Space matters: Quantifying ecosystem-mediated externalities

Anouch Missirian

Space matters:
Quantifying ecosystem-mediated externalities

Anouch Missirian

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy
in the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY

2020

ABSTRACT

Space matters: Quantifying ecosystem-mediated externalities

Anouch Missirian

Economic and ecological processes interact with one another over both spatial and temporal dimensions.This dissertation explores four socio-ecological systems where space crucially matters for both economic and ecological outcomes. In the first chapter, a windborne chemical dictates the diffusion in space of a new agricultural technology. The second chapter dissects the notion of landscape complexity to find which of its components matter for the intensity of insect pressure in agriculture, and thus the use of insecticides. In the third chapter, the location of participants in an environmental program seeking to curb deforestation points to additionality problems and anticipates the lack of measurable effects of the program. Knowing where crops are grown and temperatures less well-suited for their thriving is key to identifying in chapter four the effects of weather fluctuations on asylum applications into the European Union. The spatial dimension tends to be hard to apprehend and overlooked, but those four pieces together stress that space matters in the study of sustainable development.

Contents

List of Figures

List of Tables

Acknowledgements

Beyond measure, I am indebted to the Sustainable Development community at Columbia as a whole for their support, intellectual stimulation, and friendship. Their discreet nudges and trusting encouragements probably did as much as the sheer intellectual contributions of the sharp scholars they are to the completion of this dissertation. To all forms of support, I am ineffably grateful.

I am in particular ever grateful to Prof. Wolfram Schlenker for his guidance and his trust, and for having given me in his unique way a leg up into research. To Prof. Douglas Almond, for his sharp counsel, unfailing support, constant encouragements, and good humor. To Prof. Shahid Naeem, for his generous and thoughful advice, the enjoyment of doing research together, my full assimilation into the Naeem-Palmer Lab. To Prof. Jeffrey Shrader, first for the incredibly stimulating experience of being his teaching assistant, then for the innumerable hours spent straightening my thoughts, my sentences, my spirits. To Prof. John C. Mutter for letting me know about, letting me in, letting me stay in (and eventually letting me graduate from?) this fantastic Sustainable Development program. To Profs. Geoffrey Heal, Ruth DeFries, Scott Barrett, Rodrigo Soares, and Maria Uriarte, for their liberal dispensation of excellent advice, for the door that I always knew to be open, and for bolstering both my intellectual development and my morale. To Sustainable Development students, old and new. In particular to Eyal beyond measure, for his mentorship. To Anthony, Kimberly, Eugénie, Sol, Jesse (the Elder), for their advice and guidance. To Ana for being the best comrade in arms. To Joséphine, Claire, Florian, Jesse (the Younger), Vincent, for their indefectible enthusiasm, support, intellectual challenge, and undeterrable drive to rise my spirits. To all, for their friendship. And to Mona and Tomara, true Jiminy Crickets.

The larger Columbia community is here only artificially separated from the chemically-pure Sustainable Development one, so involved they were in the elaboration of this thesis. My program has long been enriched by rubbing shoulders with these generous allies, and I have tremendously benefited from their long-lasting benevolence and lavish provision of advice and guidance. To them all I am deeply thankful. To the Naeem-Palmer Lab, the DeFries Lab, the Wolfpack, the sustainable development and the just-development colloquia, specifically, for their community, their constructive and merciless criticism, their encouragements, their intellectual enrichments, and their pleasant company. To the rest of the E3B, DEES, Economics departments at Columbia, whose students and faculty provided support and guidance. To Andrea-Cristina Ruiz for her outstanding contributions to the progress of Chapter 3.

Many others made essential contributions to this work of various natures. For their modesty's sake my acknowledgements are terser but my gratitude no less considerable and heartfelt. To INRA's Unité d'Economie Publique in Paris, and UC Berkeley's GPL, I am thankful for hosting me a summer each, and for the ensuing inspiring interactions; for their hospitality, support, patience, and generosity, enabling two of the chapters presented here to make substantial progress enough as to exist. To Prof. Sara Tjossem, first for her invaluable mentorship and patience in my introduction to teaching, then for her friendship, and her insightful and serening counsel. To Charlotte and her father, for their support in times of dire need. To Julia and Elliot for the music. To Anastasia, Anna and Hugh for the yoga. To Pascale for the asymmetric literary exchanges. To the Hungarian Pastry Shop and its staff for tea, sympathy, and capital writing environment. To Alexandra for the rivers and the Armenian folk music. To Franz, Rita, Frédéric, and countless others, for emotional support. To Hadley Wickham, creator of the ggplot2 R package, for the sleek graphs. To my NPR station for inspiration (viz. Chapter 1), evasion, and comfort. To friends and family back home who tolerated and supported me as they could. To friends and family here who also tolerated and supported me as they could.

To Ohannes. It took me a few years too many.

Introduction[1]

Sustainable development is an applied science. As such, it seeks to solve what is perceived as "problems", to inform policymaking and serve as basis for collective action. As pointed by out Boltanski and Thévenot (1991), collective action demands justification.

Underlying any justification is one of six economies of worth ("polities", or worlds[2]) identified by the authors.[3] Together, the six form an independent and spanning set of internally coherent and legitimate economies of worth, guided by a single ordering principle enabling to rank actions, objects, people, according to how they correspond to a notion of the common good. In the domestic polity, is considered great what abides by a hierarchy established by tradition; in the market polity, competition, trade and markets are key; in the industrial polity, judgement of worth is based on process efficiency and competence; in the inspired polity, is great he or she who performs acts of original creation; in the fame polity, proof of worth is obtained by opinion polls; and in the civic polity, collective solidarity as envisioned by Rousseau constitutes the standard against which to measure greatness (Thévenot, 1996). The appeal of the justification paradigm, to me, lies in its ability to analyze conflicting yet rational and legitimate discourses. It is useful here in framing the motivation, ambition, and means of this dissertation.

For as remarked above, sustainable development is an applied discipline, and as such its role, as I perceive it, is to *justify* action, and in the particular context of this dissertation, action as it concerns addressing environmental problems. Each of the six polities mentioned above can be used to support some justification for collective action to address environmental problems. The type of action supported, the substance of the proof required to justify the action, the constituency likely to support that action, however, are likely to differ depending on the economy of worth invoked.

If the present work were to operate under domestic polity, it would certainly gather traditional knowledge, traditions and tales featuring local fauna and flora and transmitted through generations;

[1] Regretted Prof. Laurent Mermet inspired this chapter. I have failed to make any scholarly use of the enlightening insights he shared with us in his 2012 class and probably do no justice to his excellent teaching of the "theory of justification" which made such a profound impression on me. This is but a poor homage, heartfelt nonetheless.

[2] Translations found in Creppell (2007).

[3] The existence of a seventh "Green polity" has been later proposed, but it is beyond aptitude to discuss it here.

under the industrial polity, it would probably be about the engineering of clean, waste-reducing production processes or depollution technologies; under the inspired polity, it would assuredly reflect on the perceived well-being associated with the experience of biodiverse ecosystems; under the fame polity, it would likely dissect the representation of life forms in works of fiction and the reasons for their varied appeal; under the civic polity it would perhaps analyse the movements leading to the creation of national parks, or to unrest against local pollution. But this thesis is rooted in the market polity, in which what is great is what maximizes economic efficiency; it is primarily concerned with coupled human-ecological processes potentially leading to reduced welfare or inefficiencies (as in chapter 1, chapter 2 and chapter 4), and the design of policies to alleviate them (chapter 3). While a limiting perspective, the scope it leaves for the investigation of environmental problems is large.

Restricting the scope further, this collection of work is specifically concerned with land use as a medium through which ecosystems and human societies interact, and the singularly spatial nature of this interaction.

Indeed, land use is transformed in more or less subtle ways by human activities (from forest clear-cutting to the novel use of an agrochemical substance). These transformations translate for the communities that dwell there into habitat loss,which is still considered to be the most important driver of biodiversity erosion (Pimm and Raven, 2000),[4] but in turn the modification of ecological processes can alter economic outcomes and rationales driving human activities (e.g. see Frank (2017)).

The first chapter is set in the United States, where in 2015 a new genetically modified crop was introduced and adopted by farmers at a remarkable speed. While most technologies, and in particular agricultural technologies, tend to be adopted more slowly than pure profit-maximizing considerations would suggest, this new seed trait penetrated the U.S. market rapidly, and utterly defeated its competitors, old and new. I show in this chapter that part of that speediness is attributable to the consequences of the chemical properties of the complementary input (a herbicide) associated with that technology. Dicamba –the herbicide– is indeed highly volatile and prone to drifting across fields, and thus farmers who adopt the technology are liable to imposing negative externalities (dicamba damage) on those of their neighbors that have not adopted the new seed. I find that the roll-out in space of the new technology can at least in part be explained by wind patterns, with farmers in the same wind-corridor as adopters one year more likely to adopt the following year, thus confirming the role of the drift externality in the adoption of the new seeds. I further provide suggestive evidence that this adoption was overall not profitable, as yield improvements did not materialize, nor did cost

[4] Among other outcomes, e.g. food depletion.

reductions, for the crops concerned with the technology (soybean, cotton), whereas reports of damage on non-target crops and wild vegetation abound. Therefore the choice of a few to adopt forced others, through spatial wind-mediated externalities, to follow suit, overall to their detriment.

The second chapter borrows from landscape ecology the following question: does habitat configuration matter for population dynamics? and applies it in the context of the landscapes under agricultural management in the continental United States, and considers specifically the case of agricultural insect pest populations. Indeed insect damage and the use of insecticides in agriculture are privately and collectively costly (yield losses, additional expenses, health issues, etc.), and agricultural landscape simplification (for operational efficiency) and complexification have been successively and concomitantly advocated without consideration or evidence of the effects of either on pest populations – despite complexification being typically justified by its potential for pest pressure reduction. Using spatially-explicit land cover data in the conterminous United States, coauthor Eyal Frank and I build a panel of county-level landscape metrics, quantifying both composition and structure, and provide compelling (though as yet non causal) evidence supporting the importance of landscape composition but also configuration (e.g. how small or large fields are, how distributed in space) in the intensity of pest pressure. If confirmed by further work, this implies a model for collective action such that the benefits of cooperative land management be internalized and thus reaped by farmers and the neighboring communities if they outweigh the private costs.

The third chapter is concerned with the design of efficient policies to reduce deforestation in the developing tropics. The rationale for such policies is well accepted – at least because tropical forests act as a carbon sink and host remarkable biodiversity – but because deforestation typically involves private gain for collectively borne losses (hence a classical externality problem), and the drivers and actors of deforestation are many and complex (hence making classical solutions harder to implement), adequate policy options are still being sought. So-called payments for ecosystem services (PES) aspire tp make conservation actions incentive-compatible, and could in theory be designed to maximize efficiency (as opposed to incurring the loss of the public good, or conserving it at the cost of imposing sub-optimal efforts). This justifies their appeal for public policy and international endeavors in conversation. However due to the aforementioned complexity of deforestation as a phenomenon and the voluntary nature of PES schemes (i.e. mandate of incentive-compatibility), their good design and implementation is complex, but whether that weakness is a fatal one is an open debate. Here I analyze such a scheme in Ecuador, and show that not only was enrolment into the program biased towards marginal land that was at a lower risk of deforestation anyway, but also that compliance fell short. As a result, deforestation dynamics were not altered by the program. Worse still, areas benefiting from the program reacted more strongly (deforested more) when an economic downturn affected the country's finances in general, and the delivery of the payments in particular, suggesting that the PES might have shifted baselines, whether economic, symbolic or otherwise. This implies that targeting and enforcement should be a priority when considering designing a payment for ecosystem service scheme, and that considerations beyond efficacy and efficiency of PES should be further investigated.

The fourth and last chapter points to the importance of agricultural land use as a mediator between current weather fluctuations, possibly future climate change, and human activities: an "agricultural linkage" (Cai et al., 2016) between weather shocks and human migration is established. Indeed over our sample period (2000-2014) temperature and precipitation fluctuations over the maize growing season and area of the origin countries produced a U-shaped migration response (measured as applications for asylum in countries of the European Union), such that deviations from a moderate temperature optimum led to higher volumes of asylum applications, all the more so if the temperature deviation was positive (hotter temperatures). Non-spatialized temperature averages fail to exhibit a comparable response, showing the importance of taking the localized nature of human activities and weather patterns into account. Extrapolations using projections for future temperatures under two climate change scenarios highlight that all else remaining equal (e.g. no adaptation and no tipping-point), moderate to substantial climate change would lead to large increases in such distress-migration, and further add to the economic justification for action to mitigate, and adapt to, climate change.

These four chapters aim at shedding some light on land-use-mediated effects of human activities on ecosystems, and vice versa. Land use is a particularly powerful lens when the spatial aspects of these interactions are of first-order importance. It is not the only available nor relevant angle; more subtle or drastic transformations occur, other activities take place, that do not bear a mark on the land, yet profoundly affect the way ecosystems and human societies interact.

Finally, however insightful the economic framework adopted here might be, it is only relevant in justifying collective action to handle environmental problems within a "market polity". It is deficient or ill-suited to tackle other aspects environmental problems, for instance as they pertain to patrimonial values (domestic) or spiritual considerations (inspired). It can be wielded with modesty and to great effect when trade-offs and efficiency are the stated guiding principles.

The objective of each of the four chapters that follow is to serve precisely that purpose. Each seeks to leverage and advance the science on coupled ecological and economic systems, while informing the trade-offs we operate through collective action.

Chapter 1

Forced technological adoption and spatial externalities

"Life will find a way."
Jurassic Park II, The Lost World

Abstract

Diffusion of new technologies in competitive markets is often thought to be too slow relative to an optimal adoption trajectory due to learning-by-doing, learning-by-using, or network externalities. In this paper, in contrast, I study a phenomenon of hastened technology adoption facilitated by a negative spatial externality imposed by adopters on non-adopters. Focusing on new herbicide-tolerant seeds for soybean and cotton, I show that adoption by U.S. farmers was partly caused by wind carrying the drift-susceptible herbicide across plots. I estimate that being in the same wind corridor as an adopter increased the probability of adopting by about 29%. The externality also led to defensive adaptation: cropland was converted to crops able to withstand the herbicide, suggesting a form of protective land-use change to prevent damage. I then turn to broader consequences of the widespread adoption of the technology, including its overall effect on yields. *A priori*, the effect on yields is ambiguous. I find that overall, yields remained practically unchanged, despite increased crop failure. The rapid diffusion of this new technology and the consequences highlighted here call for the careful consideration of policies to address such inventions and of their accompanying side-effects.

Keywords: technology adoption, spatial externalities, land use, pesticides.
JEL Classification: Q5, Q15, Q16, O33.

Introduction

The adoption of new technologies is often slow. Sometimes inefficiently so, because of frictions such as knowledge spillovers or aversion to change (Sunding and Zilberman, 2001; Jaffe et al., 2005; Cowan and Gunby, 1996). For instance, Griliches (1957) describes the adoption of the newly developed hybrid corn seeds in the 1930s-1950s. Full adoption of these higher-yielding seed varieties within a state took anywhere from six to more than twenty years, in part because of a slow-paced social learning process. More recently, the adoption of the glyphosate-tolerant varieties of soybean took over ten years (see Figure A.1 in appendix). In contrast, dicamba-tolerant (DT) seeds have been adopted extremely rapidly. Within four years of their release in 2015, dicamba-tolerant seeds made up 80% of all U.S. soybean acreage, and more than 70% of U.S. cotton acreage. What might have been different here is the process through which the technology penetrated U.S. fields. I propose that in the case of DT seeds, negative externalities imposed by adopters on non-adopters are the cause for such a rapid adoption.

Specifically, the drift of the complementary herbicide to DT crops, dicamba, on nearby fields could have hastened adoption. DT soybean and cotton have been genetically modified to tolerate the weed-killer dicamba, and this enables farmers to continue using it in DT fields and get rid of weeds while the crops are growing. However other crops, including non-DT soybean and cotton, do not possess that trait and suffer damage if exposed to dicamba. Dicamba is also highly prone to drift. It can get carried over miles after having been sprayed, and re-volatilize even after it has first settled. Thus could dicamba drift, and the compulsion to avoid crop damage, have caused the unusually high speed of adoption of the DT technology?

In this paper, I demonstrate that spatial externalities were a driving force in the fast adoption of the new dicamba-tolerant varieties of soybean and cotton. In particular, I show that a determining factor involved in the adoption dynamics was wind patterns during the soybean growing season – the growing season being the time when soybeans are the most sensitive to damage, dicamba most useful to spray on DT crops, and when climatic conditions are most conducive to volatilization and drift of dicamba. I do so by leveraging spatial variation in wind patterns to identify their effect on the probability of adoption of DT crops at the county level. Said otherwise, does being in the same wind corridor as an adopter make a county more likely to adopt DT crops the following year? Given the potency of the herbicide dicamba, two other questions immediately follow. First, since dicamba damages many crops in addition to soybean and cotton,[1] is DT adoption accompanied by land-use change *away* from vulnerable crops, and *towards* dicamba-tolerant crops? And second, what is the effect of DT adoption on soybean and cotton yields?

I find evidence supporting the proposed adoption mechanism. Adoption was more likely in counties situated within the same wind corridor as counties with high adoption of DT seeds in the previous year. On the other hand, having high adoption in cross-wind counties did not affect the likelihood of adoption. This asymmetry clearly rules out the competing interpretation of the results as a diffusion by word-of-mouth and observation of successful neighbors (which are isotropic, i.e. are the same in

[1] Dicamba kills all dicotyledons, i.e. crops like sunflowers, hemp, (non-DT) soybean, all specialty crops. Monocotyledons, however, are relatively unharmed – crops like maize and wheat, while not immune, are far less susceptible to dicamba.

every direction of space).

Furthermore, examining the effect of DT adoption on other crops, I find that the adoption of DT seeds led to the conversion on average of over 1,800 ha (4,500 acres) per county to soybean or cotton. The decision to change crops altogether, beyond the scope of normal crop rotations, is complicated and costly. The data at hand only concerns the first three years with DT seeds (2015-2017) but the adoption trend since then has not abated, so crop switching has likely also continued. It is all the more likely that in 2017 harmed farmers might still have been in the process of updating their expectations that dicamba damage every year will be the new normal, or hoped for the intervention of their federal or state governments to limit dicamba damage. The continued trend in DT acreage suggests that damage and land-use changes might intensify after my sample ends.

Turning to soybean and cotton yields, a slightly higher crop failure associated with the arrival of the new technology confirms the harm to non-adopters, and suggests that it is indeed a tangible enough threat to motivate adoption. On the other hand, there is no overall effect on yield of cotton and soybeans at the county level. This null effect suggestions suggests that the damage endured by non-adopters was of about the same magnitude on average as the yield benefits enjoyed by adopters.

Finally, I calculate an average application rate for dicamba in-crop use on DT fields, and find it is but a fraction of the label-mandated rate, indicating that only a fraction of adopters use the seeds as prescribed. The fact that some farmers apparently choose to forego herbicide application made possible by DT seeds, even though they did switch to the new technology further goes to show that some of the adopters did not intend to adopt the DT cropping system and purchased the seeds to avoid losses. While in appearance contrary to their best self interest, it is consistent with the existence of strong social norms to address negative externalities (Ostrom, 2000). I take that finding as additional evidence that part of the success of DT seeds in penetrating the soybean and cotton seed markets was earned by forced adoption operating through fear of negative externalities and losses incurred by non-adopters.

Put simply, this paper questions the popular notions that technology adoption always goes slower than socially optimal and that faster adoption is always better.

For indeed while the hypothesized adoption process can seem problematic in itself, its acceleration of the adoption of DT seeds raises further questions that a historical detour will help to highlight. The development of the first herbicide-tolerant (glyphosate-tolerant) crops via genetic engineering in the mid 90s magnified the potency, and the use, of the broad-spectrum herbicide glyphosate. The herbicide-GM combination was obviously perceived as a boon against yield-depressing weeds, as the herbicide could be used during the growing season, killing the weeds, leaving the crops unscathed. With it also came questions. Some were ethical, pertaining to the direct human editing of genomes. Others were ecological, concerning the evolutionary, population genetics, community dynamics consequences of the introduction of these new varieties of plants. Yet others concerned human health, whether by ingestion of GM-based food or by increased environmental exposure to the herbicide.

More than twenty years after the commercial release of the first glyphosate-tolerant crops, and about a tenfold increase in glyphosate use later,[2] substantial concerns over the innocuity of glyphosate have surfaced (e.g., see Dias et al., 2019). Given the experience of U.S. agriculture with glyphosate-

[2]In the United States. In 1996, the estimated use in agriculture was 14,934 tons; in 2016, 131,673 tons (USGS).

tolerant crops, and the biochemical properties of dicamba,[3] this rapid adoption calls for an early examination of the DT seed-dicamba cropping system for potential effects on health.

Moreover, the realization that the widespread use of glyphosate as part of a GM-glyphosate cropping system had led to the prompt and widespread emergence of glyphosate-resistant weeds following a typical Red Queen Race dynamics[4] (Van Valen, 1973; Powles, 2008) is disturbing in light of the speed at which DT seeds are being adopted and dicamba use is rising. Thus the very pattern of adoption, enhanced, in speed and possibly in extent, by the side-effects of the technology may be problematic: fast and widespread enough for resistance to evolve, fast and widespread enough for significant and irreversible damage to be incurred by farmers, by the neighboring populations, and by the surrounding ecosystems, and too fast for policymaking to take place (evaluation, regulation).

This paper adds to the literature on technological change, by proposing a new mechanism for technology adoption, and one that, contrary to most of the literature so far, explains why adoption of a new technology can be remarkably fast, rather than slow. In their study on the adoption of tractors in U.S. agriculture, Manuelli and Seshadri (2014) show that the neoclassical model, adequately specified, accounts for the relatively slow adoption of tractors. The 10-90 lag[5] for tractors was about 29 years, compared to 4 to 12 years for hybrid corn (Griliches, 1957; Manuelli and Seshadri, 2014), ten years for glyphosate-tolerant soybeans and 18 years for cotton (Figure A.1). In their case, adequate specification means accounting for the continual quality improvement of the new technology, and the cost of operating the old (draft horses). While their research cautions against the overuse of frictions to explain slow adoption, frictions do matter in other contexts, as exemplified in Griliches (1957) where "passive social learning"[6] explains the non-immediate diffusion of hybrid seeds, or more recently in BenYishay and Mobarak (2019) who show the existence of learning frictions by incentivizing them away in an RCT, leading to faster adoption of the technology. Guiteras et al. (2019), on the other hand, show in another experimental setting that strategic interactions between adopters and non-adopters can play an important role in fostering technology adoption. In a similar vein, I show that protection against the negative spatial externalities generated by a new technology *by in turn adopting the new technology* can drastically increase the speed of adoption of the said technology.

This paper further relates to Heal et al. (2004) in that it exemplifies the authors' claim that crop choices are interdependent *via* spatial externalities imposed by farmers on one another. Wechsler et al. (2018) offer a dynamic perspective in interdependent production choices in agriculture, which may with time become even more relevant to the DT seed case: they show that a farmer's adoption of glyphosate-tolerant crops *ultimately* affects her neighbors' weed-control efficiency by accelerating the emergence of glyphosate-resistant weeds, whereas in the present case a farmer's adoption of DT seeds *contemporaneously* affects her neighbors' ability to grow anything else profitably. In the future, with

[3] Suspected teratogenic (induces malformations) and known irritant. See Bunch et al. (2012).

[4] The Red Queen hypothesis in Ecology states that "a set of interacting species reaches an evolutionary equilibrium at which all their rates of coevolution exactly balance each other" (Rosenzweig et al., 1987); said otherwise, species (e.g. a predator and its prey) keep developing mutually-counteracting strategies (via evolution and natural selection), e.g. for attack and escape, such that the relationship between the species remains constant in the long run despite an ever-changing set of traits.

[5] Time elapsed between the moments when the penetration rate of the technology 10% and 90%.

[6] BenYishay and Mobarak (2019).

the emergence of dicamba-resistant weeds,[7] the same dynamic case of spatial externalities could apply. The present inquiry therefore lies at the intersection of the study of the mechanisms of technology adoption and of strategic complementarity (*sensu* Bulow et al., 1985).

This investigation of spatial environmental externalities is akin in spirit to prior work on mobile natural resources. While the questions asked and methods used differ, the underlying phenomena at play are similar: the profit-maximizing activity of one individual or firm deteriorates the revenues of others by reducing or destroying their output. Costello and Polasky (2008) for instance considered the consequences of fish stock mobility on optimal harvesting rules and concluded that sustainable harvest could only be obtained under drastically more stringent policies than if the fish did not move. Similarly, seed dispersal across a landscape of private landowners complicates the eradication of an invasive weed (Costello et al., 2017). Closer to the topic of this paper, but still applying a theoretical approach, Munro (2008) wonders at the possible coexistence of GM and non-GM crops, when cross-pollination of GM onto non-GM (genetic contamination, not examined here) generates a negative externality and existential risk for non-GM growers. He concludes that absent legislation imposing draconian rules for the spatial arrangement of GM and non-GM crops, the non-GM crops are bound to disappear because of the spatial externality imposed by the dispersal of, and contamination with, GM material. Here, I show empirically that such patterns of spatial externalities can hasten technology adoption and displace activities across sectors (namely towards soybean and cotton), and do so by exploiting the chemical properties of the complementary input to the technology (the volatility of dicamba) and wind patterns.[8]

The rest of the paper is organized as follows: Section 1.1 provides a brief description of the dicamba-tolerant technology and to what extent it differs from, or is similar to, other GM crops, Section 1.2 describes the data sources, Section 1.3 details the empirical strategy, the results are presented and analyzed in Section 1.4, and finally, Section 1.5 outlines areas for future research and Section 1.6 concludes.

[7] Already reported in 2019: NPR, "As weeds outsmart the latest weedkillers, farmers run out of easy options," 11/04/2019.

[8] In that respect, I follow Schlenker and Walker (2016) and Deryugina et al. (2019) in using wind speed and direction for my identification strategy. The specifics, and the implications for identification, however, differ (see section 1.3).

1.1 Background: Dicamba-tolerant seeds in the landscape of U.S. genetically modified crops

Dicamba-tolerant (DT) seeds are the continuation of the technical and agronomic paradigm of the Green Revolution and of its later avatar, genetically modified crops, that gradually conquered the U.S. agricultural landscape, in both the figurative and the literal sense (Figure A.1), since 1996. 1996 was the year that saw the first commercial release of genetically modified (GM) seeds including glyphosate-tolerant soybean, corn, and cotton. Few other herbicide tolerances have since then been engineered in field crops[9] and in that sense the release of DT seeds in 2015 was a landmark in the history of GM crops. Tolerance to glyphosate was always the most salient trait until 2015. The 20-year experience of glyphosate-tolerant crops has produced a set of difficult ethical, legal, ecological, and epidemiological questions that can be directly transposed to dicamba-tolerant crops. But in addition to these, the very physicochemical properties of dicamba, and in particular its propensity to drift, pose other.

