

Machine Learning to Combat Wicked Problems: Applications to Climate Change and Illicit Finance

Alice Lépissier*

Brown University, Providence RI, USA

March 11, 2023

[Latest version](#)

Abstract

This paper is the prologue to my dissertation: “A methodological toolkit to understand complex policy problems: applications to climate change and illicit finance”.

Complex policy problems like climate change and illicit finance require a diverse methodological repertoire and an agnostic approach to selecting the appropriate analytical tool to accomplish discrete inferential tasks. Drawing from the disciplines of political science, economics, and statistical data science, the dissertation presented here tackles three distinct problems on causal evaluation, measurement, and missing data. This paper introduces the broad analytical lens of the dissertation and offers a perspective on empirical research for the study of complex real-world policy problems. The article argues that climate change and illicit finance can be understood as “wicked problems”, and that doing so reveals the epistemological limitations of the common inferential framework that underlies much of the policy-relevant research in applied social sciences. Instead, this paper makes the case that machine learning approaches are uniquely suited to the study of “wicked” problems which resist systematic *a priori* formulation. The inferential framework of machine learning does not require us to accept that there is a simple generative model for the problem that can be known to be true. Instead, machine learning can be deployed in conjunction with domain knowledge to generate policy-relevant insights without requiring strong assumptions on the data-generating process in nature.

*Postdoctoral Fellow, Brown University. Correspondence to: alice_lepissier@brown.edu.
PLEASE DO NOT CITE OR CIRCULATE WITHOUT AUTHOR'S PERMISSION.

Contents

1	Introduction	1
2	The “wicked problems” of climate change and illicit finance	7
3	Revisiting a tale of two cultures	13
4	How machine learning can help study wicked problems	17
5	Permissions and Attributions	24
6	Statement of Scientific Reproducibility	25

1 Introduction

Complex policy problems like climate change and illicit finance require a diverse methodological repertoire and an agnostic approach to formulating the research design that will be most appropriate to addressing the specific question at hand. Climate change and illicit finance fall under the class of “wicked problems”: problems that resist systematic *a priori* formulations and that seem impossible to solve for the social planner due to their social complexity. Wicked problems have no optimal set of solutions, only specific aspects of the problem can be ameliorated. However, certain dimensions of wicked problems can be broken down into “tame” problems that can be solved with discrete inferential tasks. Drawing from the disciplines of political science, economics, and statistical data science, this dissertation addresses three distinct problems in the study of climate change and illicit finance: a causal evaluation problem, a measurement problem, and a missing data problem.

The target of analytical inquiry differs in each case; consequently, this dissertation uses an assortment of methods to answer the questions in a principled way, while proposing specific innovations to correct for methodological difficulties pertaining to each task. In other words, the problems presented here have different estimands that call for separate estimators. In terms of substantive contributions, chapter 2 of this dissertation advances our understanding of climate change policies, while chapters 3 and 4 contribute to the scholarship on illicit finance. Chapter 2 shows how to conduct *ex post* impact evaluations of climate reforms that do not rely on simplistic Business As Usual scenarios to generate a baseline level of carbon emissions. Chapter 3 presents the construction of a new proxy measure for illicit financial flows that addresses long-standing methodological concerns about the extant methods used to detect illicit activity from discrepancies in mirror trade statistics. Finally, chapter 4 proposes a predictive approach to deal with the paucity of data on economic outcomes in developing countries, and demonstrates that machine learning models can be used to credibly augment the database presented in chapter 3. Broadly, the chapters in this dissertation can be understood as operating in the different scientific frame-

works of causal, descriptive, and predictive inference, respectively. Brief abstracts of the chapters are provided next.

Chapter 2 evaluates the causal effect of a climate mitigation policy on the carbon dioxide emissions of the United Kingdom. Here, the estimand is the Average Treatment Effect (ATE) of the climate policy, that is, the mean difference in outcomes between CO₂ emissions once the policy was in place and hypothetical emissions if the climate reform had not been passed. Impact evaluation of climate change policies is difficult because the underlying drivers of carbon emissions are complex latent factors, and the adoption of climate policies by governments is not random. Since the failure of climate governance regimes that sought to impose legally binding treaty-based obligations, the Paris Agreement relies on voluntary actions by individual countries. Yet, there is no guarantee that unilateral policies will lead to a decrease in carbon emissions. Critics worry that voluntary climate measures will be weak and ineffective, and insights from political economy imply that regulatory loopholes are likely to benefit carbon-intensive sectors. The chapter empirically evaluates whether unilateral action can still reduce carbon pollution by estimating the causal effect of the United Kingdom's 2001 Climate Change Programme (CCP) on the country's carbon emissions. Existing efforts to evaluate the overall impact of climate policies on national carbon emissions rely on Business-As-Usual (BAU) scenarios to project what carbon emissions would have been without a climate policy. Instead, the chapter uses a synthetic control estimator to undertake an *ex post* national-level assessment of the UK's CCP without relying on parametric BAU assumptions, by constructing a plausible counterfactual for the emissions trajectory of the UK in a world where the Climate Change Programme (CCP) would not be in place. Despite setting lax carbon targets and making substantial concessions to producers, the resulting estimate is that, post-treatment, the UK's CO₂ emissions per capita in 2005 were 9.8% lower relative to what they would have been if the CCP had not been passed. The findings offer empirical confirmation that unilateral climate policies can still reduce carbon emissions, even in the absence of a binding global climate agreement and in the presence of regulatory capture by industry.

Turning now to the emerging academic scholarship in illicit finance, Chapter 3 accomplishes the elemental task of creating a new proxy measure for illicit financial flows from trade misinvoicing – an illicit practice that is used to clandestinely shift money in and out of a country by manipulating the trade invoices presented to customs. Here, the estimand is the population-level quantity of misinvoiced trade – which remains unobservable because illicit financial flows are deliberately hidden. The secrecy of illicit financial flows is an emblematic characteristic of the problem and, as such, much of the academic effort in the literature is geared towards developing credible methods to detect illicit activity from official economic statistics. Thus, the problem can be reduced to asking what can be learned from an unobservable random variable given an observable one. This chapter presents an original methodology and database – the “atlas of misinvoicing” – that provides bilateral estimates of illicit trade, disaggregated by sector, for 167 countries between 2000 and 2018. Existing methods of estimating trade misinvoicing look for discrepancies in mirror trade statistics to locate instances of trade misinvoicing. Yet, these methods have been faulted for uncritically equating trade irregularities with illicit misinvoicing, and have been accused of generating phantom estimates of illicit financial flows that are an illusion created by the statistical artefacts of how countries record international trade transactions. The chapter approaches the problem in a principled way by deriving the properties that a persuasive measure of trade misinvoicing must possess in order to be both theoretically cogent and practically applicable. Then, the “atlas” method proposes several innovations designed to fulfill these criteria and to ameliorate long-standing problems in the literature; including by providing an empirical way to ascertain the trade gaps that result from benign, non-illicit, factors. The “atlas” measure estimates that developing countries lose \$500 billion of dollars in gross outflows a year, with illicit trade costing Africa \$86 billion a year, where trade in natural resource commodities is heavily misinvoiced. The implication of these findings is that combating illicit financial flows from trade misinvoicing will be crucial to allow poor countries to mobilize domestic resources to finance their own sustainable development goals.

