risk individuals to prosecute than its human counterparts did. In particular: given a set charge rate, our model would reduce rearrest rates by ve to nine percentage points. This model could have several important policy implications: it might identify defendant characteristics that are particularly 'noisy' to prosecutors; it could suggest ways of alleviating criminal caseloads without increasing crime rates; and it might provide important insights into how a prosecutor's background relates to the quality and nature of their charging decisions.

8 Demography of Ideas

The final section explores the transmission and persuasion process of novel moral theories, whose incommensurability can lead to group conflict. On a thick vein, I study the macro and demographic forces and laws that aid or hinder discrimination in legal institutions, markets, and public policy.

1. Developing countries with highly unequal income distributions, such as Brazil or South Africa, face an uphill battle in reducing inequality. Educated workers in these countries have a much lower birthrate than uneducated workers. Assuming children of educated workers are more likely to become educated, this fertility differential increases the proportion of unskilled workers, reducing their wages, and thus their opportunity cost of having children, creating a vicious cycle. Income Distribution Dynamics with Endogenous Fertility (*American Economic Review*, 89(2), 155-160, May 1999; NBER Working Paper No. w7530; M. Kremer) and Income Distribution Dynamics with Endogenous Fertility (*Journal of Economic Growth*, 7(3), 227-258, 2002; NBER Working Paper No. w7530; M. Kremer) present a model incorporating this effect, which generates multiple steady-state levels of inequality, suggesting that in some circumstances, temporarily increasing access to educational opportunities could permanently reduce inequality. Empirical evidence suggests that the fertility differential between the educated and uneducated is greater in less equal countries, consistent with the model.

2. A number of countries have begun implementing tax incentives designed to reverse the decline in fertility. Whether such incentives are effective or equitable remains an open question. During the early 20th century, France initiated an unusual tax policy to promote fertility and marriage: household income was divided by family size to obtain a final tax bracket. Similar policies whose fertility benefit increases with income are being implemented today. Using hand-collected archival data from aggregate tax returns and three natural experiments, Can Countries Reverse Fertility Decline? Evidence from France's Marriage and Baby Bonuses, 1929-1981 (*International Tax and Public Finance*, 118(3), 252-271, 2011) finds mixed evidence that these tax incentives affect fertility and marriage.

3. For those facing infertility, using assisted reproductive technology to have genetically related children is a very expensive proposition. In particular, to produce a live birth through in vitro fertilization (IVF) will cost an individual (on average) between $66,667 and $114,286 in the U.S. If forced to pay these prices out of pocket, many would be unable to afford this technology. Given, this reality, a number of states have attempted to

improve access to this technology through state-level insurance mandates that cover IVF. Several scholars, however, have worried that increasing access in this way will cause a diminution in adoptions and have argued against enactment of these mandates for that reason. In Trading Off Reproductive Technology and Adoption: Do IVF Subsidies Decrease Adoption Rates? (*Minnesota Law Review*, 95(2), 2010, I. G. Cohen), we push against that conclusion on two fronts. First, we interrogate the normative premises of the argument and expose its contestable implicit assumptions about how the state should balance the interests of existing children waiting for adoption and those seeking access to reproductive technology in order to have genetically related children. Second, we investigate the unexamined empirical question behind the conclusion: does state subsidization of reproductive technologies through insurance mandates actually reduce adoption; that is, is there a trade-off between helping individuals conceive and helping children waiting to be adopted? We call the claim that there is such an effect the "substitution theory." Using the differential timing of introduction of state-level insurance mandates relating to In Vitro Fertilization (IVF) in some states and differences in the forms these mandates take, we employ several different econometric techniques (differences-in-differences ordinary least square, two-stage least square) to examine the effect of these mandates on IVF utilization and adoption. Contrary to the assumption of the substitution theory, we find no strong evidence that state support of IVF through these mandates crowds out either domestic or international adoption.

4. On a practical level, concepts of family formation and views on sexual conduct and freedom may be fundamentally divided on religious conflict lines—some scholars contend that divisions between Western and Muslim countries are vast when it comes to attitudes towards divorce, abortion, gender equality, and gay rights. Different views of family can be decisive for the support for inter-group transfers within a nation and could decide on support for, for example, universal welfare goods, public provision of pensions and healthcare. There may also be related conflict risks (e.g., between groups where each group views the other as immoral). Different views on family and religious denomination could also be important for willingness to trust one another and support for public goods. These issues are particularly important for countries that are experiencing a potential polarization of attitudes (such as parts of Europe, where relatively large proportions of the population are highly liberal, while several important subpopulations and growing migrant groups have traditional family views).