Dicamba-tolerant cotton and soybean seeds had been in development for about ten years when the USDA approved their commercial release in January of 2015. Farmers started buying and planting them during the 2015 growing season. Quickly they were used over the majority of the U.S. soybean and cotton acreage: DT soybean represented only 2.4% of all soybean planted (82.7 million acres) in 2016, but 79.1% in 2019 (totalling 60 million acres), whereas DT cotton covered 0.5 million acres in 2015 (6.3% of U.S. upland cotton area) and about 7 million acres (73%) in 2018.

Problems were quick to emerge. As early as 2015, complaints about damage caused by off-target movement of dicamba applied on dicamba-tolerant crops surfaced. Tensions between neighbors mounted season after season, and things came to a head in October of 2017 when a dispute between Arkansan cotton farmers over DT seeds and the use of dicamba led to the murder of one of them.[10]

Yet GM crops had been part and parcel of the U.S. agricultural toolset for more than twenty years when DT seeds came about. In 2015 they had conquered an overwhelming majority of the acreage for soybean (herbicide-tolerant, HT), corn (HT, Bt[11]), and cotton (HT, Bt). Leaving aside insecticide-producing GM crops for now, the revolution operated by herbicide-tolerant varieties was that the chemical battle against weeds did not have to give way to other means of weed management (e.g. mechanical weeding) once the seeds were planted: herbicides (specifically, glyphosate) could, and were meant to, be sprayed over-the-top in HT crops *during the growing season*. All plants without a tolerance trait to the herbicide would die, and HT plants would be the only ones left standing; the benefit was such that as of 2015, 94% of all soybean cultivated in the US is HT (hence glyphosate-tolerant). DT seeds function just the same: DT cotton and soybean have been genetically modified to tolerate the broad-spectrum herbicide dicamba. Dicamba kills most weeds. All plants are somewhat sensitive, but broadleaf weeds (dicotyledons) are the main target, grasses (monocotyledons) being less sensitive to dicamba. Therefore farmers planting DT crops can continue using dicamba into the growing season to eliminate weeds while their soybean or cotton plants keep thriving.

[9] Exactly two: tolerance to glufosinate and to 2,4-D.

[10] *Arkansas Times*, "Farmer vs. farmer. The fight over the herbicide dicamba has cost one man his life and turned neighbor against neighbor in East Arkansas," 10/08/2017.

[11] Engineered to produce an insecticide protein. It is metonymically called Bt, for *Bacillus thuringiensis*, the bacterium in which the protein was first isolated.

What makes DT seeds stand out though, is that their companion herbicide, dicamba, is highly prone to drifting. This is due to the physicochemical properties of the dicamba compound itself, and of the herbicide products containing it. Drift is problematic given how deadly to broadleaf plants dicamba is. Drift comes from two separate phenomena: spray drift (during herbicide application) and post-application volatilization. With dicamba, both are important risks. Dicamba has repeatedly been identified as one of the leading causes of spray drift incidents. But perhaps most significantly, it is known to re-suspend (volatilize) up to several hours after the spray first settled because of its fairly high vapor pressure,[12] and subsequently get transported over long distances. Vapor pressure is an intrinsic property of a chemical compound or mixture that indicates its propension to evaporate (volatilize) at a given temperature. This was not only a known characteristic of dicamba, but a concern explicitly expressed in 2013 during the public comment phase of the preparation of the Environmental Impact Statement that eventually led to the approval by USDA APHIS of the two varieties of dicamba-tolerant seeds:

> "Non-target plant damage associated with herbicide spray drift and volatilization is a major concern for specialty crop growers and processors. [...] Dicamba, because of its potential to drift and volatilize, has proven to be one of America's most dangerous herbicides for non-target plant damage." (online comment APHIS-2013-0043-0030 by Kimberly Iott, Iott Ranch & Orchard)

Worries about maintaining non-GMO (genetic contamination) or organic (pesticide contamination) certifications were also expressed, but concerns about herbicide volatilization and ensuing damage were by far the most numerous and pressing.

The signs of dicamba damage are very well-known to farmers. The substance had been in use for several decades before the introduction of the dicamba-tolerant crops, and its symptoms (cupped leaves, wilting stems) readily identifiable and long-known. Dicamba's propension to drift and volatilize was also very well-known (hence the herbicide was heretoforth used before the growing season, and with extreme care regarding neighboring non-target plants even in pre-growing season applications). Therefore, the first signs of damage encountered by non-DT crop growers were easily interpreted (such marks are characteristic of auxinic herbicides and specifically of dicamba and 2,4-D[13]), all the more so that the rollout of the DT crops had been advertised and the implications in terms of dicamba use well understood (as illustrated in the farmers' testimonies collected for the Environmental Impact Statement, see above, and the concerns voiced by weed scientists at agricultural extensions throughout the relevant region). As the growing seasons passed, dicamba drift became an increasingly present and visible concern in the local and specialized (whether traditional or social) media, thus a farmer had no need to have experienced the damage herself to know that, should her crops be vulnerable, and a neighbor plant DT soybean or cotton, her yields were at risk, all the more so if past incidents or knoweldge of the prevailing winds led her to believe that her fields would be exposed to the drift.[14]

---

[12] 3.38×10^{-5} mm Hg at 25°C. Higher vapor pressure at a given temperature corresponds to more volatile compounds, e.g. compare with the vapor pressure of notoriously volatile ether, 538 mm Hg, and also contrast with that of glyphosate, 2.89×10^{-10} mm Hg.

[13] Some subtle differences may enable experts to tell them apart.

[14] See quote in 1.5. It exemplifies preventive action undertaken contemporaneously to a neighbor adopting DT soybean.

A controversial – and for now speculative – domain of potential dissimilarity between previous GM crops and DT seeds is their role in the development of resistance in pest populations. The blanket use of glyphosate throughout the United States since the release of glyphosate-tolerant crops has led to the development of several problematic populations of glyphosate-resistant weeds, and Bt crops were not spared Bt-resistance either (Bagla, 2010). Now glyphosate resistance traits are present in multiple weed species with populations large enough to wreak havoc in several regions of the U.S. (e.g. kochia in Western Kansas, marestail in Western Tennessee). The reason for this is that the massive use of a biocide substance amounts to an active kindling of what is called by evolutionary ecologists a "Red Queen race" between farmers and weeds (or insects, in the case of Bt crops). Faced with the overwhelming dominance of a single existential threat and therefore evolutionary pressure, a widely used and potent herbicide (e.g. glyphosate), multiple weeds soon develop resistance to it;[15] farmers then switch to another herbicide (e.g. dicamba) that they again use widely and at first very effectively on weeds, before some of them again evolve resistance. This coevolutionary process (arms race) between farmers and weeds is called a Red Queen race because of what the Red Queen tells Alice in Lewis Carroll's *Through the Looking-Glass*: "Now here, you see, it takes all the running you can do, to keep in the same place." Following this logic, the faster and the more widespread the diffusion of a means of pest control, the faster and the more widespread resistant specimens appear. It has been argued by DT seed proponents that weeds would not evolve resistance to dicamba. But that had been said of glyphosate too.

[15] Briefly, the mechanism is as follows: pesticide applications kill all sensitive individuals, thus application after application select (give an evolutionary advantage to) the resistant ones, that get to produce seeds and have offspring, and thus pass on their resistance traits to the ensuing generations until they become established in the population.

1.2 Data

This section describes the data sources used: the land use data with which I monitor the extent and location of soybean and cotton fields, the pesticide (dicamba) use data that serves to track adoption of the new soybean and cotton seed, the wind patterns that enable the mapping of counties according to their wind-relatedness, and the agricultural and weather data that are used to assess agricultural outcomes.

1.2.1 Spatial distribution of land cover

Spatially-explicit land cover at a 30-m resolution, the Cropland Data Layer (CDL), is available for the conterminous United States for the years 2008 through 2018 (Boryan et al., 2011; Han et al., 2012) and is provided by the National Agricultural Statistics Service (NASS) at the U.S. Department of Agriculture (USDA). I use the CDL to compute annual cropland area at the county level, and its crop-specificity enables me to track land cover change at the ($900\,\mathrm{m}^2$) pixel level, in particular conversions to and away from soybean and cotton land covers.

I further exploit the spatially-explicit nature of the CDL for the computation of indices capturing the spatial aggregation or dispersion of soybean fields within a county. These are the Patch Cohesion Index and the Aggregation Index; they are borrowed from landscape ecology, and described in McGarigal et al. (2002).

The CDL is derived from remotely sensed data through a classification process, and words of caution as to their usage apply (Lark et al., 2017). As far as this study is concerned, the thematic classes for which spatially-explicit data is leveraged correspond to common, widespread crops, and are exceedingly well identified by the classification algorithms (user accuracy above 90%[16]). Remaining classification error in producing the CDL is plausibly independent from the phenomena studied here, and can therefore be considered as white noise in the statistical analyses below. Finally, county-level estimates of crop-specific acreage obtained with the CDL track closely those obtained independently by the USDA using surveys. See for instance Figure A.11 in appendix. While that does not guarantee the quality at the pixel level, it supports the relevancy of these data on at least as fine a scale as that of the county.

1.2.2 Pesticide use

Dicamba use is obtained from the US Geological Survey (USGS). As part of the USGS's National Water-Quality Assessment (NAWQA), the Pesticide National Synthesis Project collects estimates of pesticide use in agriculture for the conterminous United States by compound (i.e. chemical species, for instance dicamba, atrazine, 2,4-D, etc.), in kilograms, at the county level and on a yearly basis since 1992. These are produced by combining proprietary pesticide-by-crop use data collected every year throughout the country with crop acreage (Baker and Stone, 2015; Thelin and Stone, 2013). 525 compounds are assessed. Pesticide use data for compounds other than dicamba, e.g. glyphosate, come from the same database.

[16] Accuracy matrices by year and state can be found here: www.nass.usda.gov.

1.2.3 Wind speed and direction

Exposure to the externality, dicamba drift, is shaped by wind patterns. It does not suffice to be neighbors with an adopter to be considered "treated": the two counties also have to be connected by winds, which is a unique feature of this technology diffusion process.

County connectivity by wind is determined using wind speed and direction data produced by the Earth System Research Laboratory at the National Oceanic and Atmospheric Administration (NOAA). The NCEP NARR data set provides gridded weather,[17] and in particular, U-wind and V-wind data in meters per second, for 1979 through 2019 (Mesinger et al., 2006). The grid is Northern Lambert Conformal Conic which in practice means that the spatial resolution varies: the mesh size is finer towards the Equator, specifically about 33 km at the southern tip of Texas, 50 km at the 49th parallel north border with Canada. The values for U-wind and V-wind in that data set are used to produce a mapping of counties, i.e. associate every county with all of its up-, down-, and cross-wind neighbors.

As illustrated in Figure 1.1 U-wind and V-wind are the components in cartesian space of a wind vector, that is to say, U-wind is the projection of the wind vector on a parallel (a positive U-wind is from the West), and V-wind that on a meridian (a positive V-wind is from the South). Denoting $\vec{W}$ the wind vector, it suffices to sum its $\vec{U} = U\vec{i}$ and $\vec{V} = V\vec{j}$ components to obtain it (U, V in meters per second, $\vec{i}$, $\vec{j}$ unity vectors pointing East and North, respectively), $\sqrt{U^2 + V^2}$ will give the wind speed, and $\operatorname{atan2}(V, U)$ the direction.[18] Summing $\pm\vec{U}$ and $\pm\vec{V}$ enables one to find points up-, down-, and cross-wind from the initial location. It then suffices to associate these locations (identified by their latitude and longitude) to obtain the set of counties located up-, down-, and cross-wind from a focal county; an illustration of the procedure is shown in Figure 1.2 (more detail in A.3.3).

Given that the initial grid is somewhat coarse, this operation potentially misses some up-, down-, and cross-wind neighbors with the original wind data. For robustness checks, I therefore interpolate (linearly) the U- and V-wind components to roughly double the resolution (about 16 km, on a rectangular grid) and thus obtain a richer and more exhaustive set of neighbors for each county of the conterminous United States (compare top bottom panels of Figure A.2).

Wind data provided by NOAA are provided at the monthly and daily frequencies. For the purpose of assigning wind-neighbor relationships between counties, I use the monthly resolution, and specifically the value reported for the month of June. Windborne dicamba particles are not acutely problematic throughout the year, and therefore, wind patterns matter for drift only part of the year. Prior to the introduction of DT seeds, dicamba was sprayed before planting. This "burndown application," was meant to clean the field. The novelty with DT seeds is that they tolerate dicamba use during the growing season, while the crop is growing. June is the month when all soybean is planted (in normal years), plants are vulnerable and damage is likely to cause not only visible signs but yield losses (McCown et al., 2018), and application of dicamba is still recommended by manufacturers on DT soybeans.[19] Monthly data are therefore well-suited for the neighbor assignment procedure described

[17] From their website at www.esrl.noaa.gov/psd. The NCEP are the National Centers for Environmental Prediction, NARR stands for North American Regional Reanalysis.

[18] The result is an oriented angle from the Equator. atan2 is the two-argument arctangent function.

[19] The R1 stage of soybean growth (i.e. beginning of bloom) occurs late June; another key period for soybeans is the R3 stage (formation of pods) which tends to happen late July, but the "application window" on DT soybeans as recommended by the seed manufacturers closes at the beginning of R1. This is also corroborated by the dates picked for

Figure 1.1: Wind data

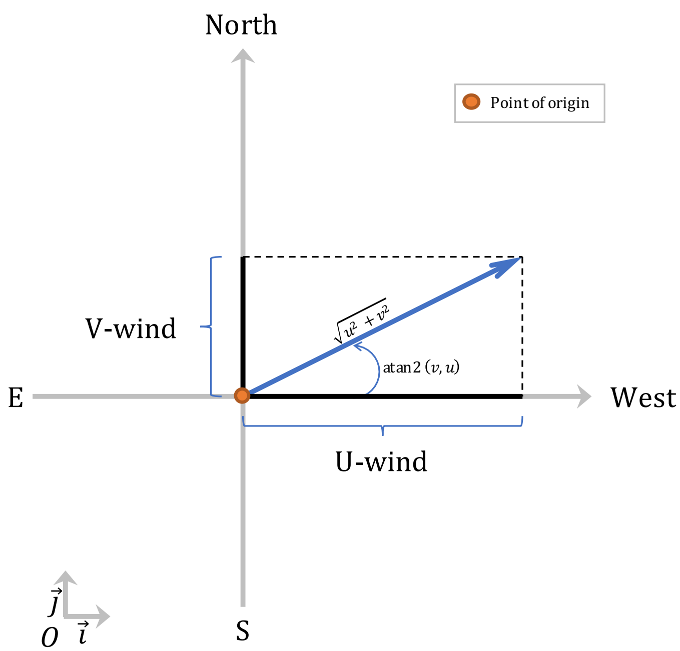

Notes: Schematic presents the NOAA wind data (U-wind, V-wind) and how to derive wind speed and direction.

above, but one may wonder monthly values adequately identify the places where herbicide could drift and settle. I explore the heterogeneity of wind direction at the daily level in the appendix (see A.3.1). I conclude that daily variations are not a threat to the identification strategy adopted in this paper, but that being "within a same wind corridor" might turn out to be more important than being down- as opposed to up-wind from an adopter county.

1.2.4 Agricultural outcomes

I use production and acreage data at the county level compiled by the USDA's National Agricultural Statistics Service (NASS), on the basis of surveys conducted monthly every year from May (small grains such as wheat) or August (row crops such as soybeans, cotton) through November by NASS on representative samples of farms, and extrapolated at the county level for each crop. Production and areas planted and harvested at the county level by crop (soybean, cotton, wheat, corn) are obtained for 2008-2017. These measures are used to construct the following agricultural outcomes at the county level: harvested-to-planted area ratio, yield by area planted, and yield by area harvested.

Since the geographical coverage is not as extensive as that of the CDL, the latter is my preferred source of acreage data in all other cases. While independently produced, the agreement between the two sources at the county-year level is excellent (see for instance the case of soybeans, Figure A.11).

most recent (2018, 2019) local (state) bans on in-crop dicamba use, meant to protect sensitive crops: most of these bans start mid- to late-June, though these are co-determined by political drivers in addition to agronomic considerations. See nonetheless robustness checks for results using July and May instead of June in appendix, tables A.7 and A.8.

Figure 1.2: Illustration of neighbor assignment method: downwind

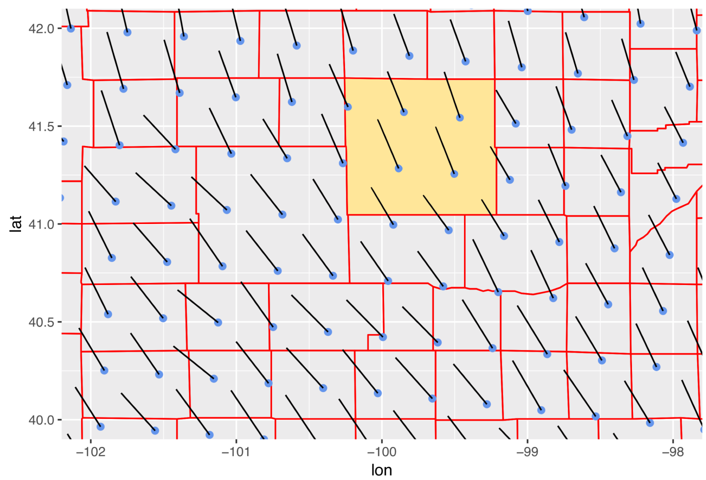

Notes: Figure displays U.S. county boundaries (red) in Nebraska overlaid with transformed wind data (June 2016) in the native resolution. The blue dots correspond to the points of origin (i.e. where U-wind and V-wind are measured); black segments are the vector sum of U- and V-winds; the unadorned tips of the black segments highlight the destination. Thus a given county is assigned as downwind neighbors all the counties other than itself in which land the unadorned tips terminating segments emerging from it. For instance, Custer county, NE (yellow), is assigned three downwind counties (North from it; from West to East, Thomas, Blaine, Loup counties) and is downwind county to three other counties (South from it; from West to East, Dawson, Buffalo, Sherman counties). Upwind counties are obviously obtained taking $-\left(\vec{U}+\vec{V}\right)$, crosswind $-\vec{U}+\vec{V}$ and $\vec{U}-\vec{V}$.

1.2.5 Weather controls

Temperature and precipitation variables at the county level are used as controls in the yield regressions. They are derived from PRISM data (PRISM Climate Group at Oregon State University, 2019) and processed as done in Schlenker and Roberts (2009).[20] These include, at the county and year level: the minimum temperature, the maximum temperature, degree days above 10°C, degree days above 29°C, and total precipitation in millimeters from 1950 to 2017.

[20] http://prism.oregonstate.edu. I am grateful to Prof. Schlenker for providing the county-level data.

1.3 Empirical strategy

I estimate the role of chemical drift in the adoption of dicamba-tolerant seeds by exploiting wind patterns occurring during the growing season preceding adoption, thus abstracting from other modes of technology diffusion. I then estimate the effects of dicamba-tolerant (DT) seed adoption on land use, and specifically quantify what I call protective land-use change, in other words, land-use change away from vulnerable crops into potentially dicamba-tolerant crops (cotton or soybean). Similarly, one could expect consequences of DT seed adoption on crop yields, and I assess that claim. Finally, I turn to consequences at the national scale for the diffusion speed of the technology.

This section details the procedure by which counties are assigned the adopter/non-adopter status in a given year. I later describe the empirical strategy leveraging wind patterns used to identify the effect of windborne externalities and in particular attribute adoption decisions to the exposure to other adopters (via the wind), and likewise attribute land conversions. Finally I turn to the identification of effects on agricultural outcomes and describe how the strategy can be modified to uncover them.

1.3.1 Adopters

Where dicamba-tolerant seeds were sown in a given year is unknown. I therefore need to infer where adoption took place. I do so by observing the use of the complementary input to DT seeds: the herbicide dicamba itself.

Before the release of DT seeds, dicamba was used routinely to "clean" fields before planting any kind of crop. Use is constrained by label prescriptions that bound per-hectare use, from above and from below.[21] It is a broad spectrum herbicide, so any crop would suffer from contact with dicamba after the seeds have germinated and broken ground, *apart* from monocotyledons (e.g. wheat), less sensitive than dicotyledons (i.e. all trees, vegetables, beans, root crops, and some grains like buckwheat and sorghum).[22]

I therefore model annual dicamba use by county (DicambaUse_{it}, in kg) *prior to the introduction of DT seeds* (2008-2014) as proportional to cropland area (Cropland_{it}, in ha) in that county:

$$\text{DicambaUse}_{it} = \beta_0 + \beta_1 \cdot \text{Cropland}_{it} + \varepsilon_{it} \tag{1.1}$$

and thus β_1 is the average dicamba use intensity in kilograms per hectare of cropland. I find $\beta_1 \simeq 0.0132$ kg/ha, and with this model can explain about 17% of the variance in dicamba use (see column (1) in Table 1.1).

Given that (a) some crops (monocots) can tolerate more dicamba than most in a given growing season, and (b) some of them are important users of dicamba (as shown in appendix, Figure A.3), I

[21] E.g. the label for dicamba sold as XTENDIMAX®. The minimum application rate is 11 fl. oz/acre, and the label further reads: "DO NOT broadcast apply *more than 44 ounces per acre for a single application* and DO NOT exceed broadcast applications of *more than 88 ounces per acre within the growing season* when a sequential application is needed for control. Use the higher rate when treating dense vegetation or perennial weeds with established root growth." (emphasis added) The label further specifies the maximum rates for each crop (Table 2, p.4) and commands to "not cultivate within 7 days after applying this product." The label is legally binding for "[i]t is a violation of Federal law to use this product in any manner inconsistent with its labeling. This product can only be used in accordance with the Directions for Use on this label."

[22] Note however that monocots are not *immune* to dicamba; for instance, Bayer is currently working on developing a dicamba-tolerant variety of corn (a monocot).

further refine the model on the margins to allow for an increased use in some crops and land uses, namely wheat and pasture/hay.[23] The coefficients associated to those variables can be interpreted as the additional dicamba per hectare expended on those crops. My preferred model is that shown in column (3) of Table 1.1, with an improved explanatory power reflecting the specificities of dicamba use in wheat fields and in pasture.

	(1)	(2)	(3)	(4)
Cropland (ha)	$\mathbf{0.013}^{***}$	$\mathbf{0.009}^{***}$	$\mathbf{0.009}^{***}$	$\mathbf{0.010}^{***}$
	(0.000)	(0.000)	(0.000)	(0.001)
Wheat (ha)		$\mathbf{0.018}^{***}$	$\mathbf{0.013}^{***}$	$\mathbf{0.012}^{***}$
		(0.001)	(0.001)	(0.001)
Pasture/hay (ha)			$\mathbf{0.003}^{***}$	$\mathbf{0.003}^{***}$
			(0.000)	(0.000)
Corn (ha)				-0.001
				(0.001)
(Intercept)	$\mathbf{248.804}^{***}$	$\mathbf{300.809}^{***}$	$\mathbf{185.130}^{***}$	$\mathbf{183.550}^{***}$
	(13.806)	(13.648)	(14.336)	(14.398)
R^2	0.16	0.20	0.22	0.22
Dep. var. mean	788.82	788.82	788.82	788.82
Num. obs.	18983	18983	18983	18983
RMSE	1463.09	1433.17	1413.00	1412.98

Notes: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. Table shows results for regression of dicamba use (in kg) at the county level against agricultural land uses to obtain in years 2008-2014 to obtain typical per-hectare dicamba use (kg/ha). The models in columns (2)–(4) further add crop acreage for two land uses that are major receivers of dicamba according to USDA statistics, namely wheat (column 2), hay/pasture (column 3), and corn (column 4). Excludes Kansas.

Table 1.1: Dicamba use models

The coefficients obtained from this regression define an average dicamba use per hectare, and I use them to extrapolate out-of-sample "typical" or expected dicamba use in the years 2015-2017 given land use figures obtained from the CDL. This extrapolation serves as a baseline, against which I compare actual dicamba use. The difference between extrapolated and actual values enables me to assign adopter status. Large positive values reveal adoption.

I use $\sigma(\hat{\varepsilon}_{it})$ as my primary threshold ("1sd"), the rationale being that anomalies larger than one standard deviation of the error term of that regression are, given the high excess kurtosis of their distribution,[24] outside the range of typical use and measurement error. Note indeed on Figure 1.3 that most of the mass sits below the continuous red line denoting $\sigma(\hat{\varepsilon}_{it})$ ($\simeq 1412$). Despite its groundedness in the statistical features of the pesticide data and the agronomic characteristics of its use, this threshold may seem somewhat arbitrary still. While the dicamba tolerance trait in DT seeds clearly confers the possibility to double (soybeans) or octuple (cotton) the quantity of dicamba used on these crops in a given growing season (see again crop-specific restrictions on dicamba labels), there is no obvious *a priori* cut-off value for the dicamba anomaly above which a county i can clearly be labelled as adopter in year t. I therefore check the robustness of my results to the use of other values (namely half ("0.5sd")

[23]Including corn area would be suggested by that logic, but it doesn't add much explanatory power (see column (4) of Table 1.1), maybe because corn area is highly correlated with cropland area (corr. = 0.71).

[24]$\text{Kurt}[\hat{\varepsilon}] \simeq 67$, and their distributions pictured year by year in Figure 1.3 are obviously leptokurtic.

and one-and-a-half ("1.5sd") $\sigma(\hat{\varepsilon}_{it})$) and report them in the corresponding appendix sections. It is just equally natural to consider the anomaly-per-hectare a measure of adoption intensity in a given county, but defining thresholds is more problematic. I therefore keep that alternative definition for the appendix, and provide results for two empirically-defined thresholds based on Figure A.4.

Dicamba use anomaly (in kilograms), and intensity of dicamba use anomaly (in kg/ha[25]) are used as alternative (continuous) measures of adoption in sections A.2.5 and A.2.5. However the main analyses presented in the paper, namely the contagion analysis in section 1.3.2 and the examination of land use changes and yields in sections 1.3.3-1.3.4, all rely on the binary variable defined above tracking adopter/non-adopter counties.