While the “atlas” database introduced in chapter 3 is the first of its kind to provide broad country

coverage, it is still missing data from countries who do not report international trade statistics, including 10 African countries. Chapter 4 presents a strategy to address the problem of missing data on economic measures in developing countries, by using machine learning models to predict illicit trade outcomes without requiring data on the observed trade flow for training. Here, the “atlas” database is taken as a measure of ground truth, and the estimand is the amount of trade misinvoicing conditional on observed country-level features denoting unilateral and bilateral characteristics. Missing or poor quality data is a prevalent problem in developing countries due to weak administrative systems for statistical reporting. This hinders the study of trade misinvoicing, which relies on recorded trade declarations by national customs authorities. The patchiness of data on commodity trade flows compounds the prejudice for African countries who are particularly vulnerable to illicit financial flows. Random Forest machine learning models are trained to predict misinvoiced trade on a sample of African countries using only variables that are easily observed, such as distance between countries, or that are readily compiled by researchers and publicly available, such as perceptions of good governance, without requiring data on the observed trade flow. The Random Forest estimators are able to explain 70% to 73% of the variation in illicit trade outcomes on an unseen test set. Placebo trials are conducted to demonstrate the statistical significance of the results, and the generalization performance of the models is characterized using an experiment that tests how well the models “travel” beyond Africa. The results show that the superior predictive performance of the machine learning models is unlikely to be the product of chance, suggesting instead that the models are able to detect meaningful structure in the dyadic nature of countries’ bilateral relationships that is predictive of illicit trade. The findings demonstrate the promise of machine learning as an imputation tool to augment existing measures of development-related outcomes in the data-scarce settings of developing countries.

The remainder of this prologue introduces the broad analytical lens of the dissertation and offers a perspective on empirical research. Analysts who study complex real-world problems and who seek to use rigorous empirical research to drive evidence-based decision-making forward might encounter the disconcerting trade-off between the need to provide relevant and timely insights

that were obtained in a principled manner, and the reluctance to make assumptions beyond what can be credibly accepted in their field. Instead, an agnostic point of view of policy-relevant empirical research emphasizes the value of producing reasoned insights while eschewing making strong assumptions on the data-generating process in nature. Recognizing the “wicked” nature of many complex policy problems suggests that there is no need to accept that there is a simple generative model for the problem that can be known to be true. Wicked problems defy a definitive formulation, and so it follows that different epistemologies will exist about the way to reach a reasoned conclusion. Instead, by breaking off facets of the wicked problem into discrete “tame” problems, this dissertation advances substantive knowledge on climate change and illicit finance using methods that are appropriate to the inferential target of the “tame” problem. The emphasis in this dissertation is on generating credible inferences that are based on substantive assumptions, derived from theory and domain knowledge, and rigorously estimated within the strictures of statistical inference. Moreover, any normative positions and assumptions relating to the social welfare consequences of these problems are transparently and clearly articulated.

Climate change is an existential peril where the urgent imperative of drastically reducing carbon emissions will necessitate a complete transformation of the ways in which our economies organize production and consumption (IPCC, 2014; Rogelj et al., 2018), thus generating inherent conflicts around the distribution of the costs and benefits of climate action (Genovese, 2019; Mildenerger, 2020). Moreover, the glacial pace of progress globally in the last 40 years relative to the starkness of the scientific record reflects disagreements over the meaning of global and intergenerational equity (Barder et al., 2015; Pickering et al., 2015). Similarly, the proliferation of illicit finance both reflects and exacerbates the unequal distribution of the gains of globalization. On the one hand, deeper financial integration has led to the emergence of secrecy jurisdictions – countries whose comparative advantage in the global marketplace is to provide legalized financial opacity, loose regulations, and low taxes – that have provided fertile ground for abuses such as concealing corruption, tax avoidance, and money laundering by powerful actors such as multinational companies, wealthy individuals, and other elites (Christensen, 2012; Shaxson, 2011; Shaxson &

Christensen, 2013). On the other hand, illicit finance threatens to disrupt the fabric of society by entrenching disparities and inflaming the political problems associated with inequality, and continues to jeopardize the prospects for sustainable development in poor countries (Baker, 2005; Lonergan & Blyth, 2020; Reuter, 2012). Therefore, the symptoms of climate change and illicit finance are related to their genesis, a hallmark of “wicked problems”. These complex policy problems seem intractable due to the complexity of competing interests from different stakeholders. Researchers wishing to make progress, piece by piece, grapple with various types of questions to ameliorate different facets of the problem. Therefore, tackling wicked problems will require a good dose of humility and a willingness to traverse methodological siloes in order to identify the best tool for the job at hand.

While shunning absolutist claims to knowledge – some tools will be useful some times to solve some aspects of the problem – it is nonetheless possible to distinguish the type of methodological instruments that are well-suited to address a specific manifestation of the problem, and some iconoclasm might even be warranted to outline the limitations of popular approaches in empirical social sciences and suggest specific ways in which other approaches from statistical data science might fill those gaps. This dissertation is an ecumenical collection of papers where the different methodological traditions in empirical social sciences and statistical data science are appreciated for the distinctive value added they confer to the analysis of climate change and illicit finance, and where the comparative advantage of each approach is correctly identified in order to select the inferential framework that will yield the most analytical leverage for the specific “tame” problem at hand. Faced with the real-world urgency of climate change and illicit finance, researchers must know how to direct a suitable methodology to the target of inquiry.

Illustrating this point, the next sections develop the concrete argument that machine learning approaches are uniquely suited to the study of wicked problems. Section 2 presents the attributes of wicked problems and demonstrates that climate change and illicit finance exhibit all the features of wicked problems. In turn, understanding climate change and illicit finance as wicked problems

reveals the epistemological limitations of the common inferential framework that much of applied social sciences is predicated on; one rooted in causal inference where the task of the analyst is to credibly estimate the causal links between independent variables and an outcome. Yet, predictive questions are conspicuous in the field of illicit finance, and as such they require a different mode of statistical inference. While econometric techniques specialize in the consistent estimation of parameters and the interpretability of the resulting coefficients, machine learning excels at predictive tasks. Section 3 further reflects on the two cultures of traditional econometrics approaches and machine learning approaches and sketches their limits and areas of complementarity. Finally, section 4 derives the unique value proposition of machine learning for the analysis of wicked problems from the features of its inferential machinery, and offers some caveats.

2 The “wicked problems” of climate change and illicit finance

Climate change and illicit financial flows (IFFs) fall under the class of “wicked problems”, that is, problems that are difficult or impossible to solve for the social planner (Conklin, 2006; Rittel & Webber, 1973). Wicked problems are hard to define due to their social complexity and are resistant to solutions. Rittel and Webber (1973) propose ten characteristics of wicked problems, which broadly map to either the problem’s formulation or the problem’s solutions. First, there is no definitive formulation of a wicked problem (Rittel & Webber, 1973). There are numerous explanations that can be provided for why a discrepancy representing a wicked problem exists. Since there is always more than one explanation for a phenomenon, the choice of explanation will determine the nature of the problem’s resolution (Rittel & Webber, 1973). The other definitional attributes of a wicked problem are that each problem is essentially unique,¹ and that every wicked

¹It is always possible to find a distinguishing property for any two problems that is trivially unique, but the authors who formalized the theory hold that a wicked problem is *essentially* unique, because one can never be sure that the idiosyncracies of the problem are not more important than any common features that it may share with a similar looking problem (Rittel & Webber, 1973).

problem can be considered to be a symptom of another problem.

Second, the defining traits of a wicked problem that pertain to the nature of its solutions are: wicked problems have no stopping rule; there are no true-or-false solutions to wicked problems, only better or worse ones; there is no immediate and no ultimate test of a solution to a wicked problem; and finally, there is not an enumerable set of solutions to attempt (Rittel & Webber, 1973). Furthermore, Rittel and Webber (1973) theorize that every attempt at a solution to a wicked problem is consequential and leaves traces that cannot be undone. There is no opportunity to learn by trial-and-error; every effort is a “one-shot operation” that changes the policy space and future solution set. Wicked problems come with an intrinsic challenge to designing policy solutions to combat them, because the process of conceiving a solution is identical to the process of understanding the nature of the problem; in Rittel and Webber (1973)'s view, “[t]he formulation of a wicked problem is the problem!” (p. 161).