8.1 Geneology of Ideology

1. Ideas Have Consequences: The Impact of Law and Economics on American Justice (E. Ash, S. Naidu) provides a quantitative analysis of the effects of the law and economics movement on the U.S. judiciary using the available universe of opinions in U.S. Circuit Courts and 1 million District Court criminal sentencing decisions linked to judge identity. We estimate the effect of attendance in the controversial Manne economics training program that 40% of federal judges attended by 1990. To isolate the effect of judges from the types of cases they face, we exploit random assignment of judges to control for court- and case-level factors, an exogenous seating network from random panel composition to trace the spread of economic reasoning in law, and ordering of cases within Circuit to identify general economic ideas that move across legal topics. We use natural language processing methods to quantify the influence of economics in written judicial opinions. Descriptively, we find that judges who use law and economics language vote for and author conservative verdicts (as coded by Songer-Auburn) in economics cases and are more opposed to government regulation. After attending Henry Manne's economics training program, participating judges use more economics language and render conservative verdicts in economics cases, rule against regulatory agencies, particularly in labor and environmental cases, get cited more and increase dissents. These results are robust to a large set of judge biographical controls, and do not exist prior to Manne program attendance, suggesting a causal effect of

economics training on judicial decisions. Further, Manne economics training is more predictive of these decisions than appointing political party. We further document a number of indirect channels of economics influence on the law beyond the direct effect on Manne program participants. Non-Manne judges exposed to Manne peers on previous cases increase their use of economics language in subsequent opinions. Further, some economics concepts are portable across legal contexts: we show that "general-purpose" economics phrases such as "capital", and "efficiency" move across legal topics within a judge. Economics reasoning diffused from regulatory domains into criminal law. Consistent with this, we show that law and economics influenced criminal decisions: Circuit Court judges that attend the Manne program are more likely to reject criminal appeals, and this effect spills over onto non-Manne judges serving on the same panel, results robust to judge biographical controls and absent prior to Manne program attendance. Moving to district courts, after attending Henry Manne's economics training program, participating judges increase sentence lengths, and using variation in judicial discretion generated by U.S. v. Booker, sentencing gaps between Manne and non-Manne judges grew by 20% (10 months) after this ruling, which allowed more judicial sentencing discretion. Finally, we document that Manne attendance is more predictive of racial and gender sentencing disparities than party of appointment, and Manne judges in both Circuit and District Courts render harsher immigration decisions, voting for enforcement of immigration regulation and longer sentences for illegal aliens. Peer effects is further explored in Is Ideology Infectious? Evidence from Repeated Random Exposure in the U.S. Courts of Appeals and polarization in Mitosis of Ideology: Integration and Assimilation or Dis-integration, Radicalization, Other-ing, and Egotism? .

2. Building towards the demography of attitudinal change is a long-term research goal. The incommensurability of novel moral theories can lead to group conflict, which I have studied using data in objective and subjective knowledge in science and law. Thomas Kuhn's Structure of Scientific Revolutions (1962) proposes a theory of knowledge— science undergoes periodic paradigm shifts—that has been interpreted to mean that there is no 'truth' in science. To test this theory, How Does Science Progress? A Statistical Approach to Postmodern Theories of Knowledge constructs "citation trees," where each node is a paper connected to all the papers it cites and the papers that cite it. Just as in evolutionary biology where a species tree has branches that get sparse or discontinued during periods of mass extinctions, I look for similar events in a citation tree. Using ISI data, I construct geology, linguistics, and literary criticism trees dating from 2001 back to 1945, 1956, and 1975 respectively. I also use the NBER Patent Citation database from 1975-1999. Markov clustering algorithms marking death of "bushes" indicate linguistics and geology are qualitatively different from literary criticism, where statistical "extinctions" occur often, contrary to the constructionist claim that all forms of discourse accumulate similarly. Paradigm shifts incommensurate enough to cause mass extinctions do not appear to have occurred in linguistics, geology, or patent citations but literary criticism may have had a paradigm shift in 1990.

3. The Genealogy of Ideology: Identifying Persuasive Memes and Predicting Agreement in U.S. Circuit Courts (*Proceedings of the ACM Conference on AI and the Law*, forthcoming; A. Parthasarathy, S. Verma) identifies memes based on the likelihood of legal phrases appearing along an edge of the citation tree but not appearing on a disconnected part of the tree. We then predicted how two judges in a particular panel align on their voting, based on the historical vote alignment of that judge with other judges, historical citation of each other, and historical use of shared legal phrases and memes. This paper was cited in the 2017 American Economic Association Presidential Address by Robert Shiller.