Figure 1.3: Distribution of dicamba anomaly

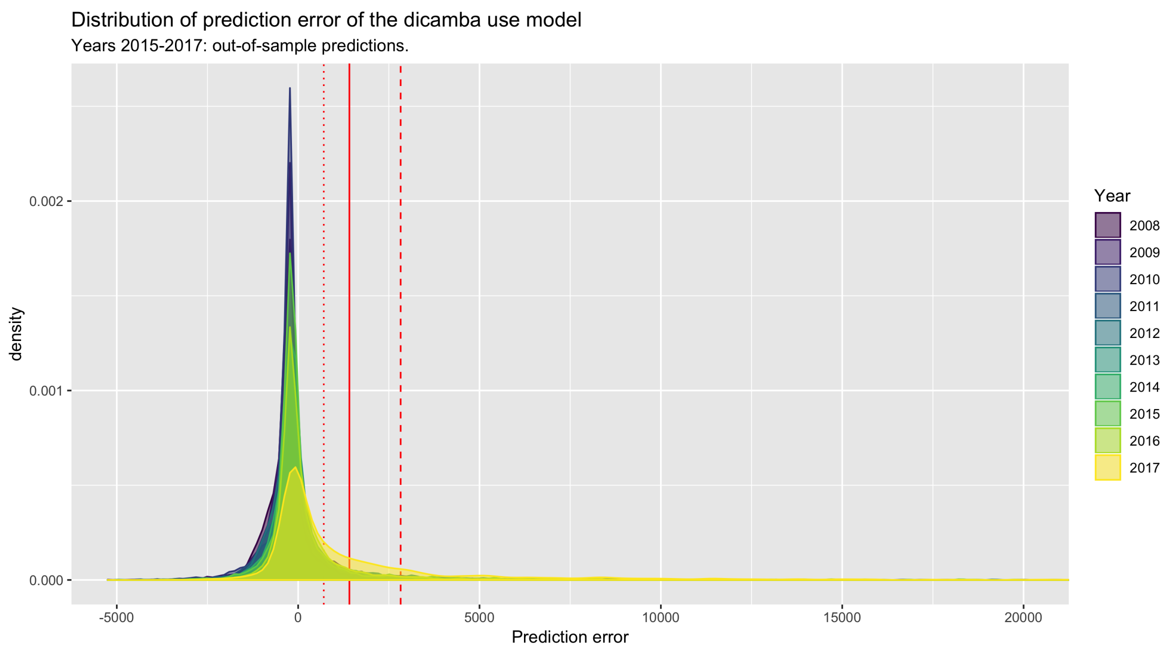

Notes: Figure displays the distribution of dicamba anomaly at the county level (homogeneous to kilograms) for every year in the data *without* surface normalization; years 2008-2014 correspond to in-sample error ($\hat{\varepsilon}_{it}$), errors for years 2015-2017 are out-of-sample errors. Note the symmetry around zero and small standard deviation of the distributions for years 2008-2014 ($\sigma < 1700$) and the positive skew, larger standard deviation of later years ($\sigma_{2016} = 2540$, $\sigma_{2017} = 5762$). The red solid line marks an arbitrary threshold at one standard deviation of the regression residuals ($\hat{\varepsilon}_{it}$ from Equation (1.1) and Table 1.1 column (3)), the red dotted line at half a standard deviation, and the red dashed line two times a standard deviation. The plot was truncated at $x = 20,000$ to facilitate reading (the full support is $[-5,253.1; 72,588.9]$).

1.3.2 Estimation of contagion

Having assigned adopter status for each county and each year (section 1.3.1) and mapped each county to its up-, down-, cross-wind neighbors (section 1.2.3), I can now estimate a wind-biased adoption

[25]Here, county area is used for surface normalization. Indeed, for county-wide consequences what matters is the pseudo-concentration of the substance, not how much more per hectare of *cropland* there is.

model, formulated in its most general form as follows:

$$Pr\left(\text{SelfAdopt}_{it}\right) = f\left(\text{UpwindAdopt}_{i,t-1}, \text{DownwindAdopt}_{i,t-1}, \text{CrosswindAdopt}_{i,t-1}\right) + \varepsilon_{it} \quad (1.2)$$

where $Pr\left(\text{SelfAdopt}_{it}\right)$ is the probability that county i is an adopter in year t (2017). It is modelled as a function of whether any county upwind ($\text{UpwindAdopt}_{i,t-1}$), downwind ($\text{DownwindAdopt}_{i,t-1}$), crosswind ($\text{CrosswindAdopt}_{i,t-1}$) from i was an adopter in year $t-1$ (2016).

As for $f(.)$, I estimate Equation (1.2) fitting a probit, and also provide results for a logit and a linear probability model in the appendix.

The fact that wind patterns are assigned exogenously to any confounding agricultural and socioeconomic features of the counties and their neighbors enables me to clearly identify the endogenous from the exogenous effects (*sensu* Manski (1993)) leading to adoption, and the role of the wind itself in the adoption process. Balance tables in appendix (Tables A.2 and A.3) reveal no substantial differences between treated and control counties.

My hypothesis is that being related by way of wind matters for DT seed adoption; therefore positive coefficients for UpwindAdopt and DownwindAdopt,[26] a non-significant coefficient on CrosswindAdopt, would be in accordance with that hypothesis.

1.3.3 Protective land-use change

What do farmers do when they incur losses (or the possibility thereof)? If they expect the cause to persist, they have two choices: persist (at the risk of taking losses) or adapt. For those farmers that were not soybean or cotton growers, one way to adapt is to become one, and specifically one that uses DT seeds. Switching from a vulnerable crop to a crop where the DT trait is available guarantees that no harm from dicamba drift will be suffered. While harm to specialty crop has been widely reported,[27] some crop changes are more unlikely than others, due to the important fixed costs and changes in farm management practices required by, for instance, converting an orchard into a soybean field (compared to converting a tomato field into a soybean field). The changes are therefore expected to be small especially in the years just following the introduction of DT seeds, for which I have data.

To measure land use changes, I track pixel values from one year to the next in the Cropland Data Layer (see the Data section) to get actual conversion from vulnerable crops to soybean and cotton – rather than picking up, for instance, the expansion of soybean into new (marginal) cropland. In years preceding the release of DT seeds, net conversions to soybean and cotton (i.e. the surface switching to soybean and cotton, minus that switching away from soybean and cotton) should be close to zero, lest there be a gradual disappearance or conversely hegemony of those crops in the county.

The hypothesized sequence, if such protective land-use changes exist, unfolds as follows: soybean and cotton growers within the (wind) neighborhood of non-DT seed users increase their use of dicamba in a given year, and neighboring non-DT crop farmers incur losses; the following year, the latter switch to a DT crop. Given the county-level resolution of the data at hand, this sequence becomes: large

[26] See section A.3.1 for a detailed discussion of why that is. In a nutshell, wind tends to blow along a certain cardinal direction (e.g. NS), but the way it blows varies (North to South, and vice versa). Therefore calling a county *downwind* (upwind) from another is probably a misnomer, as the wind will blow from (toward) it a large fraction of the time.

[27] E.g. *The Fern* reports on tomato, pecan farmers incurring seizable losses for three years in a row (*The Fern*, 11/13/2018: "Scientists warned this weedkiller would destroy crops. EPA approved it anyway").

positive dicamba use anomaly in the county increases in a given year, and the following year more conversion to DT crops is observed in that county – this corresponds to Equation (1.3).

$$Y_{i,t\to t+1} = \alpha_1 \cdot Treated_{it} \times Post_{it} + \lambda_i + \lambda_t + \varepsilon_{it} \tag{1.3}$$

with $Y_{i,t\to t+1}$ the net conversion between years t and $t+1$ to soybean and cotton from some other land use, $Treated_{it}$ a binary variable taking the value of one if county i in year t is treated, i.e. has an adopter in its wind corridor (based on dicamba use anomaly (kg, kg/ha) compared to model predictions, see 1.3.1). λ_i, λ_t are county and year fixed effects, respectively. Errors ε_{it} are clustered at the state level. Alternative formulations also consider $Treated_{it}$ to correspond directly to the focal county being itself considered an adopter (again, based on deviation from baseline dicamba use, cf. 1.3.1), or as a placebo, $Treated_{it}$ is assigned a value of one if county i has adopters cross-wind but not up- or down-wind.

1.3.4 Yields

The effect on yields of DT seed adoption is *a priori* ambiguous for cotton and soybean. On the one hand, it could lead to improved yields if the new seeds were performing better than the old, whether because they are intrinsically higher-yielding, or because they enable the farmer to destroy their weed competitors in the field. On the other hand, DT seed adoption could also depress yields if there is a trade-off between traits (i.e. the seeds tolerate the herbicide but are of a lesser yielding variant) offsetting yield gains; if the coexistence of DT and non-DT crops means significant dicamba drift and damage on non-DT plots, possibly up to crop failure; if the protective land-use change described in the previous section means that plots less suitable for soybean and cotton are sown with DT seeds, therefore adding low-yielding acres. It could also be null if the damage doesn't translate into yield losses.[28]

The effect on non-DT crops is expected to be near zero for monocots (wheat, corn), and negative for dicots (beetroots, potatoes, etc.). I only consider the former because of data limitations.

To unpack the effect of DT seeds on yields, I consider two outcome variables: yield (in physical quantity per *planted*[29] acre), and the planted-to-harvested ratio, to account for the fact that farmers observing damage beyond repair do not harvest those fields.

$$Y_{it} = \alpha_1 . Treated_{it} \times Post_{it} + \mathbf{X_{it}}\theta + \lambda_i + \lambda_t + \varepsilon_{it} \tag{1.4}$$

with Y_{it} the yield (bushel/acre for soybeans, lb/acre for cotton) or planted-to-harvested ratio for the focal crop (soybean, cotton, and two non-target crops, wheat and corn), $Treated_{it}$ is a binary variable taking the value of one if county i in year t has an adopter in its wind corridor (see 1.3.1 and above), $Post_{it}$ is a dummy signifying post-2015 (after the commercialization of DT seeds). $\mathbf{X_{it}}$ are weather controls (growing degree-days, killing degree-days, minimum and maximum temperatures, precipitations). λ_i, λ_t are county and year fixed effects, respectively. Errors ε_{it} are clustered at the state level.

[28] Indeed, visible damage (leaf cupping, stunting) can occur without consequences on yield, depending on the stage at which the damage occurs, and its extent (McCown et al., 2018).

[29] USDA NASS reports yield per harvested acre.

1.4 Results

1.4.1 Verifying dicamba use anomaly as an adequate proxy for DT seed adoption

The lack of data on where and when dicamba-tolerant seeds were actually planted made the use of proxies for DT seed adoption necessary. And while their construction relies on a straightforward logic, it is important to verify that those proxies are reasonable representations of the phenomenon.

At the county level, a good proxy for DT seed adoption should correspond to counties where soybean and/or cotton are grown (refer to map in Figure A.6), and preferably where they are the dominant crop (it is indeed unlikely that a small soybean/cotton area adopting DT seeds would produce a change large enough in dicamba use to pass the threshold I defined). A good proxy for DT adoption should also colocate with places reporting dicamba injury.[30] Finally, a good proxy should correspond to places where, DT seeds having become a divisive and fractious subject, social media activity around dicamba and DT seeds is high.

Data for the last two points are only available at the state level. Dicamba injury and complaints were compiled in 2017 and 2018 by the agricultural extension at University of Missouri; these data can only be taken as a lower bound estimate of actual damage, and the difficulty in obtaining them means that the sample is potentially selected.[31] Location data for tweets could only be obtained from user profiles (i.e. tweets couldn't be geolocated themselves) with substantial heterogeneity in precision and formatting, hence only the stated state could be identified without losing most of the database. As visible on Figure A.13, the affected region corresponds roughly, but the limitations of the corroborating data sets are such that they can only be considered as reassuring evidence rather than definitive validation.

Although soybean and cotton acreage did not enter into the determination of the adoption metric, adopter counties fall overwhelmingly in the counties with the largest soybean acreage in the United States – 47% of the 769 adopter counties (in 2017) are in the top quintile[32] and 71% in the top two quintiles for soybean acreage, and the majority of the rest falls in the top tercile (see Figure A.6) or even top quintile for cotton. Only 8 counties fail that test in that the CDL acreage in cotton and soybean at baseline (2008-2014) is zero and remains minimal in 2017.

1.4.2 Forced adoption: Evidence of wind-biased adoption

The wind-biased model of DT seed adoption is estimated using a probit; the coefficients are reported in Table 1.2 and the marginal effects of the preferred specification (column (3) of Table 1.2) are plotted on Figure 1.4. The figure also reports for comparison purposes the coefficients on a linear probability

[30]Note however that the relationship between adoption and soybean injury is complex (not even mentioning reporting and other data issues): when adoption is low, damage is low, and increases as adoption increases, until it reaches a point where there are enough adopters for very few non-DT fields to remain, and for damage on soybean to start decreasing for want of vulnerable soybean plots until full adoption.

[31]There is no centralized reporting system, not all state plant boards, not all agricultural extensions record nor report those estimates. Being based in Missouri, the team was able to conduct field visits and obtain damage data through more channels than for other states. Therefore all things otherwise equal, soybean injury in Missouri should be reported as larger just because of the bias in reporting.

[32]A quintile has 522 counties for soybean, 228 for cotton.

model using the same predictors (top, red),[33] a logit (beneath it, grey), and the marginal effects of the same probit model but with different definitions of adoption (specifically, different threshold-based rules).

	(1)	(2)	(3)
Upwind, but not downwind 2016 adopter	$\mathbf{0.73}^{**}$		$\mathbf{0.73}^{**}$
	(0.23)		(0.23)
Downwind, but not upwind 2016 adopter	$\mathbf{0.40}^{*}$		$\mathbf{0.40}^{*}$
	(0.18)		(0.18)
Crosswind 2016 adopter		0.12	0.13
		(0.34)	(0.34)
(Intercept)	$\mathbf{-0.56}^{***}$	$\mathbf{-0.55}^{***}$	$\mathbf{-0.56}^{***}$
	(0.03)	(0.03)	(0.03)
AIC	3134.69	3147.52	3136.53
BIC	3152.28	3159.24	3159.98
Log Likelihood	-1564.34	-1571.76	-1564.26
Deviance	3128.69	3143.52	3128.53
Num. obs.	2600	2600	2600

Notes: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. Table shows regression results (coefficients and standard errors) for probit models regressing adopter status in 2017 on adopter status in 2016 of up-, down-, cross-wind neighbors. Adoption is defined using the preferred dicamba use model and an absolute threshold on out-of-sample prediction error. All regressions exclude Kansas.

Table 1.2: Probability of adoption

Having an adopter county upwind or downwind increases likelihood of adoption, by about 28% and 15%, respectively, and these coefficients (Table 1.2) and marginal effects are statistically significant at the 5% level, and robust to the addition of the crosswind variable. Having an adopter county crosswind, however, doesn't affect the probability of adoption, and this is consistently observed across all models and specifications. While the point estimates are different for the up- and downwind coefficients, their confidence intervals overlap, and the estimates are not statistically significantly different from one another (Wald chi-squared test: $\chi^2 = 1.3$, df=1, $p = 0.25$). The results are unchanged (significance, sign, magnitude) if one uses a logit or a linear probability model instead of a probit (see Figure 1.4 with the preferred threshold and Figure A.5 for the full set of thresholds using a logit model).

Overall, this means that being connected by winds to a DT seed-adopter neighbor increases by 15 to 28%[34] the likelihood that DT seeds will be adopted the following year. The fact that crosswind neighborhood relationships should not matter when up- and down-wind do suggests that this is not merely a matter of diffusion by word-of-mouth and observation among neighbors, but that there is *something* in the wind – and in the present case, the something is herbicide particles.

Finally, while the comparable effects of adopters up- and down-wind could intuitively be surprising (the county should receive pesticide particles from its upwind neighbors, not from its downwind neighbors), it is easily explained. As discussed in depth in the appendix (A.3.1), daily wind patterns tend to alternate between one direction and its cardinal opposite. Therefore, the "downwind" direction given by the monthly data is actually, in the overwhelming majority of cases, the minor upwind direction

[33] Fitted values lie between 0.28 and 0.57, hence are all in $[0, 1]$.

[34] Taken together, by 29%.

Figure 1.4: Preferred specification: marginal effects

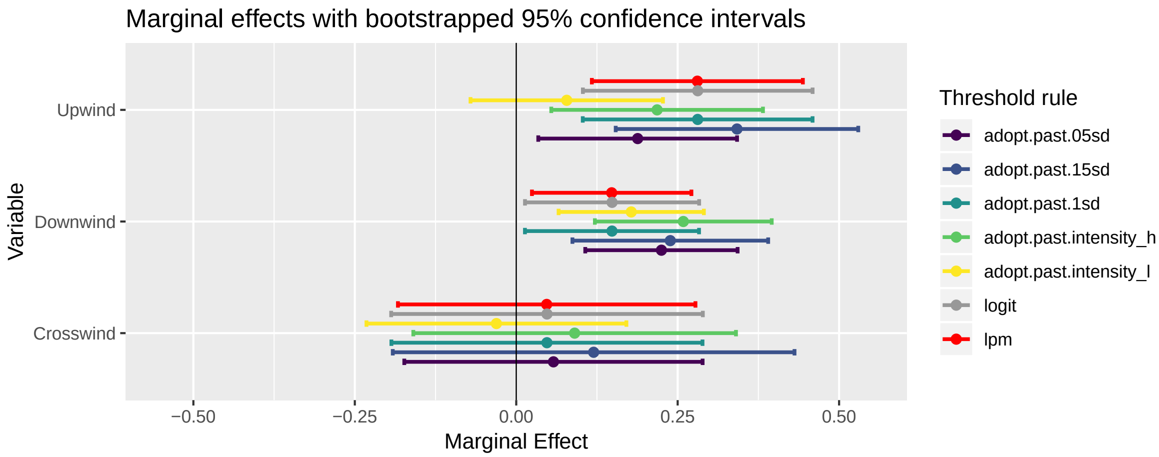

Notes: Graph displays marginal effects corresponding to the probit coefficients in column (3) of Table 1.2, i.e. the preferred specification of the adoption model. In red and in grey, the marginal effects obtained with a linear probability model ("lpm") and a logit model ("logit"), respectively, using the preferred threshold rule. The 95% confidence intervals for the logit and probit coefficients are calculated using bootstrapped standard errors.

experienced at the daily level at that location – in other words, the distribution of wind direction in a given location can be approximated with a bimodal distribution, the monthly data gives the major mode, and I find the minor mode to be about 180° from it (see A.3.1).

1.4.3 Cross-sectoral spillovers: Protective land-use change

The coefficients obtained by the estimation of Equation (1.3) are reported in Table 1.3.

	Hectares	% Cropl.	Hectares	% Cropl.	Hectares	% Cropl.
Treated (Upwind adopt.) x Post	**1852.15*****	**2.00e-04****				
	(274.18)	(0.00)				
Placebo (Crosswind adopt.) x Post			1139.32	−4.44e-05		
			(751.28)	(0.00)		
Treated (Self adopt.) x Post					**529.78****	2.99e-06
					(196.68)	(0.00)
R^2	0.00	0.00	0.00	0.00	0.00	0.00
Dep. var. mean	462.17	1.03e-04	462.17	1.03e-04	462.17	1.03e-04
Num. obs.	23473	23465	23473	23465	23473	23465

Notes: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. Table shows regression results for net land-use change soybean and cotton area (in hectares and as a percentage of baseline (2008-2014) cropland area).

Table 1.3: Protective land-use change

Both treatment variables (lines 1 and 3 in Table 1.3) show a large and positive effect of DT seed adoption on land conversion expressed in hectares of vulnerable uses to soybean and cotton, while the placebo (line 2) does not. The magnitude of the effect corresponds to the size of a few fields (per county): 1,852 ha (4,576 A) is approximately 116 (small) fields,[35] while 529 ha (1,307 A) corresponds to about 33 such fields.

The fact that on the other hand, the specifications using area converted as a proportion of total cropland as their dependent variable (columns 2, 4, 6) show positive coefficients that are less robust (significant on line 1 with treatment defined by presence of adopter in the wind, non-significant on line 3 with treatment defined by being an adopter county oneself) says probably more about the distribution and availability of vulnerable crops as a function of county size.

1.4.4 In-sector agricultural outcomes are mixed, consistent with damage for non-adopters and benefits for adopters

Harvested-to-planted area

The results of the estimation of Equation (1.4) with the harvested-to-planted ratio as the dependent variable are reported in Table 1.4. Focusing on the Cotton and Soybean columns, the effects on the harvested-to-planted area ratios for the crops where DT seeds are available seem overall negative (i.e., more crop failure), but are imprecise; in addition, situations in which the adoption of dicamba-tolerant seeds should not have affected crop failure (wheat; cross-wind placebo) produce relationships that are statistically significant but not meaningful.

	Soybean			Cotton			Wheat			Corn		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Treated x Post	-0.08			-1.71			$\mathbf{3.38}^{*}$			-0.15		
	(0.51)			(1.91)			(1.71)			(0.90)		
Placebo x Post		-1.70			$-\mathbf{9.17}^{*}$			-4.85			1.94	
		(1.39)			(4.23)			(4.53)			(2.54)	
Adopted x Post			$-\mathbf{0.94}^{***}$			-1.44			0.75			-0.07
			(0.25)			(1.46)			(1.00)			(0.48)
R^2	0.15	0.15	0.15	0.30	0.31	0.30	0.07	0.07	0.07	0.10	0.10	0.10
Dep. var. mean	97.98	97.98	97.98	92.58	92.58	92.58	81.17	81.17	81.17	86.63	86.63	86.63
Num. obs.	11903	11903	11903	2771	2771	2771	9380	9380	9380	13478	13478	13478

Notes: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. Table shows regression results for fixed-effect models regressing production outcomes on dicamba use anomaly at the county level, for four major crops, for two of which DT technology is available (cotton, soybean), the other two (wheat, corn) are relatively tolerant to dicamba, owing to their being monocotyledones. Areas are expressed in acres, production in bushels (soybean, wheat, corn) or pounds (cotton), and yields are accordingly in bushels or pounds per acre. Includes weather controls.

Table 1.4: Harvested-to-planted models

[35] Quarter-quarter section (16 ha). The quarter-quarter section is the smallest subdivision of the U.S. rectangular survey system (White, 1983, p. 90), and a typical field delineated according to the rectangular system covers either a quarter-quarter section or a quarter section (64 ha).

Soybean and cotton yields

Observing the results of the estimation of Equation (1.4) with yields (here, production per acre *planted*) as the dependent variable in Table 1.5, there seems to be no net effect of adoption, and the additional amount of dicamba used, whether in-county or in-corridor, on crop yields.[36]

	Soybean			Cotton			Wheat			Corn		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Treated x Post	-0.18			0.09			-0.05			-4.04		
	(0.79)			(0.06)			(1.41)			(2.47)		
Placebo x Post		-1.36			-0.03			$\mathbf{-7.57}^{*}$			6.82	
		(2.15)			(0.13)			(3.74)			(6.98)	
Adopted x Post			0.73			-0.09			0.27			$\mathbf{-2.57}^{*}$
			(0.38)			(0.05)			(0.83)			(1.31)
R^2	0.30	0.30	0.30	0.16	0.16	0.16	0.11	0.11	0.11	0.32	0.32	0.33
Dep. var. mean	42.09	42.09	42.09	1.70	1.70	1.70	46.38	46.38	46.38	124.39	124.39	124.39
Num. obs.	11903	11903	11903	2771	2771	2771	9380	9380	9380	13478	13478	13478

Notes: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. Table shows regression results for fixed-effect models regressing production outcomes on dicamba use anomaly at the county level, for four major crops, for two of which DT technology is available (cotton, soybean), the other two (wheat, corn) are relatively tolerant to dicamba, owing to their being monocotyledones. Areas are expressed in acres, production in bushels (soybean, wheat, corn) or pounds (cotton), and yields are accordingly in bushels or pounds per acre. Includes weather controls.

Table 1.5: Yield models

Put together with the faint negative effect on cotton and soybean harvested-to-planted ratios, this suggests that while some soybean and cotton farmers suffer from the introduction of DT seeds in their neighborhood (1.4.4) and may in addition see effects on their yield, overall thanks to the increased weed control permitted by DT seeds in DT fields, yields were stable at the county level.

Agronomic studies have investigated the action of dicamba on soybean plant development and yield, and in light of that, the results found here on yields are not surprising: the amount of damage to the plant and of yield loss depends on the growth stage at the time of exposure, the number of exposures and the quantity of dicamba. The most sensitive stages are R1 (beginning of bloom) and R3 (beginning of pod formation) as far as effects on yields are concerned; yield losses therefore range from moderate to severe for single exposure events at sublethal doses such as those typically experienced from drift (e.g. 1% to 19%, respectively, as found in McCown et al. (2018) for doses ranging from 1/256th to 1/64th the labeled rate), and have been reported to amount to as much as 50% in case of exposure in both R1 and R3 stages.[37]

[36]Except maybe on corn; an interpretation could be that counties with higher-than-normal dicamba use pre-DT seeds fare better because there are fewer weeds, or because it indicates that famers use more inputs in general, and the negative coefficient could be driven by the larger support (with more extreme positive values).

[37]This latter figure is cited by a crop specialist in DTN, July 12th 2019, "Dicamba Injury Study: Reproductive Stage Soybeans More Sensitive to Dicamba".

1.4.5 Alternative hypotheses

This section briefly overviews alternative hypotheses to the negative spatial externality-mediated forced adoption process presented here.

Land-use change: Commodity prices

The land use changes observed in section 1.4.3 could alternatively be happening because of changes in commodity prices such that after 2015, soybean and cotton would be more profitable to grow. First, for that phenomenon to explain the findings detailed in Table 1.3, they would have to differentially affect counties positioned downwind from adopter counties, which while not impossible, seems peculiar. Second, commodity prices should indeed evolve according to the hypothesis, which they do not (refer to the price time series plotted in the appendix, Figure A.7). Cotton and soybean were certainly not becoming more attractive in 2016-17.

Technical and economic superiority of the new technology

Another explanation for the rapid adoption of dicamba-tolerant seeds could be that they were technically superior to the available technologies, and satisfied a demand or need of the farmers.

DT seeds are at face value similar to the GM seeds that preexisted in soybean and cotton, only with an *additional* feature. In practice, that may not be so straightforward; for instance, new genetic traits may come at a cost for yields (fitness tradeoff or less adapted variety). And indeed, reports on the yields of DT soybeans and cotton are contrasted, with some agricultural extensions reporting no better production per hectare and even slight decreases, and some individual farmers reporting better harvests. The average effects on yields that I find are consistent with that, and see also the discussion on labor, seeds, and chemicals expenses in section 1.4.7.