In addition to complicating the design of practical policy solutions, wicked problems also entail a more subtle epistemological implication that is worth highlighting. In the definition of wicked problems proposed by Rittel and Webber (1973), the social planner “has no right to be wrong” (p. 166), which is taken to mean that the planner is liable for the consequences of the solutions they enact. A Popperian approach to social policy would construe solutions to problems as hypotheses that are put forward for refutation and thus, it follows that they should be falsifiable with evidence (Rittel & Webber, 1973). Should it be repeatedly demonstrated that the implemented solution fails to reject the null hypothesis of no effect, then the planner’s confidence in this policy solution would decrease as a result; and alternative policy interventions can then be designed and their performance can be similarly evaluated using this approach. By contrast, wicked problems do not have falsifiable solutions where the aim is to find out the truth; rather, the goal is to ameliorate some of the characteristics of the problem. Since the boundaries of wicked problems are hard to delineate, “[t]he planner who works with open systems is caught up in the ambiguity of their causal webs” Rittel and Webber (1973, p. 167). This suggests that the common epistemological

approach that underpins much of applied social sciences and public policy work – one rooted in causal inference where the main exercise consists of evaluating either the determinants or consequences of social phenomena – has some limitations. Indeed, Rittel and Webber (1973) argue that “[i]n dealing with wicked problems, the modes of reasoning used in the argument are much richer than those permissible in the scientific discourse” (p. 166). Interpreting the former as an invitation, this dissertation seeks to offer modest contributions to the analysis of climate change and illicit finance by leveraging a range of inferential tools, each designed to address specific methodological challenges that arise from the type of question at hand, grappling with policy evaluation, measure construction, and missing data in turn.

The phenomena of climate change and IFFs demonstrably meet all the criteria of wicked problems. The irreversibility of certain courses of climate action and the potential for unintended consequences, the absence of a discrete set of permissible solutions, and the multiplicity of stakeholders with competing views of the problem and of how to solve it, all conspire to make climate change a wicked problem. Climate change is perhaps the “super wicked” problem of our time, given that time is also running out, those who caused the problem are now seeking to solve it, and policy interventions irrationally discount the future (Levin et al., 2012). The analytical position that climate change is a wicked problem has long been understood, and I do not linger on it further (see, e.g., Haug et al. (2010) and Hildén (2011)). However, the academic field of illicit finance is still in its infancy, and conceptual treatments of the problem are sorely needed (Cobham & Janský, 2020; Reuter, 2012). Moreover, two out of the three chapters in this dissertation pertain to illicit finance, so I outline below the features of IFFs that identify it as a wicked problem.

There is no definitive formulation of illicit financial flows: unresolved debates remain on whether the definition of IFFs should be a narrow legalistic one (i.e., restricting the categorization to flows that stem from activities in direct contravention of laws) or a broader one where normative considerations are included (e.g., taking the position that while aggressive tax *avoidance* by multinational corporations is legal, unlike tax *evasion* which is illegal, it is harmful and so it

should be treated as an IFF, see Blankenburg and Khan (2012) and Cobham et al. (2014)). This in turn has practical implications for how policy-makers respond to the phenomenon; in fact, the setting of the United Nations Sustainable Development Goals (SDGs) left open the question of definition and measurement, simply concluding that “by 2030, [the world should] significantly reduce illicit financial flows and arms flows, strengthen the recovery and return of stolen assets and combat all forms of organized crime” (SDG 16.4, UN General Assembly (2015)).

Moreover, the fact that observed irregularities can be interpreted as representing the wicked problem differently is a particularly apt descriptor of IFFs such as trade misinvoicing. For any given gap in bilateral trade statistics that is observed, there exist plausible alternative explanations for the gap that cannot easily be parsed. This is the subject of chapter 3 which proposes a methodology to overcome this problem, by using a systematic way of distinguishing between benign discrepancies in trade gaps and ones that can be ascribed to illicit activity. The various possible ways of representing the problem in turn inform the nature of the proposed solutions. For instance, an outflow from country i to country j could be conceptualized as a problem of cash smuggling or of the embezzlement of public funds, or both. These two ways of problematizing the phenomenon would tend to suggest different policy interventions: the first one emphasizing the role of currency and capital controls, while the second would put the onus on boosting good governance.

As outlined above, wicked problems can always be viewed as symptomatic of other problems. The case of trade-based money laundering could be viewed, for example, as a symptom of grand corruption, which is the abuse of office committed by high-level public officials in pursuit of illicit enrichment. This inherent difficulty is exemplified in most of today's international anti-money laundering regimes, including the Warsaw Convention, by the fact that predicate offenses (that is, the underlying crimes) may be used to establish a charge of money laundering.² In turn, the subversion of the functions of the state associated with grand corruption can be viewed as a

²FATF Recommendation 3 calls on countries to criminalize the laundering of proceeds of all serious offences, with a view to including the widest range of predicate offences. The Warsaw Convention of 2005 (CETS 198) provides a legal framework for charging predicate offenses.

symptom of the erosion of the social contract between a government and its citizens. Therefore, as soon as one causal explanation is proposed, other causes must be accounted for; and it is hard to dispute that the causal arrow between variables goes in both directions, when we note that the depletion of state coffers and looting of domestic resources occasioned by IFFs further weaken the social contract. As a result, usual tools of causal inference, such as Directed Acyclic Graphs (DAGs) for example, cannot be used since they presuppose unidirectional causal relationships (a *directed* graph) and the absence of feedback loops between variables (an *acyclical* graph) (G. W. Imbens, 2020; Pearl, 1995, 2009).

The endogeneity of IFFs complicates the prescription of appropriate solutions, and existing policy interventions on IFFs exhibit the same features as those of the solutions to a wicked problem. Combating illicit finance is akin to playing the arcade game “whack-a-mole”, where attempts to solve the problem are often piecemeal and only result in temporary improvements: as soon as one mole is disposed of, another one will emerge from its hole. This can be seen most clearly in the case of trade misinvoicing to launder the proceeds of transnational organized crime: if one accepts the premise that crime will always exist (since most societies organize their activities around laws dictating permissible and impermissible behaviors), then there must always be some irreducible amount of illicit finance that is not caught since illicit activity, by definition, seeks to remain hidden.

Moreover, Lowery and Ramachandran (2015) have provided evidence of the unintended consequences of anti-money laundering policies, which is congruent with the idea that attempted solutions for a wicked problem are consequential. Currently, banks are required to put in place good faith efforts to stall sanctions violations, money laundering, and terrorist financing. Collectively, these rules are commonly referred to as AML/CFT (Anti-Money Laundering/Countering Financing of Terror) regimes. Given that banks are subject to large fines if they fail to perform their due diligence, many banks are altogether exiting certain sectors they deem too risky in a process known as “de-risking”, which has had dire consequences for poor countries in particular. For

example, widespread denials of banking accounts to money transfer organizations have increased the cost of remittances (the money that migrants send home) (Lowery & Ramachandran, 2015). Given that remittance flows are the largest source of external financing for low- and lower-middle income countries, above foreign direct investment and official development assistance (World Bank et al., 2016), the AML/CFT procedures put in place that were supposed to combat illicit financial outflows have aggravated the draining of resources in developing countries.