This hypothesis cannot in any case explain the findings: adoption patterns on the grounds of superiority of the new technology should be anisotropic, hence are incompatible with the wind-biased adoption uncovered here.

1.4.6 Robustness checks

I list and address below possible threats to identification, and provide corroborating evidence of the proposed mechanism.

Sensitivity to sample and treatment definition

Three educated choices have been made in defining the sample and the treatment that could affect the results obtained. Although these choices are based on agronomic and contextual knowledge and are in my view improving the quality of the signal while not adulterating it, it is important to check that they do not bear an undue responsibility in the results.

The first choice concerns sample definition, and specifically the decision to exclude Kansas from the analysis. The reason for that exclusion is detailed in appendix (A.3.2) but in short, has to do with very strong pre-trends in dicamba use at the state level: dicamba use quintupled in the five years preceding the commercial release of DT seeds, for reasons unrelated to the adoption of dicamba-tolerant seeds (since they did not exist then), as illustrated on Figure A.16. No other state displays such a behavior. This was concerning for two reasons. First, the detection of DT seed adopters relies on how closely counties follow a dicamba use model established during the "pre" DT seed period (2008-2014). The elevated use of dicamba and its ramp-up (starting 2010) are likely to compromise the precision and accuracy of the model, and therefore threaten the rest of the analyses of the paper. Second, given its elevated use of dicamba in 2014 and, again, the reliance on positive dicamba anomaly to detect adopters, it is likely that all of Kansas would have been considered adopter *even before 2015*, which is obviously incorrect. As anticipated, including Kansas in the model of dicamba use reduces its explanatory power and changes the point estimates,[38] albeit slightly (Table A.19). Accordingly, the point estimates of the probit adoption model are modified slightly, and the goodness-of-fit is reduced (compare tables A.6 and 1.2); the magnitude, significance, and direction of the effects, however, remain the same. Excluding Kansas was therefore justified, but the results do not hinge upon that decision.

The second choice affects the definition of treatment: the assignment of up-, down-, and cross-wind neighbors based on the NOAA data described in the Data section. A concern here might be that some matches are missed in a systematic way under the native resolution. Indeed while the distance between two points in the data set is typically smaller than the width of a county, not all counties are assigned a wind direction and speed, and therefore are not assigned up-, down-, and cross-wind counties (they could however be up-, down-, or cross-wind from another focal county). The reason why this could be problematic is that this omission is not entirely random: smaller counties are more likely to be missed, and so are counties in the North. The latter is caused by the reliance on a conical projection for the NOAA wind data, such that wind data points are slightly denser in the South than in the North.[39] To address this concern, I interpolate the wind data to obtain a coverage that is uniform in space, and twice as dense as the original (Figure A.2 helps visualize the change); this procedure is liable to introducing measurement error and therefore attenuation bias. Nevertheless, the results of the adoption models are stable, as reported in appendix, Table A.9, which alleviates the concern over spurious relationships found because of potentially poor county matching.

The third choice more directly affects the definition of treatment: the determination of the threshold in dicamba anomaly over which counties are considered as adopters of DT seeds. As explained in section 1.3.1, there is no obvious way to define it. I picked one that was both intuitive and relevant, such that a county is "adopter" if its dicamba use anomaly is above one standard deviation, but verifying that the dynamic is not sensitive to variations around that definition is important to ascertain the validity of the findings. I therefore contemplate a variety of other definitions based on dicamba use anomaly (visualized on figures 1.3 and A.4); as shown already in Figure 1.4 (plotting the marginal effects for various threshold values including the preferred one), it is indeed stable; the full set of corresponding regression tables is reported in the appendix, tables A.10 to A.13.

[38] For instance, Kansas being the first wheat producer of the U.S., its high dicamba use is passed along to a larger coefficient on wheat, but smaller on pasture.

[39] See section 1.2.3 of the Data section; refer to NOAA's page for more information on the Lambert Conformal Format (that in which the wind data is made available) at www.esrl.noaa.gov and to visualize the grid and the distorsions in coverage it engenders across a latitudinal gradient.

Does landscape configuration play a role in adoption?

Since the proposed mechanism for the pattern of adoption identified in section 1.4.2 is through the resuspension and drift of dicamba particles from one soybean (or cotton) field to the next (the former DT, the latter not) the arrangement in space of the fields should matter for adoption at the county level. Indeed if the soybean fields of a given county are all perfectly interspersed with other crops where the DT trait is not available, then the likelihood for damage by drifting dicamba will be minimal; conversely, if they are aggregated, dicamba drift from those among them that are DT can engender more damage among the non-DT. To evaluate this hypothesis, I leverage methods of landscape ecology: McGarigal et al. (2002) and Cushman et al. (2008) describe landscape metrics that characterise the composition (land covers) and structure (how they are arranged in space) of a landscape. I apply them on the CDL at the county level, and extract two metrics describing the dispersion of patches (here, soybean plots) in space: the Patch Cohesion Index and the Aggregation Index. While this approach is typically used to characterize suitable habitat patches for animal or plant species in a landscape, it is are relevant here as those metrics inform the propension of dicamba sprayed on a DT seed-sown field to drift to other soybean (possibly non-DT) fields.

The Patch Cohesion Index and the Aggregation Index increase as soybean plots are clustered together in space;[40] I hypothesize that including them in the adoption model should improve model fit and that their effect on probability of adoption should be positive. The results are presented in Table 1.6.

[40] Specifically, the Aggregation Index (AI $\in [0, 100]$) is "the ratio of the observed number of like adjacencies to the maximum possible number of like adjacencies given the proportion of the landscape comprised of each patch type", and attains 100 if soybean fields are a single compact patch. The Patch Cohesion Index (PCI $\in [0, 1]$) is "proportional to the area-weighted mean perimeter-area ratio divided by the area-weighted mean patch shape index", and the closer to 1, the higher the physical connectedness of soybean fields. The metrics are strongly correlated, and are therefore not used together.

	(1)	(2)	(3)	(4)	(5)	(6)
Patch Cohesion Index	$\mathbf{0.24}^{***}$					
	(0.03)					
Aggregation Index		$\mathbf{0.02}^{***}$	$\mathbf{0.02}^{***}$	$\mathbf{0.02}^{***}$	$\mathbf{0.02}^{***}$	$\mathbf{0.02}^{***}$
		(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Upwind, but not downwind 2016 adopter			$\mathbf{0.75}^{**}$		$\mathbf{0.75}^{**}$	
			(0.25)		(0.25)	
Downwind, but not upwind 2016 adopter			0.33		0.33	
			(0.20)		(0.20)	
Crosswind 2016 adopter				0.24	0.26	
				(0.35)	(0.35)	
Up- or downwind 2016 adopter						$\mathbf{0.82}^{***}$
						(0.13)
(Intercept)	$\mathbf{-2.67}^{***}$	$\mathbf{-2.02}^{***}$	$\mathbf{-2.04}^{***}$	$\mathbf{-2.02}^{***}$	$\mathbf{-2.05}^{***}$	$\mathbf{-2.09}^{***}$
	(0.24)	(0.11)	(0.11)	(0.11)	(0.11)	(0.11)
AIC	2781.86	2671.42	2663.19	2672.96	2664.66	2634.53
BIC	2793.39	2682.99	2686.33	2690.31	2693.59	2651.88
Log Likelihood	-1388.93	-1333.71	-1327.59	-1333.48	-1327.33	-1314.26
Deviance	2777.86	2667.42	2655.19	2666.96	2654.66	2628.53
Num. obs.	2356	2405	2405	2405	2405	2405

Notes: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. Table shows regression results (coefficients and standard errors) for probit models regressing adopter status in 2017 on adopter status in 2016 of up-, down-, cross-wind neighbors. Adoption is defined using the preferred dicamba use model and an absolute threshold on out-of-sample prediction error. All regressions exclude Kansas.

Table 1.6: Probability of adoption – Influence of landscape structure

Both dispersion indices are associated with a higher adoption probability, and improve model fit. In addition, the results obtained previously remain overall stable: having an adopter county crosswind the previous year does not hasten adoption, while having an adopter county upwind does; having an adopter county downwind is less precisely estimated (the coefficient becomes non significant) but the point-estimate remains fairly stable.

These results further illustrate the importance of spatial processes in the adoption of the DT technology, and corroborate the mechanism proposed here of negative spatial externalities imposed by adopters on non-adopters, that are all the more damaging that DT-eligible fields are aggregated in space.

1.4.7 On forced adoption

"Listen, I did not want to do this but I am going to be forced to go dicamba"[41]

If the adoption of DT seeds is indeed forced for some farmers, rather than a deliberate choice of a superior technology, it should be then possible that some of them adopt the seed to protect their

[41] Neighbor to soybean grower Randy Brazel. Reported by NPR (Feb 7th, 2019).

yields but refrain from using the herbicide as specified so as not to cause harm to their neighbors. That would translate into a lower-than-expected increase in dicamba use.

The standard application rate of dicamba as a post-emergence treatment (that available only for DT seeds) is 0.56 kg/ha (as specified on the label and cited in Sciumbato et al. (2004)), and farmers can make two applications during the growing season, hence an additional 1.2 kg/ha sprayed per season on DT soybean and cotton *only*.

	Add. dicamba used (kg)	Area under DT crops (ha)	Avg. application rate (kg/ha)
Method 1	5,596,694	10,521,827	0.53
Method 2	5,163,434		0.49

Notes: Methods 1 and 2 correspond to different ways of calculating the additional dicamba use attributable to DT crops; both amount to attributing dicamba use that deviates from the expected quantities to usage on DT crops. Method 1 takes the sum of all non-negative anomalies (as defined in 1.3.1), and Method 2 sums only the anomalies that are larger than one standard deviation (i.e. those that translate in the assignment of a "adopter" status to the county), both for 2017. The area under DT crops is obtained from figures cited in the press as communicated by Monsanto (specifically, 20 million acres of soybean, 6 million acres of cotton, in 2017).

Table 1.7: In-crop application rate estimates

Table 1.7 shows application rates that are of the right order of magnitude, but still well below the label specification; said otherwise, the increase in dicamba use because of DT seeds is lower than an engineering estimate based on (legally binding) pesticide use restrictions would suggest. Using DT seeds without spraying dicamba over-the-top in the fields in which they are planted defeats the purpose of the technology, and amounts to purposefully and knowingly taking losses (rather, not reaping benefits). While Table 1.7 does not rely on micro data of application rates during the soybean and cotton growing seasons, these highly stylized facts suggest that some of the growers adopted the technology intending *not* to use it, i.e. only to protect themselves from dicamba drift.

Farmers' expenses

Despite the absence of clear effect on yields, the technology could have benefited adopters in reducing expenses; indeed, by simplifying substantially the management of weeds during the growing season, the technology is *a priori* labor-saving. It might also have saved expenses on alternative herbicides, possibly enough to compensate the increased use of dicamba; the *a priori* effect on seed expenses is unclear, but discounts were offered on seed-dicamba bundles in the first years after the technology became available.

However, comparing adopters and non-adopters before and after the introduction of the new technology in two consecutive agricultural census year (2012, 2017) in Table 1.8, the counties with high adoption were at no advantage: their labor expenses (column 3) remained unchanged, and their expenditures for seeds and chemicals (columns 4 and 1, respectively) increased.

Note nonetheless that an important limitation of the data used here is its lack of specificity to soybean and cotton growers, and the findings are therefore to be interpreted with caution.

1.4.8 Limitations and caveats

As far as the costs are concerned, the monetary costs associated with yield losses and crop failures for non-DT plots are not measured; while it seems an important component in the assessment of

	Chemicals	Fertilizers	Labor	Seeds	Land
	(1)	(2)	(3)	(4)	(5)
Adopter X post	**407905***	**−1879248*****	33520	**591425****	1251
	(203235)	(295327)	(368364)	(189621)	(1306)
tMin	1613506	2374571	1840022	1293750	−9238
	(926356)	(1340664)	(1699966)	(861813)	(5909)
tMin2	**−97706***	−74370	−122460	−67986	268
	(40427)	(58532)	(74029)	(37660)	(258)
tMax	−3411177	**15471139*****	−5036363	2195726	−6122
	(1874193)	(2719112)	(3439735)	(1736024)	(11911)
tMax2	22407	**−276035*****	44618	**−59287***	113
	(32607)	(47350)	(59828)	(30229)	(207)
prec	−796	**2105***	−351	**−1288***	−7
	(701)	(1004)	(1280)	(652)	(4)
prec2	−1	−1	−1	0	0
	(1)	(2)	(2)	(1)	(0)
dday10C	**31745****	−24651	33944	2121	21
	(9646)	(13948)	(17687)	(8979)	(61)
dday10C2	−3	6	−0	2	−0
	(2)	(4)	(5)	(2)	(0)
dday29C	−2889	**38532*****	−7174	3202	−4
	(6738)	(9743)	(12278)	(6281)	(43)
dday29C2	**48*****	**−53****	5	**30***	−0
	(13)	(19)	(23)	(12)	(0)
R^2	0	0	0	0	0
Dep. var. mean	5,666,969	8,590,222	2,383,954	6,714,904	128,608.5
Num. obs.	5296	5335	5164	5296	5369

Notes: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. Table shows regression results for fixed-effect models regressing costs for chemicals, fertilizers, labor, seeds (USD), land extent (acres), on adopter status before (2012 census) and after (2017) the commercialization of the dicamba technology. All regressions include weather controls.

Table 1.8: Cost models

DT seeds as a whole, this endeavor is rendered difficult by the dearth of data. Indeed, there is no systematic accounting and monitoring of damages and losses. Efforts have been made, in particular by the agricultural extension at University of Missouri to collect and gather such data across the United States, but in addition to showing the alarming extent of the damage, they have also shown the scantness and partial nature of the data.[42] Attempts at identifying a signal of dicamba damage via remote sensing (not shown) proved fruitless: while the signs of damage such as leaf cupping (soybean plants) are unmistakable and visible from the ground, the remote sensing technology I used (Landsat, 30 m ground resolution) might be too coarse in time and space to capture them (I took GPP, gross primary productivity, as an indicator of good plant health, and computed it at the county and monthly level); and while drones have been able to detect dicamba damage over some fields,[43] it is unclear for

[42]For further information, the reader is referred to Dr. Kevin Bradley's articles on the newsletter of the Integrated Pest & Crop Management department at University of Missouri. See for instance "July 15 Dicamba injury update. Different Year, same questions" (19/07/2018).

[43]E.g. see false color image tweeted by a farmer, showing extensive dicamba damage in his field, reproduced in Appendix Figure A.12.

now that dicamba damage, even when it impacts yields, translates *necessarily* into a markedly different and characteristic spectral signal (e.g. depending on concentrations, growth stage, crop, etc.). It has been shown that dicamba's effect on soybean yields depends on the growth stage at which the soybean plant is exposed to dicamba (McCown et al., 2018), but plant scientists are just beginning to uncover the range of outcomes following an encounter with dicamba.

Further words of caution concern the procedure employed to quantify adoption at the county level. The fact that some farmers can buy DT seeds just to be protected from the drift from their neighbors' fields, without using dicamba in-crop themselves, is a known limitation of the method. Another known limitation is that the relationship between crop acreage and herbicide use is not deterministic, and while the fit obtained here for dicamba is reasonable, the construction the proxy for DT seed adoption relies on it and that necessarily introduces noise.[44]

[44] The case of Kansas provides an extreme illustration of the noise to which those estimations are subject.

1.5 Further implications, future research

This paper but scratches the surface of the transformations brought about by DT seeds. I touch here on a few foreseeable consequences to which historical analogues or anecdotal evidence point, and that in light of this study deserve further research. Only after considering these, in combination with the adoption mechanism highlighted here, will the welfare analysis of DT seeds be possible. But all have a broader relevance beyond the case of DT seeds: market structure consequences are to be anticipated in any case of forced technological adoption, while genetically modified crops in general, and in particular those tolerant to herbicide, bring about questions regarding health and ecological processes that can be addressed thanks to the DT seed adoption setting.

There are mainly four types of consequences one should be concerned with: first, consequences on the Farm structure in the U.S. because of the possible supplement of bankruptcies, second, consequences on the Structure of the soybean and cotton seed sector because of the emergence of a dominant firm, third, consequences on the Health of surrounding populations following the increased release of a volatile known irritant and suspected teratogenic substance, fourth, consequences on the Ecology and evolution of the agroecosystems where DT seeds arrive because of the heightened potential for the evolution of resistant weeds, and of the rippling effects through the ecosystem of off-target movement of dicamba on the vegetation.

Farm structure

First, dicamba damage means financial distress for specialty crop growers, it has been reported, and likely for non-DT soybean and cotton growers as well. And while farmers can for the most part withstand a year with inferior yields, several in a row may prove fatal and lead to bankruptcy, but quantifying the effect of DT seed penetration on farm bankruptcy will require longer hindsight. In the specialty crop sector, bankruptcies and decreased production are expected in the DT seed-exposed counties; bankruptcies are to be expected as well in soybean and cotton, but also farm sales, and consolidation. These hypotheses stem directly from the mechanism uncovered in this paper, and are easily testable (and quantifiable) once the data becomes available.

Structure of the soybean and cotton seed sector

> "I have a neighbor, a friend. He calls me and says, 'I am going to have to go dicamba.' [...] Then I have to get on the phone and call every other neighbor and say, 'Listen, I did not want to do this. But I am going to be forced to go dicamba.' Well, then that forces all those neighbors to call all their neighbors. And eventually what you have is a monopoly."
>
> Randy Brazel, soybean grower at the MO/TN border, to NPR (Feb 7th, 2019)

A second line of inquiry will require a longer view still: what are the consequences of the sudden dominance of a single product made by a single firm in terms of competition on the seed market? While the seed producer has reportedly offered discounts to early adopters, it could very well turn to extracting substantial rents should the dominance situation persist and intensify. The notorious opacity of the seed market[45] (at both the manufacturer and retailer levels) certainly should be no hindrance to that.

To get at the welfare consequences of the emergence of a dominant situation on the seed market, it is useful to think in terms of the model of the dominant firm with competitive fringe laid out in Carlton and Perloff (1994). Since only one firm possesses the DT technology, and that, by 2019, DT seeds made well over 70% of both the soybean seed and the cotton seed markets, that firm detains a dominant position ensured by monopoly over the DT technology and secured by the negative externalities it generates on non-adopters; therefore the model of the dominant firm is appropriate.

The paper highlights two phenomena: the within-sector spillovers that force soybean and cotton growers to adopt DT seeds, and the cross-sectoral spillovers, that force the conversion of land into soybean and cotton. There is therefore on the one hand the constitution of a less competitive sector (considering that the seed market for soybean and cotton was either perfectly competitive or an oligopoly with $n > 2$ firms), and on the other hand the transfer of land from a competitive sector (other seeds) to the newly reshaped sector with a dominant firm, soybean and cotton seeds. The total welfare loss associated with the decreased competition in the soybean and cotton seed sector will therefore be the sum of these two effects.

Health

Studies such as Dias et al. (2019) and Frank (2017) have demonstrated the effect of pesticides applied in the environment (namely, the herbicide glyphosate and insecticides as a "cocktail", respectively) on infant health. Given those precedents, and the known high volatility, and suspected teratogenicity of dicamba, effects on health are possible; they are expected in particular on birth outcomes (malformations, death) and respiratory health.

The sudden increase in dicamba use, along with its wind-biasedness would enable the causal identification of the effects of the substance on health, which is in itself of interest for policymaking purposes in general, and in the assessment of DT seeds in particular. Direct evidence the existence, nature, and magnitude of health effects of pesticides on the health of surrounding populations is lacking, were it not for the two studies cited above. Given that those studies have indeed found deleterious effects on birth outcomes (developmental issues and deaths), and given the disproportionate effects of the fetal environment on later life outcomes (Almond and Currie, 2011), it is important to produce such evidence.

Ecology and evolution of the agroecosystems

The sudden and widespread dominance of a single variety of seeds could have agro-ecological consequences too, with the onset of a Red Queen race (Van Valen, 1973; Rosenzweig et al., 1987) (roughly speaking, an evolutionary arms race) between weeds and farmers. The risk was clearly anticipated in the Environmental Impact Statement prepared to inform the decision to deregulate DT seeds:

> "Increased selection pressure caused by wide-spread adoption of HR [herbicide resistant] crops and reduction in the use of other herbicides and weed management practices,

[45] E.g. see "Silicon Valley Company Aims To Help Farmers Compare Seed Prices," NPR, 23/04/2019.

> resulted in both weed population shifts and growing numbers of HR individuals among some weed populations"
>
> (*Monsanto Petitions (10-188-01p and 12-185-01p) for Determinations of Nonregulated Status for Dicamba-Resistant Soybean and Cotton Varieties*, Final Environmental Impact Statement, December 2014, p. 81)

Said otherwise, the massive use of a single pesticide could lead to the selection of resistance traits in the pest, and the constitution of resistant populations, similar to what happened with glyphosate and glyphosate-resistant weeds (Powles, 2008). The more widespread the use, the sooner the resistance. And while resistance to dicamba has just been discovered in Palmer amaranth,[46] the extent of dicamba resistance, the role of DT seeds in fostering it, and its cost to farmers, remain open questions for future research.

Finally, unintended consequences on wildlife have been reported that are poised to be of both economic and ecological importance. Dicamba mimics a plant growth hormone (auxin), and its mode of action is to grow the plants to exhaustion. This is true not only for non-DT soybean, cotton, and specialty crops, but also for wild vegetation. Its off-target moves have reportedly caused damage on trees[47] and wild flowering plants. While in themselves preoccupying, they could translate into farther-reaching consequences for the species that depend on them (for instance for food), e.g. insect pollinators and birds. This dearth of flower resources has already been blamed for the decline (and subsequent displacement by their beekeepers) of some bee colonies, and could alter the ecosystem services provided by birds (e.g. pest suppression) should their populations decline for want of food. Both bearing on the rural economies of places where DT seeds are widely adopted. Assessing any of these ecosystem-mediated consequences will again require a longer record of the post-DT crop world.

[46] "Palmer amaranth resistance to 2,4-D and dicamba confirmed in Kansas", Kansas State University Agronomy eUpdate, Issue 734, March 1st, 2019, webapp.agron.ksu.edu.

[47] See for instance "A Drifting Weedkiller Puts Prized Trees At Risk" (NPR, 27/09/2018) on damage to protected cypress trees in Tennessee.

1.6 Conclusions

Can technology diffusion go "too fast"?

Taken in positive terms, I show that negative externalities imposed by adopters on non-adopters can indeed hasten technology diffusion. And while in their case the negative externality consisted of shame, was purely a social construct, and led to the adoption of health-improving latrines, in the case I present here, the negative externality consists of windborne particles of herbicide drifting from adopters' fields and damaging non-adopters' crops, leads to financial losses for non-adopters, and spills over to other sectors apparently unrelated to those for which the technology was developed.

From a normative point of view, and in light of the current reassessment of another GM companion herbicide, glyphosate, the analysis that precedes also suggests that in the case at hand, adoption probably went too fast, and certainly occurred before the regulators had a chance to evaluate the economic and public health consequences of the technology, let alone its consequences in terms of ecosystem ecology, conservation (direct and indirect effects on non-target species), and evolutionary ecology (emergence of resistance), which all remain to be evaluated.

The implications of this paper are twofold. The first set of consequences concerns the possibility of "forced technological adoption" because of the existence of negative externalities, the second pertains to broader implications for policy regarding herbicide tolerant GMOs of which DT crops are but an avatar. While other drivers of adoption are not ruled out (and should not be), windborne negative externalities paved the way for the new DT seed technology by deteriorating the bottom line of recalcitrant soybean and cotton growers *to the extent that they had adopters in their wind-vicinity* while leaving the adopters' unscathed. In that way, and because otherwise irrelevant wind patterns are shown here to matter for adoption, the extra help provided by windborne dicamba particles to the diffusion of the DT technology amounts to forced technological adoption. Examples of such strategic interactions leading to technological change (rather, switch) need not be restricted to agriculture: one naturally thinks about the possibility of coexistence of organic and non-organic fields, GM and non-GM crops (Munro, 2008), but a similar pattern of adoption in individual cars in the United States provides the starting point in White (2004), where drivers switch to SUVs and light trucks to be better protected in the event of a crash, especially one with a SUV or a light truck, four times deadlier than with a smaller car, thereby giving way to an "arms race" for larger, safer, yet deadlier cars (Anderson and Auffhammer, 2013). This phenomenon, whether leading to desirable (latrine adoption), undesirable (traffic fatalities), or questionable (dicamba-tolerant seeds) outcomes, seems understudied compared to other phenomena of technology adoption, despite its vast potential for policy intervention, obvious from the three examples cited here.