Problematizing illicit financial flows and climate change as wicked problems suggests that the way social science research can generate meaningful insights is to divide the wicked problem into discrete “tame” problems to ameliorate a specific dimension of the problem; this is the strategy that underpins the analytical enterprise of this dissertation. In some cases, the study of a wicked problem will benefit from working on a causal problem; this is typically the preferred approach to conducting policy impact evaluation or to analyzing the determinants of a specific manifestation of the wicked problem. Chapter 2 sets out to causally evaluate the effectiveness of a policy designed to reduce carbon emissions. Similarly, in the field of illicit finance, Allred et al. (2017) have employed the reputed gold standard in the causal inference toolkit (Gerber & Green, 2012) – a randomized experiment – to show that firms in OECD countries (compared to those in tax havens) were more willing to provide anonymous incorporations, i.e., create shell companies, and to flout international rules on financial transparency (Allred et al., 2017). However, it is not always possible for social scientists to either carefully design experiments or to identify natural experiments that provide a source of exogenous variation and allow for causal identification. When the latent factors driving a problem are complex and multiple, it is hard to provide actionable insights to policy-makers regarding the policy interventions that they can undertake to combat the problem. Pol (2020) blames the failure of existing anti-money laundering (AML) regimes, dubbed the “the world’s least effective policy experiment” (p. 73), on the mismatch between the way that outcomes are understood and the design principles of current policy prescriptions. Likewise, isolating the marginal causal effect of a particular determinant on the outcome is of limited value for public decision-making if it cannot be manipulated by the social planner. In other

cases, however, there exist discrete tasks associated with wicked problems that can be usefully accomplished in the different frameworks of descriptive and predictive inference, as demonstrated by chapters 3 and 4 of this dissertation, respectively.

3 Revisiting a tale of two cultures

In a provocative paper, Breiman (2001) issued a polemic against the overreliance of statisticians on data models which assume that the data are generated according to a stochastic model. Instead, he argues, statisticians should make space for algorithmic approaches to learning from data, where the data mechanism is treated as unknown. According to him, the almost exclusive dependence of statisticians on data models, predicated on the belief that the analyst “by imagination and by looking at the data, can invent a reasonably good parametric class of models for a complex mechanism devised by nature” (Breiman, 2001, p. 202), has led to irrelevant theory and questionable scientific conclusions. This is because any quantitative conclusion that is drawn from fitting a model to data will be a conclusion about the model’s mechanism and not about nature’s mechanism, and hence it follows that “[i]f the model is a poor emulation of nature, the conclusions may be wrong” (Breiman, 2001, p. 202). Breiman’s prescient piece foresaw the emergence of machine learning (ML) into the mainstream of many of today’s scientific enterprises.

At around the same time, a critique reminiscent of Breiman’s appeared in political science with Brady and Collier (2004)’s pushback on the “quantitative imperialism” of King et al. (1994)’s *Designing Social Inquiry*, a treatise on how qualitative research should follow the precepts of quantitative social science. Brady and Collier (2004) resisted the dogmatism of mainstream quantitative methods by pointing out that regression analysis relies on the difficult-to-test assumption that the model being estimated is correct. Though the depth of the qualitative-quantitative divide in political science is sometimes overstated when considering that both still belong to a positivist framework of social science (cf. Goertz and Mahoney (2012) and Mildemberger (2016)), this type of criticism from the qualitative culture of research is a salutary reminder that econometric work

packs a lot of assumptions on the nature of the problem in the course of its efforts to understand it.

Economists have taken up Breiman's challenge and an influential literature is emerging in applied economics that bridges the two cultures, notably by applying machine learning algorithms to causal inference problems (Athey & Imbens, 2015; Athey & Imbens, 2019; Athey et al., 2017, 2019; G. Imbens & Athey, 2021; Mullainathan & Spiess, 2017; Storm et al., 2020; Varian, 2014). Political science is also increasingly receptive to this view (Abadie et al., 2010; Brady, 2019; Diamond & Sekhon, 2013; Grimmer et al., 2021; Radford & Joseph, 2020), and some authors have proposed an agnostic analytical framework that centers on what is learnable from the world without assuming that there exists a simple generative model that can be known to be true (Aronow & Miller, 2019; Grimmer et al., 2021).

How can machine learning provide added value in social scientific settings, particularly for the study of wicked problems? As a first stage, it is useful to draw some distinctions between machine learning approaches and traditional econometric approaches in applied social sciences. To some extent, applied empirical researchers in political science and economics share much of the same methodological toolkit (cf., for example, Angrist and Pischke (2009), King et al. (1994), and Wooldridge (2010)). Though there are areas of discord and some turf wars between the two camps, applied economists and political scientists have more in common than separates them, and with the indulgence of the reader, I will refer to this as the econometric approach hereafter.

The econometric approach tends to be concerned with parameter estimation while the machine learning approach focuses on prediction (Athey & Imbens, 2019; Athey et al., 2019; Mullainathan & Spiess, 2017). Supervised machine learning techniques aim to estimate the conditional expectation $\mathbb{E}[Y|X]$ of a response Y given a set of variables X (called "features" in the ML literature). By contrast, econometric techniques seek to provide estimates of parameters β that govern the relationship between the explanatory variables X and the outcome Y . In the canonical use case, the analyst will first specify a functional form for the underlying assumed "true" relationship

between Y and X , of the general form $\mathbb{E}[Y|X] = X^T\beta + \epsilon$, and then seek to estimate the parameters $\hat{\beta}$ using a linear regression model. It is important to notice that, although the predictors X can enter the relationship in non-linear forms (including interactions with other variables or as higher-order polynomials X^k of the k th degree),³ the models are *linear in the parameters* β . This is also the case for Generalized Linear Models (GLM), another commonly used class of models in social sciences designed to deal with limited dependent variables.⁴ This already imposes a great deal of structure on the problem. Much effort in econometrics is spent on deploying estimators that have useful asymptotic properties such as consistency and normality.⁵ Statistical inference is then performed on the estimated $\hat{\beta}$ under some regularity conditions.

By contrast, machine learning approaches tend to make less assumptions about the data-generating process that is responsible for the phenomenon we observe. Some ML techniques do make distributional assumptions (e.g., LDA or QDA), but most are non-parametric methods that allow flexible functional forms, and only require that observations are independent to work.⁶ The machine learning approach specifies the relationship between X and Y in very general terms: $Y = f(X) + \epsilon$. The observed response Y is some unknown function f of the predictors plus some irreducible error ϵ . Most machine learning tasks are predictive and seek to generate predictions of the response by estimating f , but make no assumption about the functional form that f will have; the estimated function can end up being very complex. Thus, the problem is posed as

³With the only mechanical limitations on estimating complexity – absent the common sense of the analyst – being computational power and the number of degrees of freedom available. By contrast, machine learning techniques have explicit procedures to choose the amount of complexity in a model.

⁴The linear regression model is not appropriate when the range of the outcome Y is restricted, as is often the case with the variables that social scientists work with, such as count (e.g., population) or categorical (e.g., “voted”, “did not vote”) data. In that case, GLMs allow the analyst to work with limited dependent variables by relaxing the constraints on the mean of the dependent variable $\mathbb{E}[Y_i|X_i] = \mu_i$. This is accomplished by using a possibly non-linear link function $g(\mu_i) = \eta_i$ that specifies how the mean relates to a *linear* function of the explanatory variables and of the parameters β_1, \dots, β_k : $\eta_i = \sum_{k=1}^K \beta_k x_{ik}$. Since the linear predictor η_i can take any value in $(-\infty, \infty)$ while the range of Y is limited, the goal of the link function is to relax the constraints on the mean so that it maps onto \mathfrak{R} and to define the scale over which η_i is *additive*. The point here is to note that these models once again estimate a linear function of the parameters β .

⁵An estimator is consistent when, as sample size grows to infinity, its sampling distribution converges in probability to the true parameter value β^* . Normality of the errors is required to provide valid statistical inferences.

⁶The other substantive assumption that underpins machine learning is that the data in the training and test set are drawn i.i.d. from the same (unknown) distribution (Athey et al., 2019; Breiman, 2001), but I address this later.

estimating $\hat{Y} = \hat{f}(X)$. By contrast, econometric approaches impose many more assumptions on the functional form that f will take.