The second set of implications concerns the regulation of GMOs and in particular of those of the herbicide-tolerant (HT) type. What this study has shown is that DT seeds led to the increased use of a potent broad-spectrum herbicide that (a) *happens* to be very volatile, and (b) *by design* was to be sprayed during the growing season (i.e. in-crop), which led to a variety of problematic outcomes for farmers. While (a) seems to be an unfortunate chemical property of the complementary herbicide, all herbicides have some propensity to drift[48] and (b) is not an accident but a *feature* of HT crops in

[48] Despite not being just as salient, glyphosate drift was also an issue when the first HT crops were introduced in the late 90s: "In the first years after the introduction of Roundup Ready soybean it was not unusual to see 10 to 20 rows of

general. Therefore what has been happening with DT seeds since 2015 is only a mildly more vivid example of the typical consequences to be expected from herbicide-tolerant GMOs. Said otherwise, negative production externalities among farmers are a necessary consequence of HT GMO varieties. Externalities are famously inefficient, and therefore warrant policy intervention to restore efficiency. And that the damage should not be caused by the seed itself but by the herbicide is both factually correct (the herbicide is the proximal cause) and a gross fallacy (the HT seed commands the use of the herbicide during the growing season). Policy tools exist, from taxes to quotas and bans, that could prevent further inefficiencies from occurring, such as the constitution of a monopoly situation in the soybean or cotton seed market, and alleviate the public health burden borne by rural populations in the vicinity of fields planted with herbicide-tolerant crops.

corn damaged, even killed, by glyphosate drift," writes professor of agronomy Bob Hartzler of the Agricultural Extension at Iowa State University. He adds: "However, it was rare to see glyphosate injury across entire fields." ("Thoughts on the Dicamba Dilemma," 13/07/2017)

Appendix A

Appendix to Chapter 1

A.1 Additional figures

Figure A.1: GM variety uptake in soybean, cotton, corn

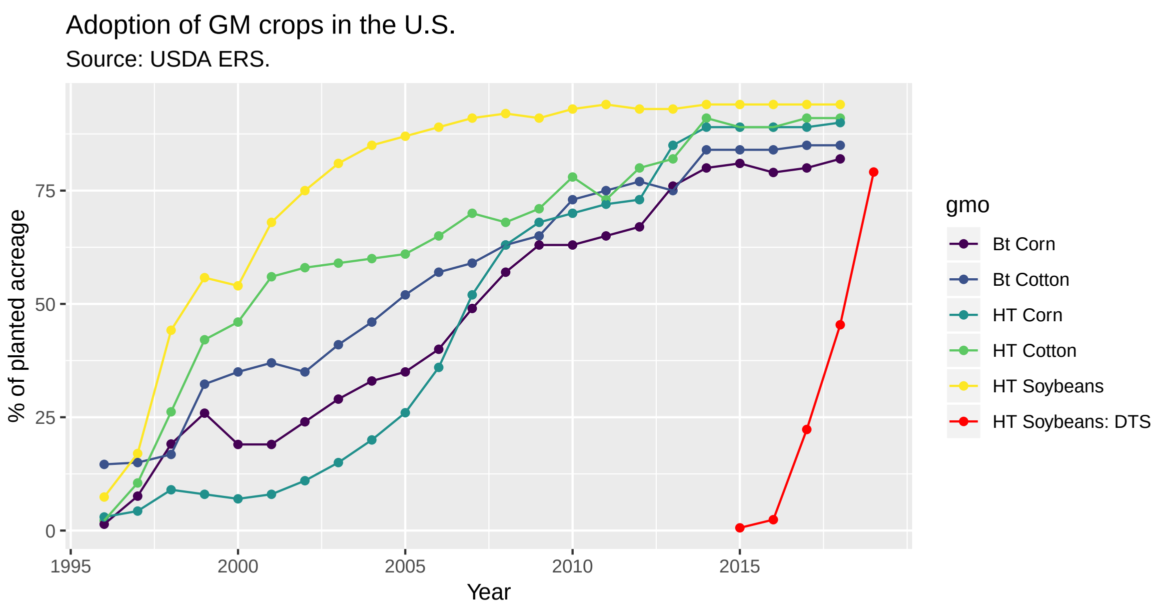

Notes: HT stands for "herbicide tolerant" (and until 2015, it meant glyphosate tolerant), Bt denotes insecticide-producing varieties (Bt stands for *Bacillus thuringiensis*, the bacterium from which the δ-endotoxin genes conferred to Bt crops are taken). Some varieties of corn and cotton are both Bt and HT (data: USDA). In red, HT soybeans that are dicamba-tolerant as a percentage of total U.S. soybean acreage (data: Monsanto/Bayer, USDA).

Figure A.2: Interpolation of the wind speed and direction data: illustration

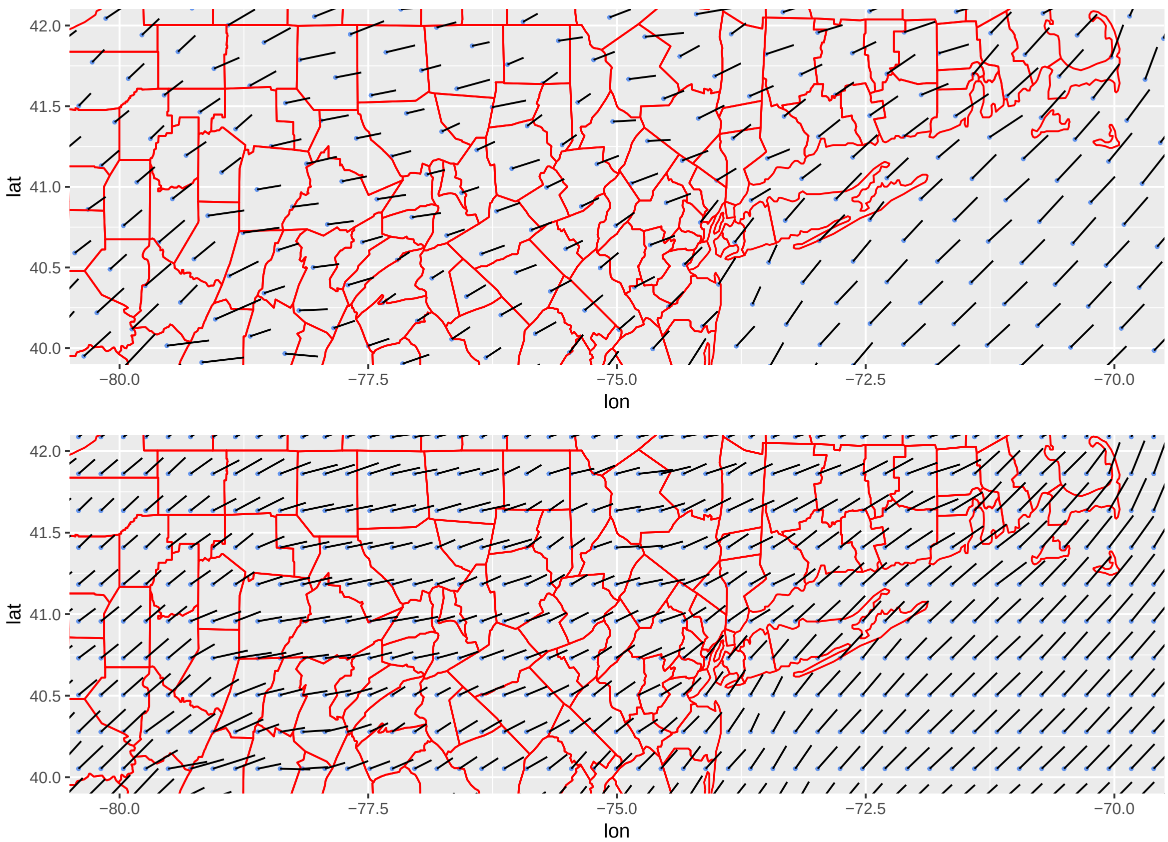

Notes: The top panel shows the original resolution of the NOAA wind data (June 2016), the bottom panel shows the interpolated data (June 2016), with points about two times denser (the distance between two points is divided by $\sim \sqrt{2}$: 32 km in the top panel, 23 km in the bottom panel). The blue tips mark the coordinates of the location to which the u-wind and v-wind values are associated, and the black segments are their projection, therefore the end of the segment would be considered "downwind". Note that the overall pattern is conserved, and that the interpolation enables more matches, but consistently misses some.

Figure A.3: Breakdown of usage by crop

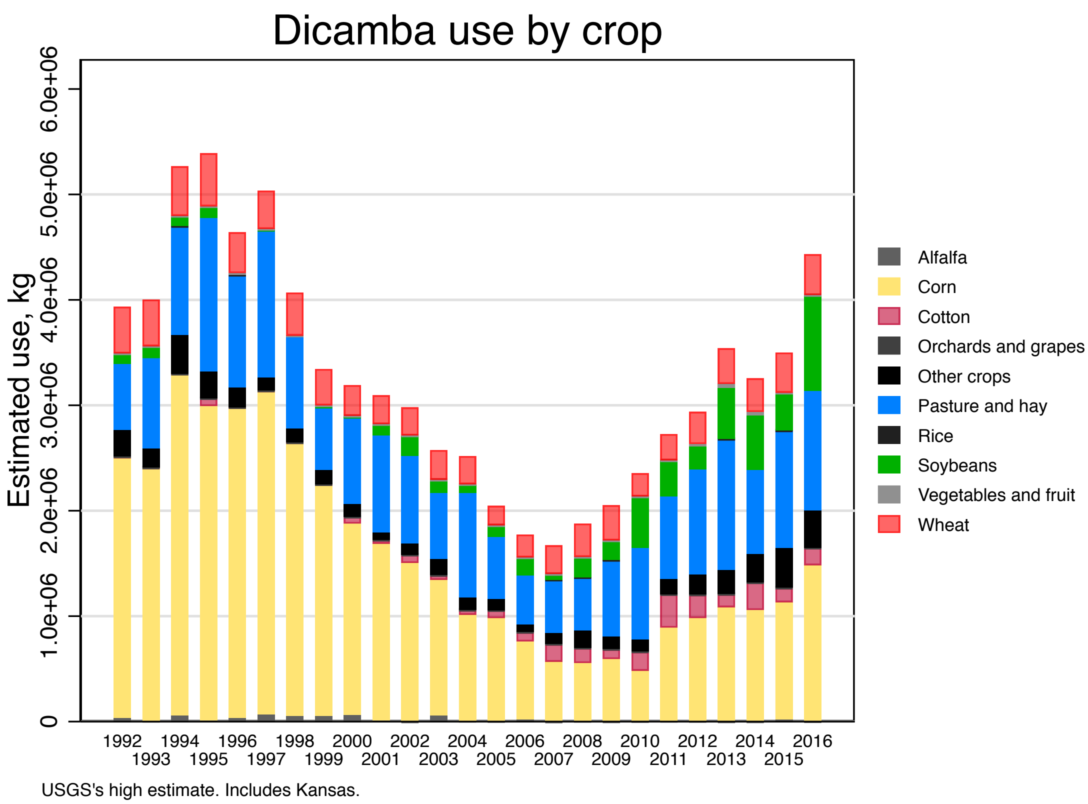

Notes: graph shows dicamba use by year for the conterminous United States, broken down by crop (data: see section 1.2.2). In our sample period (2008-2017), corn and wheat (monocots, but also two of the most important U.S. crops by acreage) are major receivers, as are pasture and hay land uses. The data by crop only goes as far as 2016 as of the writing of this version; 2017 totals for dicamba use are 9.05×10^6 kg.

Figure A.4: Distribution of dicamba anomaly, normalized per ha county area

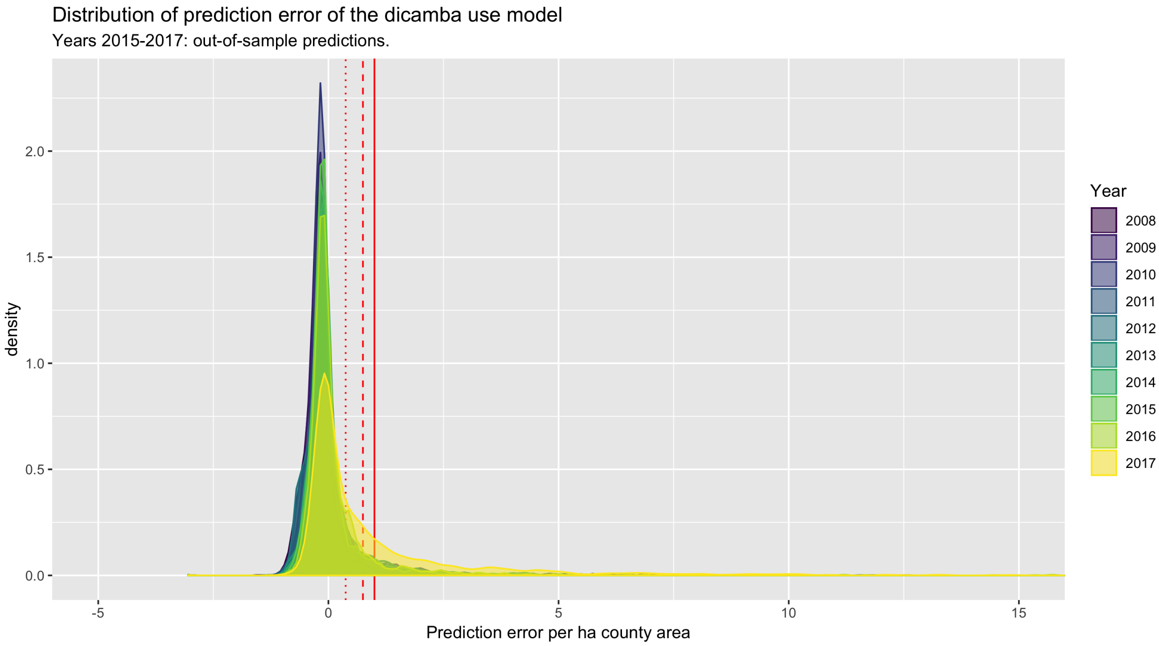

Notes: Figure displays the distribution of dicamba anomaly (homogeneous to kilograms) normalized by the surface of the county (in hectares) for every year in the data; years 2008-2014 correspond to in-sample error ($\hat{\varepsilon}_{it}$), errors for years 2015-2017 are out-of-sample errors. Note the symmetry around zero and small standard deviation of the distributions for years 2008-2014 ($\sigma < 0.9$) and the positive skew, larger standard deviation of later years ($\sigma_{2016} = 1.36$, $\sigma_{2017} = 2.82$). The red dotted line marks an arbitrary threshold at 0.375, the red dashed line at 0.75, the red solid line at 1.00. The plot was truncated at $x = 15$ to facilitate reading (the full support is $[-3.05; 41.47]$ with values beyond 14 kg/ha attained only in 2016 and 2017).

Figure A.5: Adoption models: marginal effects – Logit

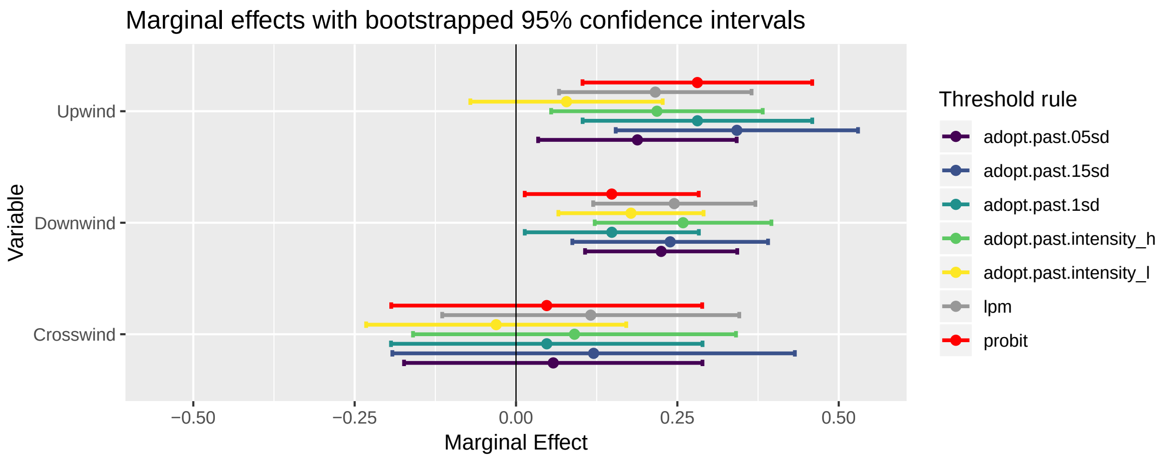

Notes: Plot shows the marginal effects for the logit approach to estimating the adoption model (probability of adoption in 2017 as a function of neighbors' adoption in 2016), with added for reference, the marginal effects associated with the probit model (in red) and the coefficients of the linear probability model (in grey) with the preferred threshold (moderate) already presented in the main text (Figure 1.4).

Figure A.6: Spatial distribution of soybean and cotton growing areas

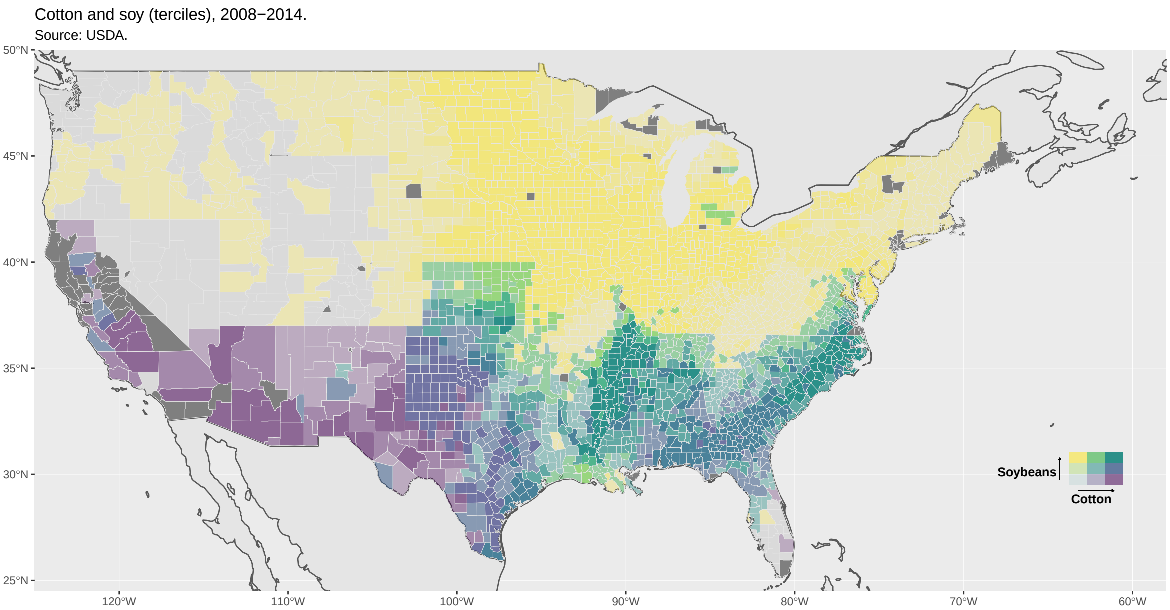

Notes: Maps shows arrangement in space of cotton and soybean growing areas in the pre-DT period (2008-2014) by tercile, based on yearly CDL estimates. Darker shades of purple indicate counties with larger cotton growing areas, more intense shades of yellow counties with larger soybean growing areas, darker blue-green hues counties with jointly larger soybean and cotton growing areas. Note that the total surface under soybeans is about 7.5 times larger than that under cotton in the United States (in 2017 figure), and that there are more counties with any soybean than counties with any cotton; tercile values and sizes differ therefore vastly for the two crops.

Figure A.7: Soybean and cotton prices

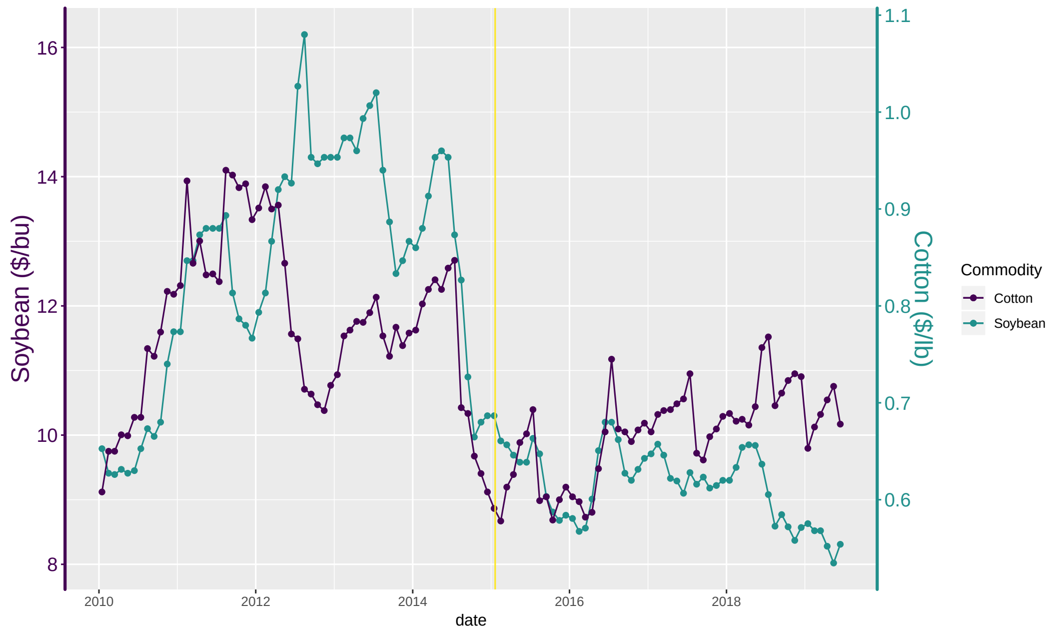

Notes: Graph plots monthly prices received for U.S. soybean (purple, left y-axis) and cotton (teal, right y-axis). Source: USDA. The yellow line corresponds to January 20th, 2017, and the deregulation by the USDA-APHIS of dicamba-tolerant seeds, marking the onset of their commercialization.

Figure A.8: Wheat and corn prices

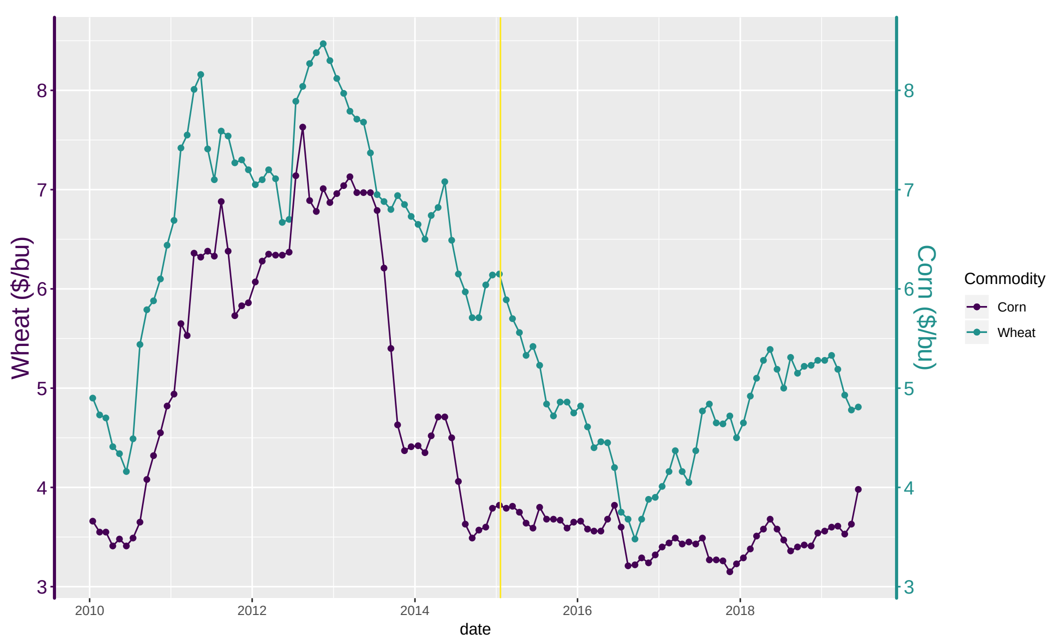

Notes: Graph plots monthly prices received for U.S. wheat (purple, left y-axis) and corn (teal, right y-axis). Source: USDA. The yellow line corresponds to January 20th, 2017.

Figure A.9: Cotton, soybeans, corn, wheat: Relative prices

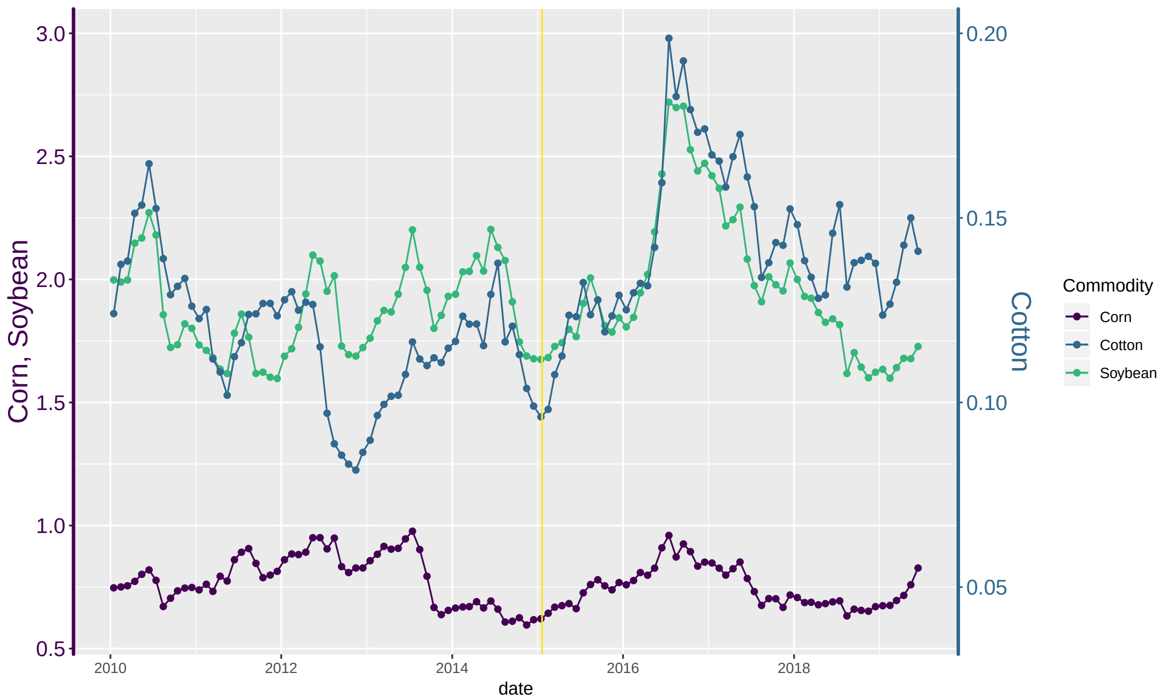

Notes: Graph plots monthly prices received for U.S. corn (purple, left y-axis), soybeans (green, left right y-axis) and cotton (teal, right y-axis) relative to that of wheat in the same month. Source: USDA. The prices of wheat, corn, soybeans are in dollars per bushel, that of cotton is in dollars per pounds.The yellow line corresponds to January 20th, 2017.