As I illustrated above with the canonical linear regression model, the assumption in many of the most commonly used econometric models in social sciences is that f is linearly additive in its parameters, yet that is not often surfaced. Specifically, linear regression models assume $\mathbb{E}[Y|X] = \beta_0 + \beta_1 X_1 + \dots + \beta_k X_k + \epsilon$. Many policy-oriented problems in social sciences revolve around identifying the causal effect β_k of a treatment X_k . Policy-oriented analyses are interested in providing estimates of the *marginal* effect of the treatment of interest. That is, we seek to know what the causal effect is of increasing X_k by one unit, holding other variables X_1, \dots, X_{k-1} constant. But it could be the case that the model f is misspecified, and the true data-generating process contains both non-linearities and interactions, which would complicate the interpretation of $\hat{\beta}_k$ as a marginal effect.⁷ When policy-relevant estimands, e.g., the Average Treatment Effect (ATE), are defined in terms of a linear model, the consequences of misspecifying the model f might lead to non-negligible welfare loss if treatment effect estimates are used to guide policy decisions.

The goals of the two frameworks are fundamentally different. Parameter estimation approaches set out to find an unbiased estimator $\hat{\beta}$ that minimizes in-sample error, while machine learning approaches aim to minimize the prediction error of \hat{Y}_i for a *new* data point i . Thus, econometric approaches are focused on unbiasedness by construction whereas machine learning techniques provide an empirical way to manage the bias-variance trade-off.⁸ The difference between both frameworks is visible in how they approach validation: econometric approaches will rely on in-sample goodness-of-fit measures to judge analytical power, while machine learning approaches assess performance by evaluating predictive accuracy on an out-of-sample observation. Since

⁷For example, it could be the case that the true data-generating process contains interactions, e.g., $Y = \beta_0 + \beta_1 X_1 + \dots + \beta_k (X_1, \dots, X_{k-1}) X_k + \epsilon$ where the treatment effect of interest β_k is a (possibly very complicated) function of all other covariates, or the true model could be non-linear where the effect of X_k changes with the level of X_k (so that not all unit changes are equal). I thank Alex Franks for pointing this out to me.

⁸Note that $\hat{\beta}_{OLS}$ is the Best Linear Unbiased Estimator and so by construction makes a choice in the bias-variance trade-off by ensuring zero bias; whereas ML optimizes with respect to both in-sample and out-of-sample error (Kleinberg et al., 2015).

$\hat{Y} = \hat{f}(X)$ is a deterministic function of \hat{f} given X , the problem of providing the best prediction of \hat{Y}_i for a new observation i reduces to finding the functional form of f that minimizes a chosen loss function $\mathcal{L}(Y_i, \hat{f}(X_i))$ from a set of functions $\hat{f} \in \mathcal{F}$.⁹ This is possible because we can observe prediction quality, whereas in econometric approaches we need to make assumptions about f to ensure consistency (Mullainathan & Spiess, 2017). In other words, econometric approaches *assume* the functional form of f , while machine learning approaches *learn* the representation of f from the data.

4 How machine learning can help study wicked problems

Machine learning (ML) approaches in general, and deep learning techniques such as neural networks in particular,¹⁰ are well-suited to deal with problems where the underlying nature of the phenomenon is complex, dynamic, and resists *a priori* and systematic characterizations; all of which are properties of wicked problems. In other fields, machine learning has had spectacular success in accomplishing tasks where it is similarly difficult to come up with a set of hard-wired rules to follow: object recognition, translating natural languages, generating artificial but photo-realistic images, etc. What these tasks have in common is that they seem to require some degree of “intelligence” because it is not possible to pre-specify exhaustively and comprehensively the set of procedures that an agent must follow in order to correctly accomplish the task. For example, to accurately identify and distinguish between images of dogs and cats, there is no set recipe to follow that would, say, direct us to first look for triangular ears and then to look at

⁹Crucially, the set of candidate functions \mathcal{F} is not restricted to a set of linear predictors. The loss functions that are often picked are the Mean Square Error (MSE) in regression problems and the misclassification error rate in classification problems.

¹⁰In this paper, I do not spend time on definitional debates about what constitutes artificial intelligence (AI), machine learning (ML), and deep learning (DL), and what distinguishes them from statistics. Instead, I adopt the simplest operational definitions used across computer science (Goodfellow et al., 2016) and computational social science (Brady, 2019). AI is a broad concept that refers to computer-led intelligent tasks. ML is a subset of AI where tasks are accomplished by feeding data to an algorithm that improves the more data it is exposed to. DL is a subset of ML where learning occurs through a hierarchy of concepts, where each concept is defined through its relation to a simpler concept, and where the computer learns the complicated concept by building on the simpler one (Goodfellow et al., 2016; LeCun et al., 2015).

the length of the snout before arriving at a decision. Mullainathan and Spiess (2017) contend that the real breakthrough occurred when analysts stopped approaching intelligence tasks procedurally, but instead approached them empirically. For example, in the case of computer vision and image recognition – where models called Convolutional Neural Networks have distinguished themselves (Goodfellow et al., 2016; LeCun et al., 2015) – the algorithm will first learn a set of low-dimensional features such as edges and shadows which are then encoded into progressively more complex elements (e.g., an ear). The point is that the analyst does not attempt to deduce the rules that would help discriminate between a dog and a cat; instead the algorithm lets the data find the rules that work best to accomplish the task at hand.

Though in social sciences we tend to work with structured tabular data rather than pixels, the principle is the same. Complex socio-political phenomena are difficult to capture and represent as a set of equations. Certainly, the role of theory is to come up with a parsimonious model of that phenomenon, but it can only ever be a stylized representation of the “true” state of the world, and models more often than not impose strong assumptions on the nature of the problem. However, with a wicked problem, we do not actually know the definitive nature of the problem until we have solved it. In other words, we struggle to specify the true data-generating process, because it might contain non-linearities and complex interactions in many dimensions. Yet machine learning techniques are exceptionally good at fitting flexible functional forms. A neural network, for instance, can approximate any continuous function arbitrarily well (Hastie et al., 2017). Machine learning tools estimate the conditional value of an outcome variable given a set of independent variables, but without making many assumptions about the structure of the relationship. And in the case of intractable wicked problems, I submit that this is a desirable attribute. More worryingly, by explicitly attempting to specify the complex relationships that underpin a wicked problem, we will – according to the definition of a wicked problem – misspecify the model. Machine learning techniques are thus particularly good at finding generalizable structure in a problem because the functional form is determined by the data. That is, model selection occurs automatically in the process of finding the best predictions of the outcome.

There are two other situations that are relevant to wicked problems where machine learning methods may be advantageous: (1) in the case of heterogeneous treatment effects, and (2) in high dimensional settings. Machine learning can illuminate heterogeneous “treatment” effects where there exists heterogeneity in responses with respect to observed covariates (Athey et al., 2019; Storm et al., 2020). In a linear regression setting, heterogeneous effects might be estimated via the coefficients on interaction terms (Mullainathan & Spiess, 2017), yet they can also be construed as a prediction problem of mapping unit-level attributes to individual effect estimates. This view underpins many existing marketing and recommendation systems in business applications (e.g., a web page renders customized advertisements according to the browsing habits of the internet user) and is being increasingly used in precision medicine to provide individualized treatment plans (Ge et al., 2020; Obermeyer & Emanuel, 2016). Predictive tasks can also be helpful to tackling policy problems. I demonstrate this in chapter 4 of the dissertation where the problem of missing data in developing countries can be abated by generating machine learning predictions that can be used to augment the measure of illicit finance that is developed in chapter 3.

Another attractive feature of machine learning is that it offers a principled way of dealing with high dimensional settings, which is where there are a large number of covariates K , sometimes to the point where $K \gg N$ and the number of variables is much higher than the number of observations. High dimensionality poses several challenges. First, when $K > N$ estimation is infeasible, since we have more regressors than data points and hence negative degrees of freedom. Further, when there are many covariates that we think might interact to predict the outcome but we don’t know how, it would not be sensible to, say, include all pairwise interactions of the explanatory variables in our model. By contrast, model selection with machine learning is data-driven and systematic, because ML searches for those interactions automatically. A key way in which machine learning deals with high dimensional settings is through regularization, a procedure that penalizes variables that are not informative. In standard econometric approaches, high dimensionality will also manifest as multicollinearity, where some variables are highly correlated with each other, which would be reflected by large standard errors on those coefficients.

Because most ML techniques do not provide standard errors, multicollinearity in ML settings is effectively dealt with through regularization. The data-driven approach of ML in dealing with high dimensional settings is useful for wicked problems, where the covariate space is plausibly large. Given that in wicked problems, the formulation of the problem is inextricable from the planner's (subjective) view of how to solve it, it might be productive to let the data speak for itself and use a transparent and empirical approach to extracting the predictors that have the most informational value for explaining the variation in the outcome.

A particularly useful heuristic for the study of wicked problems that can be borrowed from machine learning approaches is the concept of regularization. Regularization in ML is the process of selecting important variables in a data-driven way. These techniques can help us extract information from a messy and noisy dataset and “distill the essence of the data” (Grimmer et al., 2021, p. 402). Implicit in this view is the notion that we can represent a high-dimensional covariate space in lower-dimensional space because there are some key latent factors that are responsible for the phenomenon we observe. The core idea behind regularization – that a high dimensional space can be expressed in low-dimensional terms – offers us a lot of analytical traction when dealing with wicked problems, because it allows analysts to delineate the scope of the problem by identifying the variables or dimensions that explain a lot of the variance in the outcome. Wicked problems encompass so much social complexity that it seems quixotic to attempt to completely characterize them with a parametric model, yet we can start to make progress by thinking of them as being constituted of latent factors that operate in unknown and complex ways to produce the phenomenon that we observe. Those latent factors operate in a low-dimensional space that can be interpreted as representing an underlying property of the wicked problem (Grimmer et al., 2021).¹¹ By definition, there is no definitive formulation of a wicked problem: we do not know f . The generative process in nature that gives rise to illicit finance Y is a black box that contains some predictive features, some parameters, and noise. Machine learning techniques are criticized

¹¹Note that regularization on high-dimensional data does necessitate the assumption that the true model is “sparse” (Belloni et al., 2014): that it is possible to reconstruct the data relatively well using a low-dimensional representation of the covariate space.

for precisely this reason: the function $f(X)$ that maps the feature space X to the outcome Y is a black box. The way that ML deals with the “black box” problem is to let the data tell us what variables are important through regularization. This then allows for dimension reduction, which is critical to drawing meaningful conclusions from the data. Machine learning approaches take the position that it is better to use a very flexible model constrained by regularization than to constrain the model *ex ante* by using few predictors (Grimmer et al., 2021).

The classic bias-variance trade-off means that a very complex model might overfit the training data and generalize poorly out-of-sample. ML methods deal with this by adding a penalty parameter to the model that discourages model complexity. The severity of the penalty parameter is determined empirically by tuning the model (usually through a process known as cross-validation) to find the amount of flexibility in a model that yields the best predictive performance. The data-driven way in which ML tunes model parameters is well-suited for wicked problems, since model selection is accomplished empirically, rather than by a subjective choice of the analyst. However, one common critique levelled against machine learning is that the estimated representation of the data may be difficult to interpret. However, there is no guarantee that the tuned model will be simple or easy to interpret. Occam’s razor holds that there is a trade-off between parsimony and accuracy: more complex models are usually better predictors than simple and interpretable ones. This is a valid concern, but Breiman (2001) proposes an alternative outlook on the topic that can be a useful frame for the study of wicked problems: the goal is to extract information, and interpretability is the means to that end. A model does not have to be simple to provide reliable information about the relationship between X and Y .

Machine learning is not without its drawbacks. Next, I highlight three challenges in the application of machine learning to the study of wicked problems: the difficulty in incorporating uncertainty around parameter estimates, the fact that coefficient estimates $\hat{\beta}$ are rarely consistent, and the pitfalls of naively inferring structure about the data-generating model from machine learning outputs.

First, one drawback of ML techniques is that it is difficult to obtain correct standard errors on the *coefficient* estimates because the data was used for model selection. In the econometric approach, researchers will rely on statistical theory to estimate confidence intervals for their estimated parameters (Athey et al., 2019). By contrast, ML methods have struggled with providing valid confidence intervals, even if only asymptotically (Grimmer et al., 2021). This is because it is hard to know how to incorporate uncertainty when the data itself has been used for model selection (Athey & Imbens, 2015; Athey et al., 2019; Mullainathan & Spiess, 2017). Since ML techniques use the “training” data to learn the representation of f that yields the best predictive performance, any parameter estimates that are generated in the course of estimating the model will have to reflect uncertainty around model selection itself. Mullainathan and Spiess (2017) point out that this is a ubiquitous problem in machine learning: not only do the lack of standard errors make it hard to make inferences on parameters *after* model selection, but this is also a problem of the consistency of the model selection itself. Therefore, ML algorithms have to be modified to provide valid confidence intervals for estimated parameters when the data is used to select the model (Athey et al., 2019). One approach proposed by Athey and Imbens (2016) relies on what they call an “honest” approach to estimation, which is accomplished through sample-splitting. With an application to decision trees, they use one sample to construct the partition (i.e., fit the model), and the other sample to estimate treatment effects. They show that confidence intervals built around the estimates in the second sample will have nominal coverage (Athey & Imbens, 2016).

Even though the empirical approach to model selection in ML makes it a challenge to provide standard errors, there is one ancillary benefit that algorithmic model selection offers. It is a systematic, data-driven way of selecting a model that provides both superior performance and is reproducible (Athey et al., 2019). Athey et al. (2019) makes the point that, in practice, applied researchers may test a variety of econometric specifications behind the scenes when they are performing model selection, yet only report a few specifications as a robustness check. According to Athey et al. (2019), this practice is rampant yet researchers are not often honest about it

because it would invalidate many of the reported confidence intervals around their coefficients due to the multiple comparison problem. Note that this is still a problem, even if researchers are otherwise honest and do not engage in p-hacking or cherry-picking specifications (Ferman et al., 2020). Compared to this, the algorithmic approach to specification searches has the advantage of being systematic, transparent, and reproducible.

Second, another drawback of machine learning concerns the consistency of the $\hat{\beta}$ s. Even when a ML model produces coefficient estimates $\hat{\beta}$ in the course of making its predictions, these estimates are rarely consistent (Mullainathan & Spiess, 2017). Even if a model generates predictions \hat{Y} of the outcome that are robust and of good quality (as measured by predictive accuracy in a test dataset), the coefficients $\hat{\beta}$ on the variables X used to arrive at these predictions can vary with small perturbations of the dataset. Mullainathan and Spiess (2017) perform an experiment that illustrates this problem. They set out to estimate house prices from a variety of predictors (e.g., square footage) from the American Housing Survey. They randomly split the sample into 10 different partitions of 5,000 observations. On each partition, they use a LASSO regression to estimate house prices with a fixed penalty parameter. The LASSO (Least Absolute Shrinkage and Selection Operator) (Tibshirani, 1996) is a regression model that induces sparsity by shrinking some of the coefficient estimates to exactly 0, with the amount of shrinkage determined by a given penalty parameter. Therefore, when estimating a LASSO model, not every variable ends up being explanatory of the outcome (if the corresponding estimated $\hat{\beta}$ is 0). By estimating a LASSO on 10 different sub-samples of the data, Mullainathan and Spiess (2017) show that different explanatory variables end up being used each time, such that no stable patterns are detected, even though the R^2 denoting prediction quality remains constant from partition to partition. This is a problem of consistency of the model selection itself. Breiman (2001) refers to this as the Rashomon Effect of the multiplicity of good models:¹² there exist multiple models of the data $f(X)$ that are just as good to predict the outcome Y . That is, different functional

¹²Named after a movie where the plot revolves around 4 people testifying at a trial about their recollection of the same crime that they witnessed. Even though the facts of the crime that the 4 people report are the same, the stories that they recount about what happened are very different (Breiman, 2001).

forms that combine the variables X in multiple ways can yield the same prediction error rate. This effect is likely to be at play with wicked problems too: there are multiple ways in which the problem can be characterized that are just as good, “good” depending on the conceptualization of the planner.