Figure A.10: Glyphosate use by state

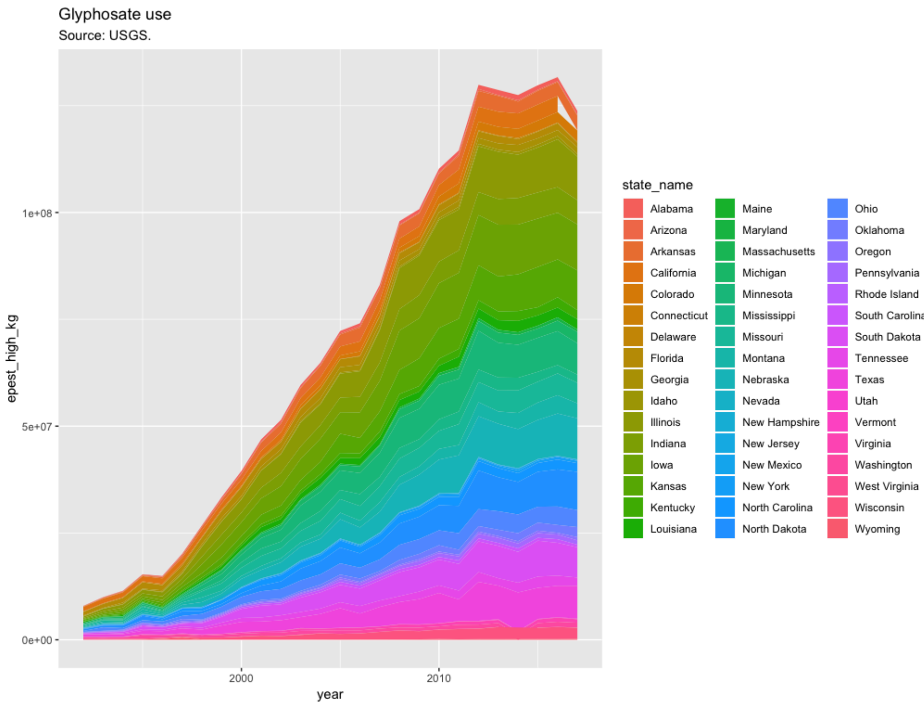

Notes: Graph plots stacked time series of glyphosate use (kg) by U.S. state (one state, one color) using USGS (NAWQA, see Section 1.2 data (high estimate). At the time of writing, the estimates for pesticide use in 2017 in California have not been released.

Figure A.11

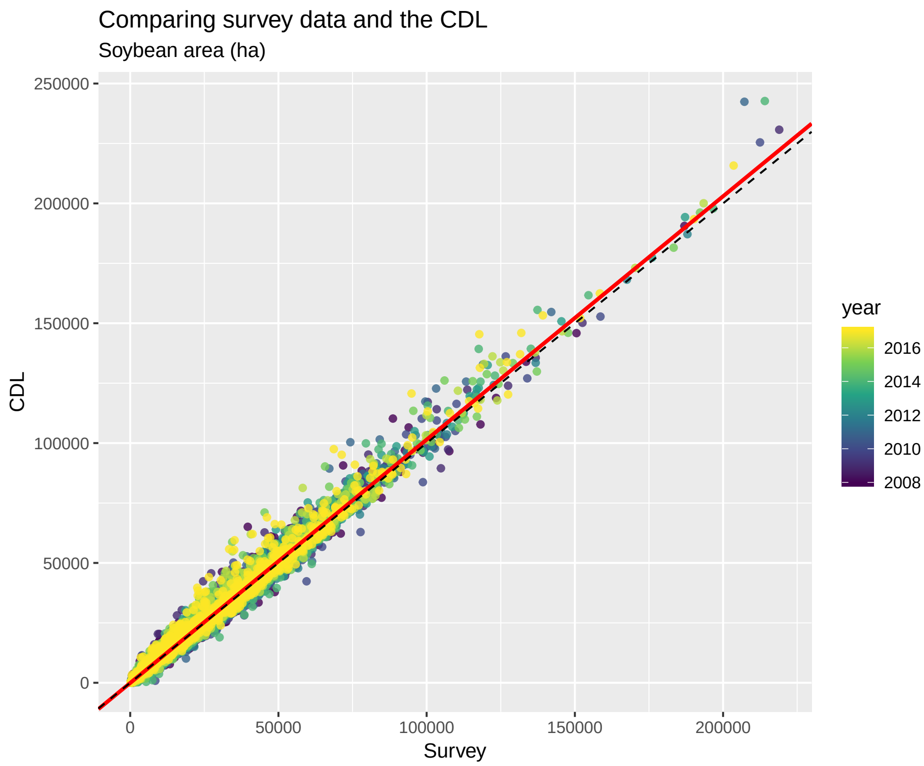

Notes: Graph plots county area planted in soybean (in hectares, i.e.10^4m) as calculated by NASS (USDA) from surveys, against area calculated from the CDL, 2008-2017. Color denotes survey/layer year. Slope (red sold line): 1.02 (p-value < 0.01), R^2: 0.98. The black dashed line corresponds to the 45° line where the CDL and survey areas coincide perfectly.

Figure A.12: Evidence of dicamba drift and damage by RGB drone imaging

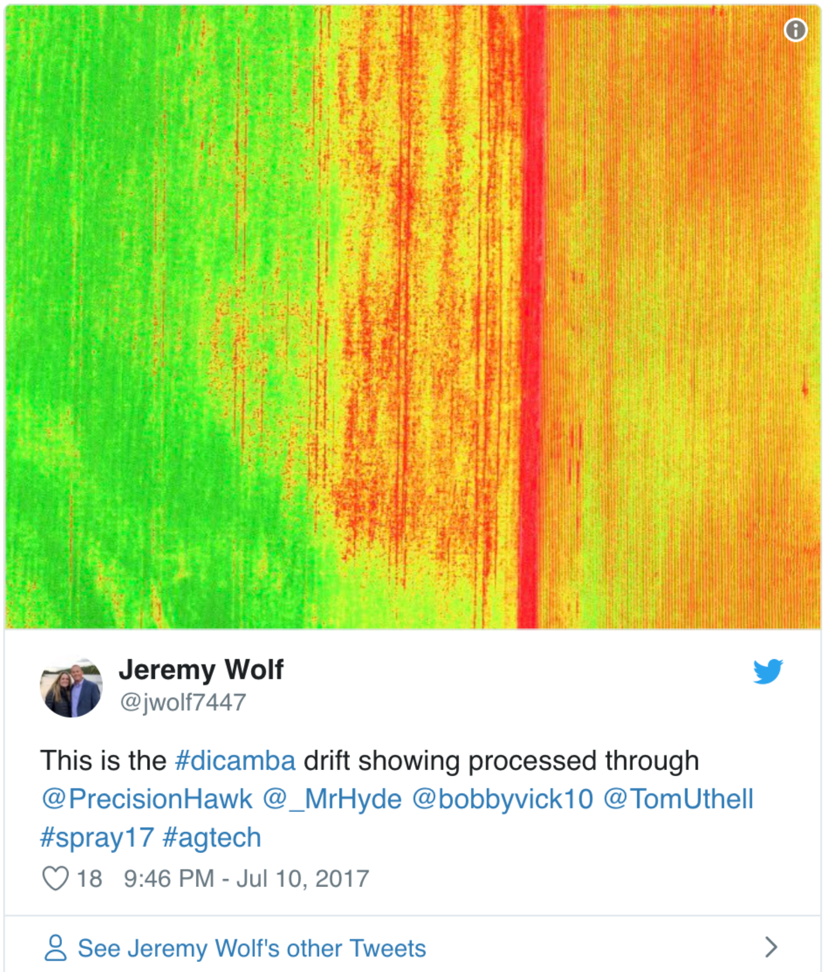

Notes: Tweet by farmer Jeremy Wolf (@jwolf7447) of Homer, Illinois (10/07/2017). Green corresponds to photosynthetic material (healthy plants), red to non-photosynthetic surfaces (road, dead plants). Reproduced with permission.

Figure A.13: State maps

(a) Adoption (% of counties)

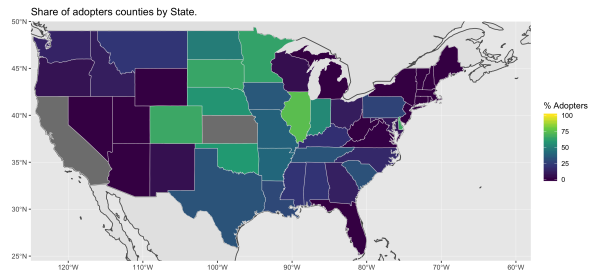

(b) Soybean injury

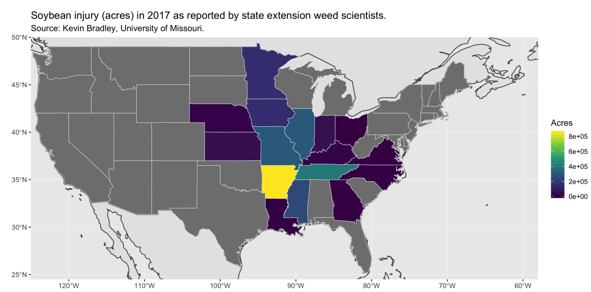

(c) Tweet volume

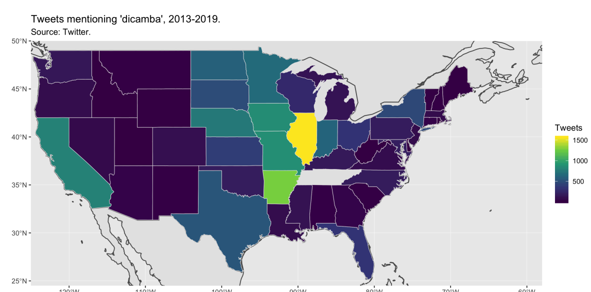

Notes: Maps show various proxies for adoption at the state level. In (a) the share of counties found as "adopters" by the method described in 1.3.1. (b) plots the injured soybean areas found by Dr. Kevin Bradley from University of Missouri for 2017; note that the brightest state is Missouri. (c) shows tweet volume for tweets that containted "dicamba".

A.2 Additional tables

A.2.1 Dicamba use models: alternative samples

	(1)	(2)	(3)	(4)
Cropland (ha)	**0.016*****	**0.009*****	**0.009*****	**0.004*****
	(0.000)	(0.000)	(0.000)	(0.001)
Wheat (ha)		**0.029*****	**0.024*****	**0.031*****
		(0.001)	(0.001)	(0.001)
Pasture/hay (ha)			**0.002*****	**0.002*****
			(0.000)	(0.000)
Corn (ha)				**0.011*****
				(0.001)
(Intercept)	**252.906*****	**337.354*****	**230.416*****	**244.805*****
	(18.377)	(17.971)	(19.054)	(19.092)
R^2	0.140	0.192	0.202	0.205
Dep. var. mean	928.7	928.7	928.7	928.7
Num. obs.	19718	19718	19718	19718
RMSE	1960.445	1900.347	1888.061	1884.575

Notes: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. Table shows results for regression of dicamba use (in kg) at the county level against agricultural land uses to obtain in years 2008-2014 to obtain typical per-hectare dicamba use (kg/ha). The models in columns (2)–(4) further add crop acreage for two land uses that are major receivers of dicamba according to USDA statistics, namely wheat (column 2), hay/pasture (column 3), and corn (column 4). Includes Kansas.

Table A.1: Dicamba use models – Including Kansas

A.2.2 Comparing counties in the wind corridor and cross-wind from adopters

Table A.2: Comparing counties that have an adopter upwind, crosswind (2013)

	Variable	Upwind	Crosswind	t-stat
1	Area (m^2)	2.13e+09	2.11e+09	0.1577
2	Population	3.10e+04	3.22e+04	-0.1191
3	Income per capita (USD)	2.44e+04	2.41e+04	0.3453
4	Median Income (USD)	4.77e+04	4.70e+04	0.4298
5	Unemployment (%)	2.95	3.29	-1.3970
6	Population employed in natural resources sectors	6.77e+02	7.80e+02	-0.8402
7	Population employed in natural resources sectors (%)	13.4	12.0	0.9364
8	Dicamba use (kg)	3.28e+03	2.73e+03	0.9084
9	Glyphosate use (kg)	7.37e+04	7.18e+04	0.1780
10	Dicamba use (kg/ha)	1.51e-02	1.30e-02	0.9271
11	Wheat (acres planted)	6.42e+03	5.57e+03	0.3949
12	Corn (acres planted)	6.45e+04	6.24e+04	0.1610
13	Soybean (acres planted)	4.85e+04	4.61e+04	0.1947
14	Cotton yield (lb/acre)	8.88e+02	8.83e+02	0.0740
15	Soybean yield (bu/acre)	44.1	42.9	0.5250

Notes: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. Table presents t-tests comparing key economic, demographic, and agricultural variables (taken in year 2013, in the pre-DT period) for counties upwind and crosswind from DT adopters. Data from the US Census Bureau, the USDA, the USGS.

Table A.3: Comparing counties that have an adopter downwind, crosswind (2013)

	Variable	Downwind	Crosswind	t-stat
1	Area (m^2)	2.02e+09	2.11e+09	-0.613
2	Population	3.76e+04	3.22e+04	0.518
3	Income per capita (USD)	2.48e+04	2.41e+04	0.907
4	Median Income (USD)	4.83e+04	4.70e+04	0.871
5	Unemployment (%)	3.18	3.29	-0.440
6	Population employed in natural resources sectors	7.64e+02	7.80e+02	-0.097
7	Population employed in natural resources sectors (%)	11.8	12.0	-0.161
8	Dicamba use (kg)	2.60e+03	2.73e+03	-0.287
9	Glyphosate use (kg)	7.65e+04	7.18e+04	0.449
10	Dicamba use (kg/ha)	1.38e-02	1.30e-02	0.408
11	Wheat (acres planted)	5.77e+03	5.57e+03	0.097
12	Corn (acres planted)	6.87e+04	6.24e+04	0.494
13	Soybean (acres planted)	5.31e+04	4.61e+04	0.566
14	Cotton yield (lb/acre)	8.46e+02	8.83e+02	-0.611
15	Soybean yield (bu/acre)	43.9	42.9	0.450

Notes: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. Table presents t-tests comparing key economic, demographic, and agricultural variables (taken in year 2013, in the pre-DT period) for counties downwind and crosswind from DT adopters. Data from the US Census Bureau, the USDA, the USGS.

A.2.3 Adoption models: alternative specifications and samples

Logit models

	(1)	(2)	(3)	(4)	(5)
Upwind, but not downwind 2016 adopter	$\mathbf{1.18}^{**}$		$\mathbf{1.18}^{**}$		
	(0.37)		(0.37)		
Downwind, but not upwind 2016 adopter	$\mathbf{0.65}^{*}$		$\mathbf{0.65}^{*}$		
	(0.28)		(0.28)		
Crosswind 2016 adopter		0.19	0.22		0.26
		(0.55)	(0.55)		(0.55)
Up- or downwind 2016 adopter				$\mathbf{1.24}^{***}$	$\mathbf{1.24}^{***}$
				(0.19)	(0.19)
(Intercept)	$\mathbf{-0.91}^{***}$	$\mathbf{-0.88}^{***}$	$\mathbf{-0.91}^{***}$	$\mathbf{-0.95}^{***}$	$\mathbf{-0.95}^{***}$
	(0.04)	(0.04)	(0.04)	(0.04)	(0.04)
AIC	3134.69	3147.52	3136.53	3105.98	3107.77
BIC	3152.28	3159.24	3159.98	3117.70	3125.36
Log Likelihood	-1564.34	-1571.76	-1564.26	-1550.99	-1550.88
Deviance	3128.69	3143.52	3128.53	3101.98	3101.77
Num. obs.	2600	2600	2600	2600	2600

Notes: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. Table shows regression results (coefficients and standard errors) for logit models regressing adopter status in 2017 on adopter status in 2016 of up-, down-, cross-wind neighbors. Adoption is defined using the preferred dicamba use model and an absolute threshold on out-of-sample prediction error. All regressions exclude Kansas.

Table A.4: Probability of adoption: logit models

Linear probability models

	(1)	(2)	(3)	(4)	(5)
Upwind, but not downwind 2016 adopter	**0.28*****		**0.28*****		
	(0.08)		(0.08)		
Downwind, but not upwind 2016 adopter	**0.15***		**0.15***		
	(0.06)		(0.06)		
Crosswind 2016 adopter		0.04	0.05		0.05
		(0.12)	(0.12)		(0.12)
Up- or downwind 2016 adopter				**0.29*****	**0.29*****
				(0.04)	(0.04)
(Intercept)	**0.29*****	**0.29*****	**0.29*****	**0.28*****	**0.28*****
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
R^2	0.01	0.00	0.01	0.02	0.02
Adj. R^2	0.01	-0.00	0.01	0.02	0.02
Num. obs.	2600	2600	2600	2600	2600
RMSE	0.45	0.46	0.45	0.45	0.45

Notes: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. Table shows regression results (coefficients and standard errors) for linear probability models regressing adopter status in 2017 on adopter status in 2016 of up-, down-, crosswind neighbors. Adoption is defined using the preferred dicamba use model and an absolute threshold on out-of-sample prediction error. All regressions exclude Kansas.

Table A.5: Probability of adoption: linear probability models

Main specification, including Kansas

	(1)	(2)	(3)	(4)	(5)
Upwind, but not downwind 2016 adopter	$\mathbf{0.70}^{**}$		$\mathbf{0.71}^{**}$		
	(0.23)		(0.23)		
Downwind, but not upwind 2016 adopter	$\mathbf{0.62}^{***}$		$\mathbf{0.62}^{***}$		
	(0.18)		(0.18)		
Crosswind 2016 adopter		0.29	0.32		0.29
		(0.39)	(0.39)		(0.39)
Up- or downwind 2016 adopter				0.13	0.13
				(0.31)	(0.31)
(Intercept)	$\mathbf{-0.66}^{***}$	$\mathbf{-0.64}^{***}$	$\mathbf{-0.67}^{***}$	$\mathbf{-0.64}^{***}$	$\mathbf{-0.64}^{***}$
	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)
AIC	3088.37	3107.22	3089.71	3107.60	3109.03
BIC	3106.08	3119.03	3113.32	3119.41	3126.74
Log Likelihood	-1541.18	-1551.61	-1540.86	-1551.80	-1551.52
Deviance	3082.37	3103.22	3081.71	3103.60	3103.03
Num. obs.	2705	2705	2705	2705	2705

Notes: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. Table shows regression results for probit models regressing adopter status in 2017 on adopter status in 2016 of up-, down-, cross-wind neighbors. Includes Kansas.

Table A.6: Probability of adoption – Including Kansas

Main specification, alternative wind months

	(1)	(2)	(3)	(4)	(5)
Upwind, but not downwind 2016 adopter	$\mathbf{0.63}^{*}$		$\mathbf{0.63}^{*}$		
	(0.31)		(0.31)		
Downwind, but not upwind 2016 adopter	$\mathbf{0.51}^{*}$		$\mathbf{0.50}^{*}$		
	(0.25)		(0.25)		
Crosswind 2016 adopter		-0.30	-0.29		-0.27
		(0.64)	(0.64)		(0.64)
Up- or downwind 2016 adopter				$\mathbf{0.84}^{***}$	$\mathbf{0.84}^{***}$
				(0.16)	(0.16)
(Intercept)	$\mathbf{-0.56}^{***}$	$\mathbf{-0.55}^{***}$	$\mathbf{-0.55}^{***}$	$\mathbf{-0.57}^{***}$	$\mathbf{-0.57}^{***}$
	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)
AIC	3141.47	3147.41	3143.26	3119.02	3120.83
BIC	3159.06	3159.14	3166.71	3130.74	3138.42
Log Likelihood	-1567.73	-1571.70	-1567.63	-1557.51	-1557.41
Deviance	3135.47	3143.41	3135.26	3115.02	3114.83
Num. obs.	2600	2600	2600	2600	2600

Notes: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. Table shows regression results (coefficients and standard errors) for probit models regressing adopter status in 2017 on adopter status in 2016 of up-, down-, cross-wind neighbors. *County relationships with respect to the wind are determined based on May 2016 data.* Adoption is defined using the preferred dicamba use model and an absolute threshold on out-of-sample prediction error. May 2016 winds determine up-, down-, and cross-wind relationships. All regressions exclude Kansas.

Table A.7: Probability of adoption – May 2016 wind

	(1)	(2)	(3)	(4)	(5)
Upwind, but not downwind 2016 adopter	0.28		0.28		
	(0.27)		(0.27)		
Downwind, but not upwind 2016 adopter	**0.49***		**0.49***		
	(0.21)		(0.21)		
Crosswind 2016 adopter		−0.30	−0.29		−0.27
		(0.64)	(0.64)		(0.64)
Up- or downwind 2016 adopter				**0.77*****	**0.77*****
				(0.14)	(0.14)
(Intercept)	**−0.56*****	**−0.55*****	**−0.56*****	**−0.58*****	**−0.58*****
	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)
AIC	3142.94	3147.41	3144.73	3114.84	3116.66
BIC	3160.53	3159.14	3168.19	3126.57	3134.25
Log Likelihood	-1568.47	-1571.70	-1568.37	-1555.42	-1555.33
Deviance	3136.94	3143.41	3136.73	3110.84	3110.66
Num. obs.	2600	2600	2600	2600	2600

Notes: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. Table shows regression results (coefficients and standard errors) for probit models regressing adopter status in 2017 on adopter status in 2016 of up-, down-, cross-wind neighbors. *County relationships with respect to the wind are determined based on July 2016 data.* Adoption is defined using the preferred dicamba use model and an absolute threshold on out-of-sample prediction error. All regressions exclude Kansas.

Table A.8: Probability of adoption – July 2016 wind

Main specification, interpolated wind data

	(1)	(2)	(3)	(4)	(5)
Upwind, but not downwind 2016 adopter	$\mathbf{0.60}^{**}$		$\mathbf{0.60}^{**}$		
	(0.20)		(0.20)		
Downwind, but not upwind 2016 adopter	$\mathbf{0.48}^{**}$		$\mathbf{0.48}^{**}$		
	(0.16)		(0.16)		
Crosswind 2016 adopter		-0.22	-0.20		-0.17
		(0.33)	(0.33)		(0.33)
Up- or downwind 2016 adopter				$\mathbf{0.82}^{***}$	$\mathbf{0.82}^{***}$
				(0.11)	(0.11)
(Intercept)	$\mathbf{-0.57}^{***}$	$\mathbf{-0.54}^{***}$	$\mathbf{-0.57}^{***}$	$\mathbf{-0.60}^{***}$	$\mathbf{-0.60}^{***}$
	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)
AIC	3132.20	3147.17	3133.84	3087.99	3089.72
BIC	3149.79	3158.90	3157.29	3099.71	3107.31
Log Likelihood	-1563.10	-1571.59	-1562.92	-1541.99	-1541.86
Deviance	3126.20	3143.17	3125.84	3083.99	3083.72
Num. obs.	2600	2600	2600	2600	2600

Notes: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. Table shows regression results (coefficients and standard errors) for probit models regressing adopter status in 2017 on adopter status in 2016 of up-, down-, cross-wind neighbors. Adoption is defined using the preferred dicamba use model and an absolute threshold on out-of-sample prediction error. Interpolated wind data. All regressions exclude Kansas.

Table A.9: Probability of adoption

Alternative definitions of adoption based on "level" (anomaly)

	(1)	(2)	(3)	(4)	(5)
Upwind, but not downwind 2016 adopter	$\mathbf{0.47}^{*}$		$\mathbf{0.48}^{*}$		
	(0.20)		(0.20)		
Downwind, but not upwind 2016 adopter	$\mathbf{0.57}^{***}$		$\mathbf{0.57}^{***}$		
	(0.16)		(0.16)		
Crosswind 2016 adopter		0.12	0.15		0.18
		(0.30)	(0.30)		(0.30)
Up- or downwind 2016 adopter				$\mathbf{0.82}^{***}$	$\mathbf{0.83}^{***}$
				(0.10)	(0.10)
(Intercept)	$\mathbf{-0.29}^{***}$	$\mathbf{-0.26}^{***}$	$\mathbf{-0.29}^{***}$	$\mathbf{-0.32}^{***}$	$\mathbf{-0.32}^{***}$
	(0.03)	(0.02)	(0.03)	(0.03)	(0.03)
AIC	3478.75	3495.22	3480.50	3428.29	3429.92
BIC	3496.34	3506.94	3503.95	3440.02	3447.51
Log Likelihood	-1736.37	-1745.61	-1736.25	-1712.15	-1711.96
Deviance	3472.75	3491.22	3472.50	3424.29	3423.92
Num. obs.	2600	2600	2600	2600	2600

Notes: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. Table shows regression results (coefficients and standard errors) for probit models regressing adopter status in 2017 on adopter status in 2016 of up-, down-, cross-wind neighbors. Adoption is defined using the preferred dicamba use model and an absolute threshold on out-of-sample prediction error. All regressions exclude Kansas.

Table A.10: Probability of adoption – Adoption based on level (low threshold)

	(1)	(2)	(3)	(4)	(5)
Upwind, but not downwind 2016 adopter	$\mathbf{0.93}^{***}$		$\mathbf{0.93}^{***}$		
	(0.24)		(0.24)		
Downwind, but not upwind 2016 adopter	$\mathbf{0.67}^{***}$		$\mathbf{0.67}^{***}$		
	(0.20)		(0.20)		
Crosswind 2016 adopter		0.33	0.36		0.38
		(0.43)	(0.43)		(0.43)
Up- or downwind 2016 adopter				$\mathbf{0.98}^{***}$	$\mathbf{0.98}^{***}$
				(0.14)	(0.14)
(Intercept)	$\mathbf{-0.79}^{***}$	$\mathbf{-0.76}^{***}$	$\mathbf{-0.79}^{***}$	$\mathbf{-0.80}^{***}$	$\mathbf{-0.81}^{***}$
	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)
AIC	2737.90	2760.95	2739.23	2710.39	2711.66
BIC	2755.49	2772.68	2762.69	2722.11	2729.25
Log Likelihood	-1365.95	-1378.48	-1365.62	-1353.19	-1352.83
Deviance	2731.90	2756.95	2731.23	2706.39	2705.66
Num. obs.	2600	2600	2600	2600	2600

Notes: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. Table shows regression results (coefficients and standard errors) for probit models regressing adopter status in 2017 on adopter status in 2016 of up-, down-, cross-wind neighbors. Adoption is defined using the preferred dicamba use model and an absolute threshold on out-of-sample prediction error. All regressions exclude Kansas.

Table A.11: Probability of adoption – Adoption based on level (high threshold)

Alternative definition of adoption based on "intensity" (anomaly per hectare)

	(1)	(2)	(3)	(4)	(5)
Upwind, but not downwind 2016 adopter	0.20		0.20		
	(0.19)		(0.19)		
Downwind, but not upwind 2016 adopter	**0.45****		**0.45****		
	(0.15)		(0.15)		
Crosswind 2016 adopter		-0.10	-0.08		-0.05
		(0.27)	(0.27)		(0.27)
Up- or downwind 2016 adopter				**0.66*****	**0.65*****
				(0.10)	(0.10)
(Intercept)	$\mathbf{-0.20}^{***}$	$\mathbf{-0.18}^{***}$	$\mathbf{-0.20}^{***}$	$\mathbf{-0.23}^{***}$	$\mathbf{-0.23}^{***}$
	(0.03)	(0.02)	(0.03)	(0.03)	(0.03)
AIC	3546.10	3554.25	3548.01	3507.48	3509.45
BIC	3563.69	3565.98	3571.47	3519.20	3527.04
Log Likelihood	-1770.05	-1775.13	-1770.01	-1751.74	-1751.72
Deviance	3540.10	3550.25	3540.01	3503.48	3503.45
Num. obs.	2600	2600	2600	2600	2600

Notes: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. Table shows regression results (coefficients and standard errors) for probit models regressing adopter status in 2017 on adopter status in 2016 of up-, down-, crosswind neighbors. Adoption is defined using the preferred dicamba use model and an absolute threshold on out-of-sample prediction error. All regressions exclude Kansas.

Table A.12: Probability of adoption – Adoption based on intensity (low threshold)

	(1)	(2)	(3)	(4)	(5)
Upwind, but not downwind 2016 adopter	$\mathbf{0.56}^{**}$		$\mathbf{0.56}^{**}$		
	(0.21)		(0.21)		
Downwind, but not upwind 2016 adopter	$\mathbf{0.67}^{***}$		$\mathbf{0.67}^{***}$		
	(0.18)		(0.18)		
Crosswind 2016 adopter		0.22	0.24		0.34
		(0.33)	(0.33)		(0.33)
Up- or downwind 2016 adopter				$\mathbf{0.79}^{***}$	$\mathbf{0.80}^{***}$
				(0.12)	(0.12)
(Intercept)	$\mathbf{-0.49}^{***}$	$\mathbf{-0.47}^{***}$	$\mathbf{-0.50}^{***}$	$\mathbf{-0.51}^{***}$	$\mathbf{-0.59}^{***}$
	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)
AIC	3240.28	3258.74	3241.74	3214.62	3102.58
BIC	3257.87	3270.47	3265.19	3226.35	3120.17
Log Likelihood	-1617.14	-1627.37	-1616.87	-1605.31	-1548.29
Deviance	3234.28	3254.74	3233.74	3210.62	3096.58
Num. obs.	2600	2600	2600	2600	2600

Notes: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. Table shows regression results (coefficients and standard errors) for probit models regressing adopter status in 2017 on adopter status in 2016 of up-, down-, cross-wind neighbors. Adoption is defined using the preferred dicamba use model and an absolute threshold on out-of-sample prediction error. All regressions exclude Kansas.

Table A.13: Probability of adoption – Adoption based on intensity (high threshold)

A.2.4 Land use change models: alternative specifications

Treatment intensity

Farmers' decisions are probably sensitive to the density of adopters surrounding them rather than to a binary indication of adoption. To explore the effect of the intensity of the treatment, I therefore consider:

$$Y_{i,t\to t+1} = \alpha_1.d_{it} + \alpha_2.\left(d_{it} \times \mathbb{1}\left\{t \geq 2015\right\}\right) + \lambda_i + \lambda_t + \varepsilon_{it} \tag{A.1}$$

with $Y_{i,t\to t+1}$ the net conversion between years t and $t+1$ to soybean and cotton from some other land use, d_{it} the dicamba use anomaly (in kg, kg/ha) compared to model predictions (see 1.3.1). λ_i, λ_t are county and year fixed effects, respectively. Errors ε_{it} are clustered at the state level. d_{it} is interacted with a post-DT seed release dummy, $\mathbb{1}\left\{t \geq 2015\right\}$, taking the value of 1 if t is 2015 or later (corresponding to herbicide applications made in 2015 or later, and land-use changes occurring between years 2015 and 2016 and subsequently), 0 otherwise.

	(1)	(2)	(3)	(4)	(5)	(6)
Anomaly x Post	-0.01					
	(0.03)					
Per-hectare Anomaly x Post		80.59				
		(56.03)				
Dicamba x Post			$\mathbf{0.08}^{**}$			
			(0.03)			
Dicamba per hectare x Post				$\mathbf{22268.42}^{***}$		
				(5391.32)		
Log Dicamba x Post					$\mathbf{352.51}^{***}$	
					(31.01)	
Log Dicamba per hectare x Post						$\mathbf{342.81}^{***}$
						(32.40)
R^2	0.00	0.00	0.00	0.00	0.01	0.01
Adj. R^2	-0.12	-0.12	-0.12	-0.12	-0.12	-0.12
Num. obs.	23473	23473	23473	23473	23465	23465

Notes: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. Table shows regression results for net change in soybean and cotton area (in hectares).

Table A.14: Net conversion with continuous treatment

A.2.5 Crop outcome models

Other outcomes: production and yield per harvested acre

In addition to crop failure and yields per planted area, two other outcomes are of interest: yields per harvested area (corresponds to the yield reported by the USDA) and production. The effect of being subjected to neighboring counties ("Treated") or farmers ("Adopted") adopting DT seeds on these two variables is reported in Table A.15 and Table A.16, respectively.

	Soybean			Cotton			Wheat			Corn		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Treated x Post	-0.25			$\mathbf{0.12}^{*}$			-1.34			$\mathbf{-5.41}^{*}$		
	(0.77)			(0.06)			(1.22)			(2.41)		
Placebo x Post		-0.99			0.15			$\mathbf{-9.19}^{**}$			4.95	
		(2.10)			(0.13)			(3.24)			(6.81)	
Adopted x Post			$\mathbf{0.93}^{*}$			-0.07			-0.29			$\mathbf{-5.19}^{***}$
			(0.38)			(0.04)			(0.71)			(1.28)
R^2	0.30	0.30	0.30	0.08	0.08	0.08	0.10	0.10	0.10	0.32	0.32	0.32
Dep. var. mean	42.78	42.78	42.78	1.80	1.80	1.80	55.00	55.00	55.00	141.60	141.60	141.60
Num. obs.	11903	11903	11903	2771	2771	2771	9380	9380	9380	13478	13478	13478

Notes: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. Table shows regression results for fixed-effect models regressing production outcomes on dicamba use anomaly at the county level, for four major crops, for two of which DT technology is available (cotton, soybean), the other two (wheat, corn) are relatively tolerant to dicamba, owing to their being monocotyledones. Areas are expressed in acres, production in bushels (soybean, wheat, corn) or pounds (cotton), and yields are accordingly in bushels or pounds per acre. Includes weather controls.

Table A.15: Yield models (harvested)

	Soybean			Cotton			Wheat			Corn		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Treated x Post	13901.50			**24498.51*****			−77737.24			−387586.72		
	(84234.37)			(4081.19)			(82163.94)			(251919.31)		
Placebo x Post		−255568.60			4533.60			−341977.13			160166.86	
		(230326.90)			(9115.05)			(218166.51)			(711788.83)	
Adopted x Post			**712071.80*****			**12010.80*****			**−146176.72****			**554196.19*****
			(40254.85)			(3133.91)			(48200.53)			(133021.48)
R^2	0.09	0.09	0.14	0.18	0.17	0.18	0.09	0.09	0.09	0.09	0.09	0.09
Dep. var. mean	2490910	2490910	2490910	43751	43751	43751	929694	929694	929694	8089558	8089558	8089558
Num. obs.	11903	11903	11903	2771	2771	2771	9380	9380	9380	13478	13478	13478

Notes: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. Table shows regression results for fixed-effect models regressing production outcomes on dicamba use anomaly at the county level, for four major crops, for two of which DT technology is available (cotton, soybean), the other two (wheat, corn) are relatively tolerant to dicamba, owing to their being monocotyledones. Production in bushels (soybean, wheat, corn) or pounds (cotton). Includes weather controls.

Table A.16: Production models

Alternative specifications

Similar to the analysis presented in the body of this paper, I consider chiefly two outcome variables: yield (in physical quantity per *planted*[1] acre), and the planted-to-harvested ratio, to account for the fact that farmers observing damage beyond repair do not harvest those fields. The difference with the specifications presented above resides in the de-discretization of the treatment, i.e. the use of the anomaly itself *in the county of interest*, as the regressor:

$$Y_{it} = \alpha_1.\mathrm{d}_{it} + \alpha_2.\left(\mathrm{d}_{it} \times \mathbb{1}\left\{t \geq 2015\right\}\right) + \mathbf{X_{it}}\theta + \lambda_i + \lambda_t + \varepsilon_{it} \tag{A.2}$$

with Y_{it} the yield (bushel/acre for soybeans, lb/acre for cotton) or planted-to-harvested ratio for the focal crop (soybean, cotton, and two non-target crops, wheat and corn), d_{it} alternatively: the dicamba use anomaly (kg, kg/ha) compared to model predictions (see 1.3.1), dicamba use (in kg, in levels or logs), and dicamba use per hectare (in levels or logs). $\mathbb{1}\left\{t \geq 2015\right\}$ is a dummy variable that is given a value of 1 in 2015 and thereafter. $\mathbf{X_{it}}$ are weather controls. λ_i, λ_t are county and year fixed effects, respectively. Errors ε_{it} are clustered at the state level.

The reason why the interaction with a "post" dummy still makes sense in these cases is that dicamba use, as levels or as anomalies, is expected to be different pre- and post-DT because the anomaly post-DT is attributed to the in-crop use required by DT seeds (i.e. spraying dicamba during the growing season) while pre-DT it should be confined (by law and for agronomic reasons) to the pre-planting period; their magnitudes also differ vastly.

The results are reported in the tables A.17-A.18 below.

[1] USDA NASS reports yield per harvested acre.

	Soybean	Soybean	Soybean	Soybean	Soybean	Soybean	Cotton	Cotton	Cotton	Cotton	Cotton	Cotton
Anomaly x Post	$\mathbf{-0.00}^{***}$						-0.00					
	(0.00)						(0.00)					
Per-hectare Anomaly x Post		$\mathbf{-0.17}^{**}$						-0.21				
		(0.06)						(0.26)				
Dicamba x Post			$\mathbf{-0.00}^{***}$						-0.00			
			(0.00)						(0.00)			
Dicamba per hectare x Post				$\mathbf{-18.49}^{***}$						-17.16		
				(5.45)						(25.27)		
Log Dicamba x Post					$\mathbf{-0.32}^{***}$						-0.07	
					(0.06)						(0.39)	
Log Dicamba per hectare x Post						$\mathbf{-0.33}^{***}$						-0.08
						(0.06)						(0.35)
R^2	0.13	0.13	0.14	0.13	0.14	0.14	0.29	0.29	0.29	0.29	0.29	0.29
Dep. var. mean	97.98	97.98	97.98	97.98	97.98	97.98	92.58	92.58	92.58	92.58	92.58	92.58
Num. obs.	13196	13196	13196	13196	13196	13196	3064	3064	3064	3064	3064	3064

	Wheat	Wheat	Wheat	Wheat	Wheat	Wheat	Corn	Corn	Corn	Corn	Corn	Corn
Anomaly x Post	0.00						0.00					
	(0.00)						(0.00)					
Per-hectare Anomaly x Post		0.15						-0.04				
		(0.20)						(0.11)				
Dicamba x Post			0.00						0.00			
			(0.00)						(0.00)			
Dicamba per hectare x Post				20.40						-5.52		
				(19.64)						(10.65)		
Log Dicamba x Post					$\mathbf{0.49}^{*}$						0.01	
					(0.24)						(0.11)	
Log Dicamba per hectare x Post						$\mathbf{0.50}^{*}$						0.02
						(0.25)						(0.11)
R^2	0.05	0.05	0.05	0.05	0.05	0.05	0.09	0.09	0.09	0.09	0.09	0.09
Dep. var. mean	81.17	81.17	81.17	81.17	81.17	81.17	86.63	86.63	86.63	86.63	86.63	86.63
Num. obs.	10270	10270	10270	10270	10269	10269	14920	14920	14920	14920	14919	14919

Notes: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. Table shows regression results for fixed-effect models regressing production outcomes on dicamba use anomaly at the county level, for four major crops; for two of them (soybean, cotton), the DT technology is available and absent the DT trait, they are extremely vulnerable to dicamba, while for the other two (wheat, corn) the DT technology is not available but they are relatively tolerant to dicamba. Includes weather controls.

Table A.17: Harvested-to-planted models

	Soybean	Soybean	Soybean	Soybean	Soybean	Soybean	Cotton	Cotton	Cotton	Cotton	Cotton	Cotton
Anomaly x Post	**−0.00*****						−0.00					
	(0.00)						(0.00)					
Per-hectare Anomaly x Post		**−0.26****						0.01				
		(0.09)						(0.01)				
Dicamba x Post			**−0.00****						0.00			
			(0.00)						(0.00)			
Dicamba per hectare x Post				−10.74						0.90		
				(8.59)						(0.84)		
Log Dicamba x Post					**0.58*****						**0.04****	
					(0.09)						(0.01)	
Log Dicamba per hectare x Post						**0.75*****						0.02
						(0.10)						(0.01)
R^2	0.27	0.27	0.27	0.27	0.27	0.28	0.12	0.12	0.12	0.12	0.12	0.12
Dep. var. mean	42.09	42.09	42.09	42.09	42.09	42.09	1.70	1.70	1.70	1.70	1.70	1.70
Num. obs.	13196	13196	13196	13196	13196	13196	3064	3064	3064	3064	3064	3064

	Wheat	Wheat	Wheat	Wheat	Wheat	Wheat	Corn	Corn	Corn	Corn	Corn	Corn
Anomaly x Post	0.00						**−0.00*****					
	(0.00)						(0.00)					
Per-hectare Anomaly x Post		−0.08						**−2.07*****				
		(0.17)						(0.31)				
Dicamba x Post			0.00						**−0.00*****			
			(0.00)						(0.00)			
Dicamba per hectare x Post				9.80						**−183.52*****		
				(16.77)						(29.83)		
Log Dicamba x Post					**0.81*****						−0.28	
					(0.20)						(0.30)	
Log Dicamba per hectare x Post						**1.21*****						0.39
						(0.21)						(0.32)
R^2	0.07	0.07	0.07	0.07	0.07	0.07	0.29	0.29	0.29	0.29	0.29	0.29
Dep. var. mean	46.38	46.38	46.38	46.38	46.38	46.38	124.39	124.39	124.39	124.39	124.39	124.39
Num. obs.	10270	10270	10270	10270	10269	10269	14920	14920	14920	14920	14919	14919

Notes: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. Table shows regression results for fixed-effect models regressing production outcomes on dicamba use anomaly at the county level, for four major crops; for two of them (soybean, cotton), the DT technology is available and absent the DT trait, they are extremely vulnerable to dicamba, while for the other two (wheat, corn) the DT technology is not available but they are relatively tolerant to dicamba. Areas are expressed in acres, production in bushels (soybean, wheat, corn) or pounds (cotton), and yields are accordingly in bushels or pounds per acre. Includes weather controls.

Table A.18: Yield models

A.3 Other additional material

A.3.1 On wind direction

Wind direction is central to this paper. Indeed, the proposed adoption mechanism relies on the physical transportation of herbicide particles from a source to an endpoint, where they may or may not cause observable effects, which may or may not action by farmers. Determining whether farmers indeed take action is the crux of the paper. Therefore handling the wind data properly and in accordance with their empirical manifestation is crucial.

While the effect of counties identified as "downwind" from the focal county on their adoption of the DT seeds can at face value be surprising, a closer look at wind patterns reveals a possible explanation. Indeed, the wind-neighboring relationships were obtained using monthly data, which conveniently provides a single value for wind speed and direction at every location, but observation of daily data over the course of a month, a season (see Figure A.14), a year, provides a more complex picture: in a large share of the counties, the distribution of wind direction is bimodal.

Figure A.14: Example: a location in Arkansas

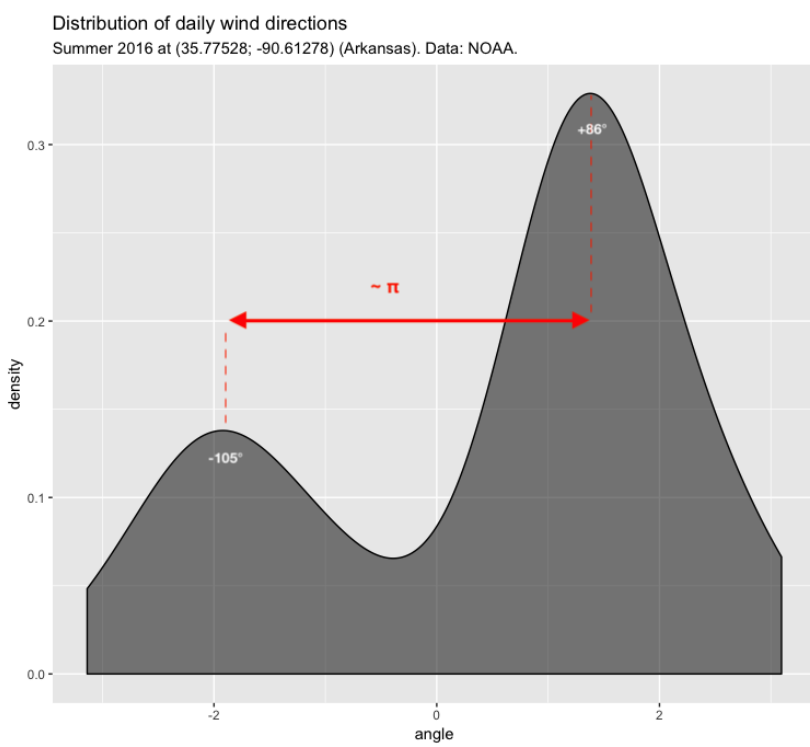

Notes: Graph shows the daily distribution of wind directions during the summer of 2016 (May through August, total of 123 days) at a data point of the NOAA dataset (latitude $\sim 36°$, longitude $\sim -91°$). The data were then fitted to a mixture of two normal distributions, whose means are $\hat{\mu}_1 \simeq -1.8$ and $\hat{\mu}_2 \simeq 1.5$ (converted to degrees on the graph); the difference between the means of the fitted underlying normal distributions is about π or $180°$. Which means that at this location in Arkansas, the wind tends to circulate along a specific corridor, most of the time in a certain direction, and in the opposite direction a significant share of the time.

For instance, in 2016, the distribution of direction was bimodal at about 42.4% of the points in the NOAA dataset and located within the boundaries of a U.S. county (in yellow on Figure A.15). I fitted the distributions at each point over the course of the year to a mixture of two normal distributions, and calculated the difference between the estimated means: as shown on Figure A.15, the overwhelming majority of them differed by π (even when the distribution was not formally diagnosed as bimodal, purple curve), that is to say, while the dominant winds may blow in a direction α most of the time (e.g. 1.5 radians on Figure A.14), they also blow in the exact opposite direction $\alpha \pm \pi$ (-1.8 radians) a significant part of the time.

Therefore that adoption in downwind counties should matter to counties upwind from them shouldn't come as a surprise.

Figure A.15

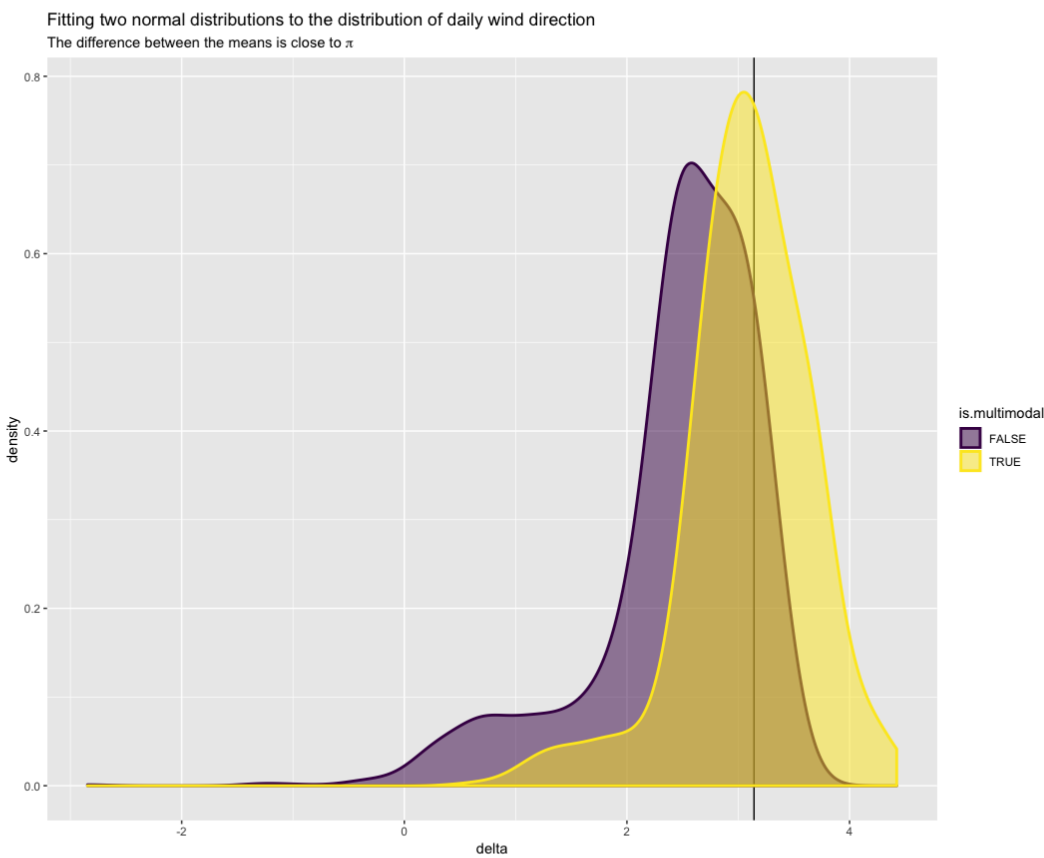

Notes: Graph shows kernel density (band width= 0.2) of $\hat{\mu}_2 - \hat{\mu}_1$ for both points diagnosed with a multimodal distribution of their daily wind directions (yellow), and the others (purple). Both wind direction distributions (in radians) were fitted with two normal distributions, whose means ($\hat{\mu}_1$, $\hat{\mu}_2$), standard deviation ($\hat{\sigma}_1$, $\hat{\sigma}_2$) and weights ($\hat{w}_1 + \hat{w}_2 = 1$) were estimated. The vertical black line indicates pi, that is to say, the point where the wind typically oscillates between a certain direction $alpha°$ and $alpha + 180°$.

A.3.2 On reasons for excluding Kansas from all analyses

Most analyses exclude Kansas; while the results hold with the inclusion of Kansas, the fact that the entire state appears as an outlier as far as dicamba is concerned *before* the advent of the DT seed technology, as evidenced by Figure A.16, exhorts to prudence.

Figure A.16: Dicamba use by state, total for all crops, selection of states

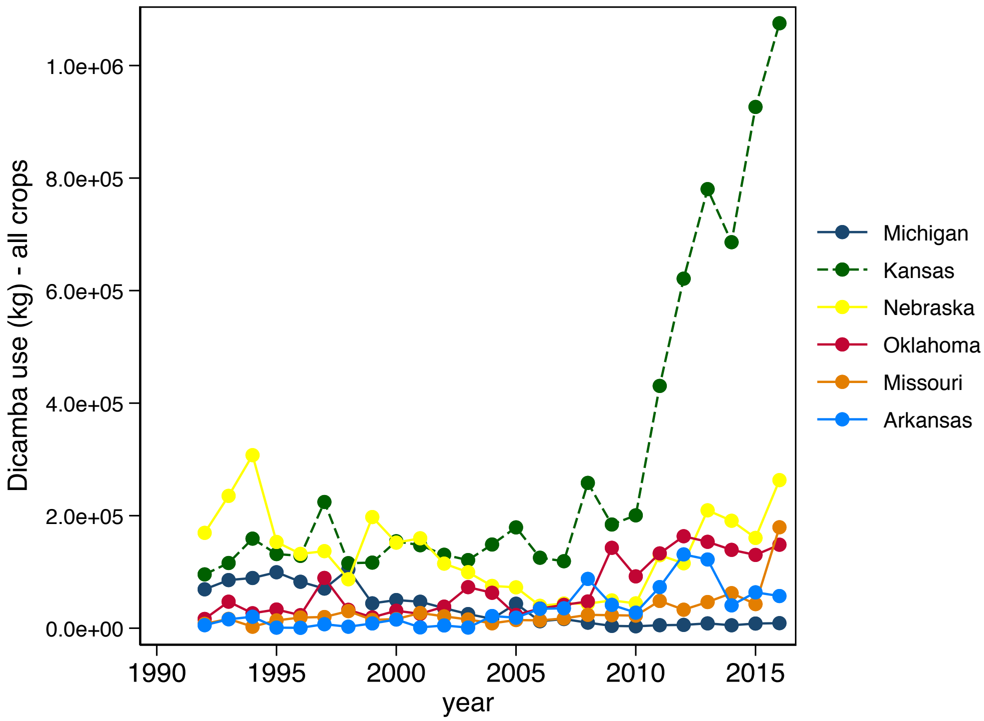

Notes: Graph shows total dicamba use as reported by the USGS (see section 1.2.2) by state and by crop, aggregated over all crops, for a selection of Midwestern states at the heart of the DT seed boom. It appears clearly that while the other states are only starting in 2016 to see, for some of them, an uptake in dicamba use due to the commercialization of DT seeds, dicamba use in Kansas amorced a steep increase from 2010 onwards, with a fivefold increase over the course of 5-6 years, with no change in land-use to match (not shown).

I therefore excluded Kansas from the dicamba use model. While the results are not very different from those obtained without Kansas (see Table 1.1), including Kansas (Table A.19) seems to nudge the coefficients upwards in particular for wheat (of which Kansas is the U.S.'s largest producer) and to degrade the quality of the fit – both consistent with an atypically large use of dicamba in the state.

Absent obvious explanations I questioned the agricultural extension at Kansas State University: glyphosate-resistant kochia (a weed) has appeared 2007, and has been reported in Kansas since about 2010, and dicamba has been a prominent part of the arsenal against it (Burton et al., 2014; Kumar et al., 2019) from the beginning; the problem has been particularly prevalent in the western part of the state, with fallow land and no-till practices being particularly favorable to the proliferation of glyphosate-resistant kochia. Note that the increased use of dicamba to overcome glyphosate-resistant weeds *is not* accompanied by a decrease in glyphosate use (i.e., no substitution): glyphosate, being a broad-spectrum herbicide, is used to kill all unwanted plants, and other herbicides are *added* to take

	(1)	(2)	(3)	(4)
Cropland (ha)	**0.016*****	**0.009*****	**0.009*****	**0.004*****
	(0.000)	(0.000)	(0.000)	(0.001)
Wheat (ha)		**0.029*****	**0.024*****	**0.031*****
		(0.001)	(0.001)	(0.001)
Pasture/hay (ha)			**0.002*****	**0.002*****
			(0.000)	(0.000)
Corn (ha)				**0.011*****
				(0.001)
(Intercept)	**252.906*****	**337.354*****	**230.416*****	**244.805*****
	(18.377)	(17.971)	(19.054)	(19.092)
R^2	0.140	0.192	0.202	0.205
Dep. var. mean	928.7	928.7	928.7	928.7
Num. obs.	19718	19718	19718	19718
RMSE	1960.445	1900.347	1888.061	1884.575

Notes: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. Table shows results for regression of dicamba use (in kg) at the county level against agricultural land uses to obtain in years 2008-2014 to obtain typical per-hectare dicamba use (kg/ha). The models in columns (2)–(4) further add crop acreage for two land uses that are major receivers of dicamba according to USDA statistics, namely wheat (column 2), hay/pasture (column 3), and corn (column 4). Includes Kansas.