A final caveat with the application of ML to wicked problems is that it is hard to recover the structure of the model from the estimated coefficients. Mullainathan and Spiess (2017) warn against the temptation to use the estimated $\hat{\beta}$ parameters to try and learn something about the underlying data-generating process. After all, when predictive performance is high, some structure in \hat{Y} must have been found. However, the lack of consistency of the parameter estimates will prevent us from making credible inferences on the underlying structure. On the other hand, Mullainathan and Spiess (2017) make the point that making some assumptions on the data-generating process would allow us to take the $\hat{\beta}$ more seriously. Thus, the challenge of the analyst studying wicked problems is to decide whether to try and recover some of the structure of the problem by placing some assumptions on the data-generating process that delineate the problem *ex ante*, or whether to stick to predictive tasks only. Ultimately, this is a judgement call for the analyst depending on the question at hand. In chapter 4 of this dissertation, the latter option is chosen because the task is a purely predictive exercise that seeks to generate reliable predictions in order to address a missing data problem.

5 Permissions and Attributions

1. The contents of chapter 2 and appendix A are the result of a collaboration between Alice Lépissier and Matto Mildenerger, and have previously been published as: Lépissier, A. Mildenerger, M. (2021). “Unilateral climate policies can substantially reduce national carbon pollution”. *Climatic Change*, 166:31. DOI:10.1007/s10584-021-03111-2. The article is made available under a Creative Commons Attribution 4.0 International License.¹³

¹³To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

2. Some of the contents of chapter 3 are the result of a collaboration between Alice Lépissier, William Davis, and Gamal Ibrahim. The methodology section of this chapter draws on an unpublished methodological note prepared for the United Nations Economic Commission for Africa (UNECA) by Alice Lépissier. Part of this work was carried out while Alice Lépissier was employed as a consultant for UNECA during portions of 2018 and 2019. The methodology has since been substantially refined. Results based on an early version of this methodology are included in the *Financing for Sustainable Development Report 2019* of the Inter-agency Task Force on Financing for Development (New York, NY: United Nations), available at <https://developmentfinance.un.org/sites/developmentfinance.un.org/files/FSDR2019.pdf>. The contents of the methodology and findings sections in this chapter are reproduced with the permission of William Davis and Gamal Ibrahim. The methodology is developed in the authors' personal capacity and does not necessarily reflect the views of their respective institutions. Alice Lépissier gratefully acknowledges financial support from the United Nations Economic Commission for Africa.

6 Statement of Scientific Reproducibility

Every effort has been made to provide an entirely reproducible analytical pipeline; starting with the acquisition of raw data, data cleaning, statistical analyses, and generating results and data visualizations in the final step. The code base for each project in this dissertation is publicly available in online repositories. The input data required to produce results for each project is available online (data acquisition is either automated or details on the sources are provided in the relevant script files). Additional data products are available upon request.

1. The code for chapter 2 is available at <https://github.com/walice/synth>. Raw data and results are available at <https://doi.org/10.5281/zenodo.4566803>.
2. The code for chapter 3 is available at <https://github.com/walice/Trade-IFF>. The full "atlas" database is available at <https://doi.org/10.5281/zenodo.3610557>.

3. The code for chapter 4 is available at <https://github.com/walice/illicitAI>.

References

- Abadie, A., Diamond, A., & Hainmueller, J. (2010). Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program. *Journal of the American Statistical Association*, *105*(490), 493–505. <https://doi.org/10.1198/jasa.2009.ap08746>
- Allred, B. B., Findley, M. G., Nielson, D., & Sharman, J. C. (2017). Anonymous shell companies: A global audit study and field experiment in 176 countries. *Journal of International Business Studies*, *48*(5), 596–619. <https://doi.org/10.1057/s41267-016-0047-7>
- Angrist, J. D., & Pischke, J.-S. (2009). *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press.
- Aronow, P. M., & Miller, B. T. (2019). *Foundations of Agnostic Statistics*. Cambridge University Press. <https://doi.org/10.1017/9781316831762>
- Athey, S., & Imbens, G. (2015). A Measure of Robustness to Misspecification. *American Economic Review*, *105*(5), 476–480. <https://doi.org/10.1257/aer.p20151020>
- Athey, S., & Imbens, G. (2016). Recursive partitioning for heterogeneous causal effects. *Proceedings of the National Academy of Sciences*, *113*(27), 7353–7360. <https://doi.org/10.1073/pnas.1510489113>
- Athey, S., Imbens, G., Pham, T., & Wager, S. (2017). Estimating Average Treatment Effects: Supplementary Analyses and Remaining Challenges. *American Economic Review*, *107*(5), 278–281. <https://doi.org/10.1257/aer.p20171042>
- Athey, S., & Imbens, G. W. (2019). Machine Learning Methods That Economists Should Know About. *Annual Review of Economics*, *11*(1), 685–725. <https://doi.org/10.1146/annurev-economics-080217-053433>

- Athey, S., Tibshirani, J., & Wager, S. (2019). Generalized random forests. *The Annals of Statistics*, 47(2), 1148–1178. <https://doi.org/10.1214/18-AOS1709>
- Baker, R. W. (2005). *Capitalism's Achilles Heel: Dirty Money and How to Renew the Free-Market System*. John Wiley & Sons.
- Barder, O., Lépiessier, A., & Evans, A. (2015). A global emissions budget. In *Climate 2020 - Facing the future* (pp. 121–123). United Nations Association - UK.
- Belloni, A., Chernozhukov, V., & Hansen, C. (2014). High-Dimensional Methods and Inference on Structural and Treatment Effects. *Journal of Economic Perspectives*, 28(2), 29–50. <https://doi.org/10.1257/jep.28.2.29>
- Blankenburg, S., & Khan, M. (2012). Governance and Illicit Flows. In P. Reuter (Ed.), *Draining Development? Controlling Flows of Illicit Funds from Developing Countries* (pp. 21–68). World Bank Group.
- Brady, H. E. (2019). The Challenge of Big Data and Data Science. *Annual Review of Political Science*, 22(1), 297–323. <https://doi.org/10.1146/annurev-polisci-090216-023229>
- Brady, H. E., & Collier, D. (2004). *Rethinking social inquiry: Diverse tools, shared standards*. Rowman & Littlefield.
- Breiman, L. (2001). Statistical Modeling: The Two Cultures. *Statistical Science*, 16(3), 199–215.
- Christensen, J. (2012). The hidden trillions: Secrecy, corruption, and the offshore interface. *Crime, Law and Social Change*, 57(3), 325–343. <https://doi.org/10.1007/s10611-011-9347-9>
- Cobham, A., & Janský, P. (2020). *Estimating illicit financial flows: A critical guide to the data, methodologies and findings*.
- Cobham, A., Janský, P., & Prats, A. (2014). Estimating Illicit Flows of Capital via Trade Mispricing: A Forensic Analysis of Data on Switzerland. *Center for Global Development Working Paper*, 350.
- Conklin, J. (2006). *Dialogue Mapping: Building Shared Understanding of Wicked Problems*. John Wiley & Sons.

- Diamond, A., & Sekhon, J. S. (2013). Genetic Matching for Estimating Causal Effects: A General Multivariate Matching Method for Achieving Balance in Observational Studies. *Review of Economics and Statistics*, 95(3), 932–945. https://doi.org/10.1162/REST_a_00318
- Ferman, B., Pinto, C., & Possebom, V. (2020). Cherry Picking with Synthetic Controls (B. Barnow, Ed.). *Journal of Policy Analysis and Management*, 39(2), 510–532. <https://doi.org/10.1002/pam.22206>
- Ge, Q., Huang, X., Fang, S., Guo, S., Liu, Y., Lin, W., & Xiong, M. (2020). Conditional Generative Adversarial Networks for Individualized Treatment Effect Estimation and Treatment Selection. *Frontiers in Genetics*, 11, 585804. <https://doi.org/10.3389/fgene.2020.585804>
- Genovese, F. (2019). Sectors, Pollution, and Trade: How Industrial Interests Shape Domestic Positions on Global Climate Agreements. *International Studies Quarterly*, 63(4), 819–836. <https://doi.org/10.1093/isq/sqz062>
- Gerber, A. S., & Green, D. P. (2012). *Field Experiments: Design, Analysis, and Interpretation*. W. W. Norton & Company.
- Goertz, G., & Mahoney, J. (2012). *A Tale of Two Cultures: Qualitative and Quantitative Research in the Social Sciences*. Princeton University Press.
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning (textbook)*. MIT Press.
- Grimmer, J., Roberts, M. E., & Stewart, B. M. (2021). Machine Learning for Social Science: An Agnostic Approach. *Annual Review of Political Science*, 24(1), 395–419. <https://doi.org/10.1146/annurev-polisci-053119-015921>
- Hastie, T., Tibshirani, R., & Friedman, J. (2017). *The Elements of Statistical Learning - Data Mining, Inference, and Prediction*. Springer Series in Statistics.
- Haug, C., Rayner, T., Jordan, A., Hildingsson, R., Stripple, J., Monni, S., Huitema, D., Massey, E., van Asselt, H., & Berkhout, F. (2010). Navigating the dilemmas of climate policy in Europe: Evidence from policy evaluation studies. *Climatic Change*, 101(3-4), 427–445. <https://doi.org/10.1007/s10584-009-9682-3>

- Hildén, M. (2011). The evolution of climate policies – the role of learning and evaluations. *Journal of Cleaner Production*, 19(16), 1798–1811. <https://doi.org/10.1016/j.jclepro.2011.05.004>
- Imbens, G., & Athey, S. (2021). Breiman’s Two Cultures: A Perspective from Econometrics. *Observational Studies*, 7(1), 127–133. <https://doi.org/10.1353/obs.2021.0028>
- Imbens, G. W. (2020). Potential Outcome and Directed Acyclic Graph Approaches to Causality: Relevance for Empirical Practice in Economics. *Journal of Economic Literature*, 58(4), 1129–1179. <https://doi.org/10.1257/jel.20191597>
- IPCC. (2014). *Climate Change 2014: Synthesis Report. Contribution of Working Groups I, II and III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change* (R. K. Pachauri & L. A. Meyer, Eds.).
- King, G., Keohane, R. O., & Verba, S. (1994). *Designing Social Inquiry: Scientific Inference in Qualitative Research*.
- Kleinberg, J., Ludwig, J., Mullainathan, S., & Obermeyer, Z. (2015). Prediction policy problems. *American Economic Review*, 105(5), 491–495. <https://doi.org/10.1257/aer.p20151023>
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444. <https://doi.org/10.1038/nature14539>
- Levin, K., Cashore, B., Bernstein, S., & Auld, G. (2012). Overcoming the tragedy of super wicked problems: Constraining our future selves to ameliorate global climate change. *Policy Sciences*, 45(2), 123–152. <https://doi.org/10.1007/s11077-012-9151-0>
- Lonergan, E., & Blyth, M. (2020). *Angrynomics*. Agenda Publishing Limited.
- Lowery, C., & Ramachandran, V. (2015). *Unintended Consequences of Anti-Money Laundering Policies for Poor Countries - A CGD Working Group Report*. Center for Global Development.
- Mildenberger, M. (2016). Process-tracing, counterfactual comparisons and causal inference: A new bridge over the qualitative-quantitative divide.

- Mildenberger, M. (2020). *Carbon Captured: How Business and Labor Control Climate Politics*. MIT Press.
- Mullainathan, S., & Spiess, J. (2017). Machine learning: An applied econometric approach. *Journal of Economic Perspectives*, 31(2), 87–106. <https://doi.org/10.1257/jep.31.2.87>
- Obermeyer, Z., & Emanuel, E. J. (2016). Predicting the Future - Big Data, Machine Learning, and Clinical Medicine. *New England Journal of Medicine*, 375(13), 1212–1216. <https://doi.org/10.1056/NEJMp1606181.Predicting>
- Pearl, J. (1995). Causal diagrams for empirical research. *Biometrika*, 82(4), 669–688. <https://doi.org/10.1093/biomet/82.4.669>
- Pearl, J. (2009). *Causality*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511803161>
- Pickering, J., Jotzo, F., & Wood, P. J. (2015). Sharing the Global Climate Finance Effort Fairly with Limited Coordination. *Global Environmental Politics*, 15(4), 39–62.
- Pol, R. F. (2020). Anti-money laundering: The world's least effective policy experiment? Together, we can fix it. *Policy Design and Practice*, 3(1), 73–94. <https://doi.org/10.1080/25741292.2020.1725366>
- Radford, J., & Joseph, K. (2020). Theory In, Theory Out: The Uses of Social Theory in Machine Learning for Social Science. *Frontiers in Big Data*, 3(18), 13. <https://doi.org/10.3389/fdata.2020.00018>
- Reuter, P. (2012). *Draining development? Controlling flows of illicit funds from developing countries*. World Bank Group.
- Rittel, H. W. J., & Webber, M. M. (1973). Dilemmas in a General Theory of Planning. *Policy Sciences*, 4, 155–169.
- Rogelj, J., Shindell, D., Jiang, K., Fifita, S., Forster, P., Ginzburg, V., Handa, C., Kheshgi, H., Kobayashi, S., Kriegler, E., Mundaca, L., Séférian, R., & Vilariño, M. (2018). Mitigation Pathways Compatible with 1.5°C in the Context of Sustainable Development. In V. Masson-Delmotte, P. Zhai, H.-O. Pörtner, D. Roberts, J. Skea, P. Shukla, A. Pirani, W. Moufouma-Okia, C. Péan, R. Pidcock, S. Connors, J. Matthews, Y. Chen, X. Zhou, M.

- Gomis, E. Lonnoy, T. Maycock, M. Tignor, & T. Waterfield (Eds.), *Global Warming of 1.5°C. An IPCC Special Report on the impacts of global warming of 1.5°C above pre-industrial levels and related global greenhouse gas emission pathways, in the context of strengthening the global response to the threat of climate change, sustainable development, and efforts to eradicate poverty* (pp. 93–174). Cambridge University Press.
- Shaxson, N. (2011). *Treasure islands: Uncovering the damage of offshore banking and tax havens*. Palgrave Macmillan.
- Shaxson, N., & Christensen, J. (2013). *The Finance Curse*. Tax Justice Network.
- Storm, H., Baylis, K., & Heckelei, T. (2020). Machine learning in agricultural and applied economics. *European Review of Agricultural Economics*, 47(3), 849–892. <https://doi.org/10.1093/erae/jbz033>
- Tibshirani, R. (1996). Regression Shrinkage and Selection via the Lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, 58(1), 267–288.
- UN General Assembly. (2015). *Transforming our world: The 2030 Agenda for Sustainable Development*.
- Varian, H. R. (2014). Big data: New tricks for econometrics. *Journal of Economic Perspectives*, 28(2), 3–28. <https://doi.org/10.1257/jep.28.2.3>
- Wooldridge, J. M. (2010). *Econometric Analysis of Cross Section and Panel Data*. MIT Press.
- World Bank, Ecofys, & Vivid Economics. (2016). *State and Trends of Carbon Pricing*. <https://doi.org/10.1596/978-1-4648-0268-3>