Table A.19: Dicamba use models – Including Kansas

care of those that resist.[2]

[2] Personal communication, August 12th, 2019.

A.3.3 Technical appendix

Units

Useful units and orders of magnitude:

- hectare (ha): by definition, a surface equivalent to 100 m x 100 m, i.e. 10^4 m^2 or about 2.47 acres. A quarter-quarter section (small field) corresponds to 16 ha.
- kilogram per hectare (kg/ha): 1 kg.ha^{-1} $\simeq$ 0.89 lb.A^{-1}. Typically dicamba is applied at a rate of 0.5 to 1 lb ae (acid equivalent.) per acre (or 11 to 22 fl oz per acre), hence about 0.56 to 1,12 kg/ha.

Determination of wind-neighbors procedure, detail

The data obtained from NOAA consists in a set of points identified by their latitude and longitude and associated at each time period with a value of u-wind and v-wind (in meters per second). To produce the one-to-many mapping of focal counties with their up-, down- and cross-wind neighbors, I proceeded as follows.

First, I projected the vectorial sum of the $\vec{U}$ and $\vec{V}$ vectors from the original latitude and longitude: this gives me the latitude and longitude of the place where a particle would land if it was lifted off by the wind at the origin location. This requires converting distances originally expressed in meters (the units of the wind data) to degrees of latitude and longitude.

The same procedure is repeated for $-\left(\vec{U}+\vec{V}\right)$ (upwind), $\vec{U}-\vec{V}$ and $-\vec{U}+\vec{V}$ (crosswind).

It then suffices to find in which counties (if any) the origin point and the up-, down-, and cross-wind destinations are located. Inevitably, some of these will fall into the ocean, abroad or in the same county as the origin point. These instances are removed, and the rest of the destinations are combined, and associated with their corresponding origin county to obtain the set of counties located up-, down-, and cross-wind from a focal county.

Chapter 2

Agricultural landscape complexity and pest pressure: A Less simplistic approach to landscape complexity

with Eyal G. Frank

"Flying to Minneapolis from the West, you see it as a theological problem.

The great flat farms of Minnesota are laid out in a ruled grid, as empty of surprises as a sheet of graph paper. Every gravelled path, every ditch has been projected along the latitude and longitude lines of the township-and-range-survey system. The farms are square, the fields are square, the houses are square; if you could pluck their roofs off from over people's heads, you'd see families sitting at square tables in the dead centre of square rooms. Nature has been stripped, shaven, drilled, punished and repressed in this right-angled, right-thinking Lutheran country. It makes you ache for the sight of a rebellious curve or the irregular, dappled colour of a field where a careless farmer has allowed corn and soybeans to cohabit.

But there are no careless farmers on this flight path. The landscape is open to your inspection – as to God's – as an enormous advertisement for the awful rectitude of the people. There are no funny goings-on down here, it says; we are plain upright folk, fit candidates for heaven.

Then the river enters the picture – a broad serpentine shadow that sprawls unconformably across the checkerboard. Deviously winding, riddled with black sloughs and green cigar-shaped islands, the Mississippi looks as if it had been put here to teach the god-fearing Midwest a lesson about stubborn and unregenerate nature. Like John Calvin's bad temper, it presents itself as the wild beast in the heart of the heartland."

Jonathan Raban, "Mississippi Water" (excerpt, *In*: *Granta*, issue #45, 1993)

Abstract

The use of insecticides in agriculture raises public health concerns. Evidence from landscape ecology suggests that insect abundance might be lower in complex landscapes, i.e. with diverse landcovers (composition) and irregular arrangements in space (structure). Thus increasing complexity could enable a better control of agricultural pests while reducing the quantities of insecticides used, but there is no evidence that landscape complexity affects pest pressure at the relevant scale and in the relevant context, nor is the magnitude of the putative effect known. We explore these questions to assess the relevance of landscape complexity in pest management in the case of the conterminous United States over 2008-2017. We use spatially-explicit landcover data to quantify complexity, and proxy pest pressure with insecticide use. Of the several measures of landscape complexity, we a few that are consistently associated with insecticide use, and thus potentially to insect pressure. These findings offer promising venues in future explorations of the ways in which the structure of an agricultural landscape affects pest pressure, and how to leverage them to minimize it.

Keywords: agriculture, pests, spatial externalities, land use, landscape ecology.
JEL Classification: Q5, Q15, Q16.

Introduction

A British geographer once remarked that,

> "In some quarters it has become axiomatic to regard subdivision in general, and irrespective of the crop grown, as the blackest of evils, to be prevented by legislative action as one would attempt to prevent prostitution or blackmail" (Farmer, 1960)

Indeed in the twentieth century notions of economies of scale, but also the recent advent of engine-powered farm machinery, made the persistence of small, misshapen fields seem like a nuisance to agricultural productivity, to economic efficiency, to progress. This gave rise to land consolidation policies that would not only reduce the scattering of parcels, but also increase their size and make their shape more amenable to mechanization (square) (Bentley, 1987). By design, these policies fuelled a trend of landscape simplification around the world. Some started as early as the late 19th century (Great Britain), but consolidation truly became a fixture of national agricultural policy and international development programs after the Second World War – see for instance Jacoby (1959) in Europe, Robinson and Sutherland (2002) for Great Britain, and Deo & Swanson (1990) for the developing world. James C. Scott, while contesting the existence of such productivity advantages in the United States, noted that larger farms eventually had the upper hand thanks to their facilitated and heightened access to credit (Scott, 1998, p. 197).

Regardless of the debates on the economic merits of the consolidation policies and whether they met expectations (see also Bentley (1987)), agricultural landscapes underwent several decades of consolidation and simplification, and simplification and consolidation are still promoted as a necessary first step towards agricultural development and modernization.

Absent from these discussions were the effects of landscape simplification on the ecological processes at work in agricultural systems. Recent development in landscape ecology have ascertained the importance of habitat configuration for certain species (e.g., Jonsen and Fahrig (1997); Olimpi et al. (2020)), and consequently purported that the simplification of agricultural landscape was liable to causing higher levels of agricultural pest pressure.

The reasoning is the following. Large fields constitute large habitat patches for crop pests, increasing their food supply and the carrying capacity of landscape for them, while making them harder to hunt for their "natural enemies", which may in addition be species that require the proximity of non-cropped habitat and/or are edge-lovers (see summary schematic in Figure 2.2). From their point of view smaller patches make the search for prey easier, intricate patch shapes provide more edge (for a given patch area) to dwell in, and the interspersion of hunting (crop) patches with habitat (non-crop) patches similarly increases the number of individuals a landscape could support.

If ascertained, that would be an heretoforth unrecognized cost to simplification, borne individually by each farmer (insecticide expenditures), and possibly inhabitants (health costs) as a result of the behavior of farmers as a whole.

Intriguingly, while neither the effect itself nor its magnitude have been established, the reasoning detailed above has been the rationale underlying some of the recent environmental policies in Europe attempting to reverse decades of landscape simplification in agriculture, in particular with the agri-environmental schemes supported by the Common Agricultural Policy of the European Union (Tscharntke et al., 2005).

There is a consensus in the ecological literature that heterogeneity, and in particular spatial heterogeneity, is an important factor in population dynamics, species interactions, and the ensuing ecosystem-level processes and properties (Wiens, 2000). But what features matter, at what scale, for which species, and whether that translates specifically into measurable effects in agricultural landscapes, remains an open question.

The field of landscape ecology is notoriously refractory to experiment because of obvious practical and ethical difficulties (Kareiva, 1990; McGarigal and Cushman, 2002). There is thus a dearth of experimental studies, and those that exist are often limited by confounding factors, external validity issues (because of the scale at, or system in which they've been conducted), in addition to falling short of measuring landscape features in a spatially-explicit way. Nonetheless of note are Kareiva (1987); Martin (2013), the former manipulating habitat fragmentation to observe its effects on ladybird-aphid interactions, the latter taking landscape complexity as a given and manipulating components of a trophic network.[1]

Observational studies have been few, and have to our knowledge failed so far to account for the spatial nature of habitats. Most ambitious among these were Larsen (2013); Larsen et al. (2015); Larsen and Noack (2017), and Meehan et al. (2011); Meehan and Gratton (2015, 2016), the two groups coming to opposite conclusions with increasing nuance over time. While none addresses the spatial nature of agricultural systems, all these studies tackle the question of complexity and pest pressure in agricultural systems at the scale relevant the human management of agricultural landscapes.

This is the approach we seek to emulate, while recognizing in addition the spatial nature of ecological interactions.

In this paper, we leverage spatially-explicit land cover data in a panel study to examine the relationship between landscape complexity and pest pressure on agricultural land, proxied by insecticide use.

Our study stands out as the first to take landscape complexity as seriously as its theoreticians did, i.e. as a spatial phenomenon, and by its critical use of metrics (FRAGSTATS metrics suite) and data (the Cropland Data Layer) well-suited for that purpose, in combination.

In doing so, it overcomes the dearth of evidence at a management-relevant scale of effects of compositional and structural complexity of agricultural landscape. Being conducted on a continental scale (the conterminous United States), it also aims at providing more general insights spanning a variety of agricultural contexts and climatic zones.

Finally, and from a policymaking perspective, it provides evidence of the effects the simplification trend of the past few decades on pest pressure and ecosystem-mediated pest suppression that informs future policies whether favoring further simplification or incentivizing re-complexification.

We pinpoint dimensions of landscape complexity that covary with pesticide use, and thus possibly matter for the control of pest pressure via the arrangement in space of landscapes. Two of these dimensions are not identifiable through the use of non-spatial methods.
The rest of the paper is organized as follows: Section 2.1 provides a brief overview of the ecological theory and evidence pertaining to the role of landscape factors in pest abundance in agriculture, Sec-

[1] Jenerette and Shen (2012) provide a useful review of experimental studies landscape ecology.

tion 2.2 describes the data sources, Section 2.3 details the empirical strategy, the results are presented and analyzed in Section 2.4, and finally, Section 2.5 outlines areas for future research and concludes.

2.1 Background: Landscape complexity influences pest pressure, in theory

Wiens (1976) is touted as "the beginning of a landscape approach to population dynamics" (Turner, 2005). It indeed summarizes a considerable volume of evidence spanning several realms and systems, establishing to the importance of "discontinuities in time and space" for population dynamics, whether as single populations, interacting species, or whole ecosystems. The availability of multiple habitat patches might be key to maintaining a population, while another will thrive better in a single continuous habitat patch of the same area, or in a situation where the edge-to-area ratio is maximized (edge-lovers) or minimized (edge-avoiders); the way heterogeneity operates are detailed in (Wiens, 2000). Thus landscape characteristics such as the number and size of habitat patches, their shape (edge effects), and arrangement in space (interspersion) typically matters ecologically.

But while theoretical, observational, and experimental studies point to the existence of an effect of landscape fragmentation on population dynamics and species diversity (Kareiva, 1990; Bender et al., 1998),[2] implications for the management of agricultural landscapes and the control of crop-damaging insect populations are still wanting Fahrig et al. (2011). The reaction of pest species and their "natural enemies" (e.g. predators) to spatial heterogeneity is indeed poorly known (but see Kareiva (1987)). The view that complex landscapes favor pest control is nonetheless widely held.

Importantly, none of the studies we could find addressed the spatial nature of the landscape complexity-pest pressure relationship. Experimental studies (Chaplin-Kramer et al., 2011; Rusch et al., 2016; Karp et al., 2018, e.g.,) typically consider the amount of "natural habitat" in the landscape, hence no consideration of landscape structure, and a very coarse measure of landscape composition. Recently, two sets of studies have sought to overcome the limitations in terms of extent and external validity of the experimental studies by taking an observational approach at a continental scale (Larsen, 2013; Larsen et al., 2015; Larsen and Noack, 2017; Meehan et al., 2011; Meehan and Gratton, 2015, 2016) with conflicting results, yet without addressing the spatial nature of agricultural landscapes. This is problematic because, as illustrated in Figure 2.1, all these approaches amount to equating a variety of vastly different situations.

[2] Bender et al. (1998) conducts a meta-analysis of the 25 existing suitable studies evaluating an effect of patch size on population trends, differentiating the effects across edge-avoider, edge, and generalist species of birds, mammals, and insects. Almost trivially, they find that edge-avoiders suffer from patch reduction, edge species thrive in small patches, and generalists are not affected by patch size. 9 of the 12 insect species surveyed were interior (edge-avoider) species (the remaining 3 were generalists), while birds (113 species) were evenly spread across the three categories.

Figure 2.1: What we mean by simplification of the agricultural landscape

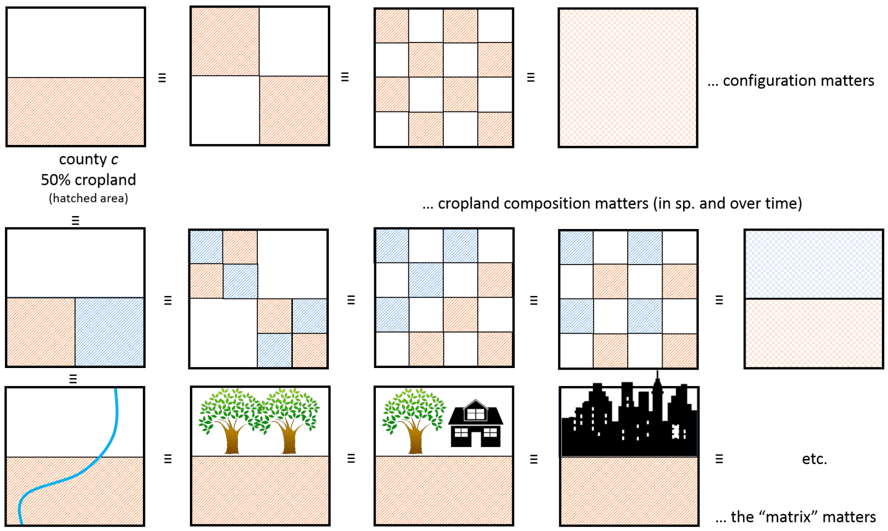

Notes: Colors orange and blue indicate cropland (represent two types of crops), no color indicates other uses ("the matrix"). Previous work studying the effects of "landscape simplification" on pesticide use (Meehan et al., 2011; Larsen, 2013; Larsen et al., 2015; Meehan and Gratton, 2016; Larsen and Noack, 2017) based their analyses on the share of cropland in any given county as their measure of landscape complexity. This would amount to considering all the above situations as equivalent, which has been refuted by research in landscape ecology.

The hypothesis we set out to test in this study is the following (illustrated schematically in Figure 2.2): agricultural pest populations are hindered by complex landscapes, because they exhibit edge avoiding behaviors and benefit from extensive stores of food sources (large monoculture fields) and because on the contrary their predators avoid open areas and benefit from the availability of edges to hide and hunt (i.e. benefit from complex landscapes). Implicit in that thesis statement is that not only the composition, but also the spatial arrangement of the landscape matters, and our approach explicitly addresses that.

Figure 2.2: Hypothesis

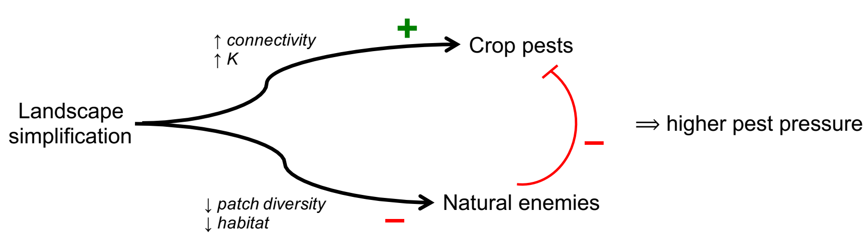

Notes: Natural enemies are the species that predate on crop pests (e.g. other insects, birds). Crop pests here are herbivorous insects feeding on plants of agricultural interest (e.g. aphids). "+"s indicate a positive effect, "–"s indicate a negative effect.

2.2 Data

This section describes the data sources used, which chiefly pertain to spatially-explicit land use maps and their processing, and to pesticide use as a proxy for pest pressure.

2.2.1 Spatial distribution of land cover

This paper hinges on the spatial nature of ecological processes and in particular the configuration of species' habitat in space. A panel of spatially-explicit land cover surveyed over an extensive period of time is therefore a key component of the analysis. We use the Cropland Data Layer (CDL) in the years when it spans the entirety of the conterminous United States (2008-2018). It is derived from remotely sensed data classified into 135 land covers by the National Agricultural Statistics Service (NASS) at the U.S. Department of Agriculture (USDA) (Boryan et al., 2011; Han et al., 2012), 77 of which are agricultural (excluding pasture); it has a ground resolution of 30 meters, which is more than enough to delineate individual fields, but is still larger than the typical scale used in landscape structure experiments (e.g., see Kareiva (1987)), and too coarse to detect, for instance, such agri-environmental features as hedgerows.

Words of caution as to the usage of remotely-sensed land cover, and in particular its thematic classification, apply (Lark et al., 2017). While common, widespread crops are exceedingly well identified by the classification algorithms (e.g., cotton, soybean, wheat have user accuracy above 90%[3]), it is a well-known fact that the CDL performs poorly on others, notably those akin to pasture and hay farming (Lark et al., 2017; Scott, 2013). In the aggregate, county-level estimates of crop-specific acreage obtained with the CDL track closely those obtained independently by the USDA using surveys (see for instance Figure A.11 in the Appendix to Chapter 1). However this study is concerned with arrangement in space (rather than, for example, transition between land covers in a given pixel over time), and therefore is exposed to a particular risk of measurement error: the misclassification of a pixel can introduce artificial complexity in otherwise "simple" shapes (e.g. a wheat pixel within a wheat field, misclassified as barley, adds a new patch, adds edge to the wheat patch, and increases the complexity of the wheat shape (non convex), and potentially alters substantially the typical distance between barley patches).

Nonetheless we deem the classification error in the CDL is plausibly independent from the phenomena studied here, and devise strategies to circumvent or alleviate the most typical sources of classification errors.

2.2.2 From rasterized land cover to landscape composition and structure: the FRAGSTATS metrics

The FRAGSTATS metrics suite (McGarigal et al., 2002) was developed to quantify elements of landscape structure and composition. It comprises 38 metrics characterizing various dimensions of land cover composition (e.g., land cover richness, number of patches, Shannon diversity...[4]) and structure

[3]Accuracy matrices by year and state can be found here: www.nass.usda.gov.

[4]Said otherwise: number of distinct land cover class (R), number of contiguous patches of any given landcover (fields), Shannon's diversity index: $H' = -\sum_{i=1}^{R} p_i \ln p_i$ with p_i the proportion occupied by land cover type i in the

at the patch (e.g., perimeter-to-area ratio, fractal dimension, ...) and at the class (e.g., aggregation index, contagion index, ...) levels.

This is an innovation from previous studies as this set of metrics enables to take into account the very spatial nature of agricultural landscapes (the shape of the patches/fields and how fields of a similar crop are arranged in space).

Some of these metrics are redundant, whether by construction or empirically (see Figure B.1 and Figure B.2); to avoid using metrics that are correlated, we propose to follow Cushman et al. (2008) and identify the non-redundant dimensions of the landscape with a principal component analysis to use only these in our analyses (whether we use the eigenvectors, independent by design, or the most important principal components, easier to interpret).

The `SDMTools` package[5] in R (in particular the functions `PatchStat` and `ClassStat`) enables to apply those metrics on spatially-explicit land cover data, in our case, on the CDL at the county-year level.

2.2.3 Pesticide use

Similar to 1.2.2 in the previous chapter, pesticide use is obtained from the US Geological Survey (USGS). Its Pesticide National Synthesis Project collects estimates of pesticide use in agriculture for the conterminous United States by compound (i.e. chemical species, for instance chlorpyrifos, atrazine, 2,4-D, etc.), in kilograms, at the county-year level since 1992. These are produced by combining proprietary pesticide-by-crop use data collected every year throughout the country with crop acreage (Baker and Stone, 2015; Thelin and Stone, 2013). 525 compounds are assessed, among which some are used as herbicides (weed-killers), others as insecticides, yet others as fungicides, or other (nematicide, etc.). Each compound is associated to its main use following the same procedure as in Frank (2017).

2.2.4 Agricultural census

Data for the replication of previous studies are taken from the U.S. agricultural censuses of 1987, 2002, 2007, 2012, 2017. These are specifically: acres of cropland, harvested acres, acres treated with insecticide, all at the county level.

2.2.5 Weather controls

Identical to chapter 1, temperature and precipitation variables at the county level are used as controls in the pesticide use regressions, as wetter or hotter years might influence the prevalence of pests. They are derived from PRISM data (PRISM Climate Group at Oregon State University, 2019) and processed as done in Schlenker and Roberts (2009).[6] These include, at the county and year level: the minimum temperature, the maximum temperature, the average temperature, degree days above 10°C, degree days above 29°C, and total precipitation in millimeters from 1950 to 2017.

landscape (county).

[5] rdocumentation.org/packages/SDMTools.

[6] http://prism.oregonstate.edu. I am grateful to Prof. Schlenker for providing the county-level data.

2.3 Empirical strategy

2.3.1 Naive approach

The general approach followed in this paper is to exploit panel variation over the 2008-2017 period, and estimate:

$$y_{it} = \beta_0 + \beta_1 \text{Landscape\_metric}_{it} + f(X_{it}) + \varepsilon_{it} \tag{2.1}$$

with y_{it} pesticide use at the county i year t level, Landscape_metric$_{it}$ a single or composite measure of landscape complexity (see this section below), X_{it} controls (weather, county and year fixed-effects), ε_{it} errors clustered at the county level.

Refinements on the strategy, and in particular concerns for identification, are discussed in 2.5.

As explained in 2.2, the choice of Landscape_metric$_{it}$ is not trivial nor innocuous; of particular concern is the correlation between landscape metrics. The following sections detail two procedures.

2.3.2 Selecting landscape metrics

A solution to the problem of correlation between metrics is to test for correlation and pick or eliminate on that basis.

This is done in Figure B.1. Once aggregated at the county level, few metrics appear prohibitively correlated (compare before aggregation over land cover classes, Figure B.2), thus only a few metrics are discarded in the "natural" regressions presented below in 2.4.2. 10 county-level variables are dropped, 10 remain: Landcover richness, Shannon diversity index, Avg. patch area, Avg. landscape shape index, Avg. perimeter-to-area ratio, Avg. patch cohesion index, Maximum patch area,[7] % agriculture, % wheat, % fruits and vegetables (of total county area) – see appendix section B.3.1 for more details on the metrics.

The advantage of such an approach is that the components of landscape complexity that do matter for pest pressure can be identified in a legible manner. The obvious drawback is the potential loss of information and the risk of keeping redundant information.

2.3.3 Operating on principal components

A more systematic solution to the problem of correlation is to follow Cushman et al. (2008) and run a principal components analysis (PCA) on the landscape metrics, extract the orthogonal dimensions, and run the regression on the first few principal components.

The drawback of that approach is that principal components do not have as tangible and legible a meaning as the raw metrics; in regressions, interpreting signs is impossible (they are arbitrary). The obvious advantage over the previous method is its systematicity and the better exploitation of the information contained in the metrics.

Taking an arbitrary threshold at 80% of the variance (Figure 2.3), we keep the first five principal components for the regressions on the rotated data (results presented in 2.4.3).

[7] Note that these statistics aggregate over all the classes present in the "landscape" (county). E.g. the Patch cohesion index is derived for every class (land cover) present in a given county in a given year, and then averaged to obtain Avg. patch cohesion index.

Figure 2.3: Scree plot

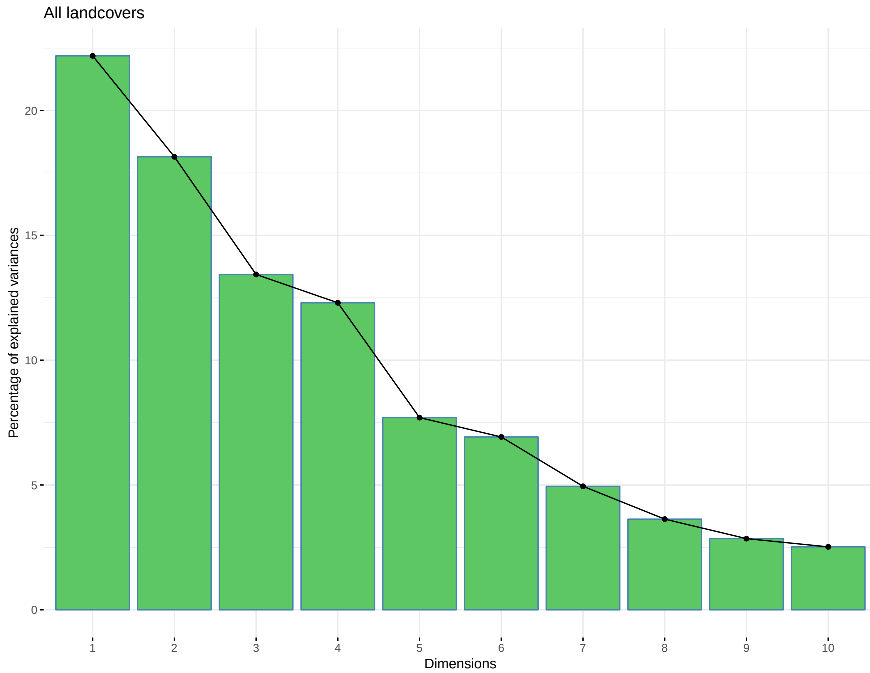

	PC1	PC2	PC3	PC4	PC5	PC6	PC7	PC8
Standard deviation	2.0532	1.8568	1.5975	1.5286	1.20943	1.14693	0.96947	0.83077
Proportion of Variance	0.2219	0.1815	0.1343	0.1230	0.07699	0.06923	0.04947	0.03633
Cumulative Proportion	0.2219	0.4033	0.5377	0.6606	0.73763	0.80686	0.85633	0.89266

Notes: Plot shows the proportions of the variance explained by each of the first ten principal components (line 2 of the table).

The strongest loadings in the rotation matrix typically serve to interpret what dimensions of the data the components capture, though interpretability is always limited, particularly so when the input variables are many (here, 19). In PC1, for instance, the strongest loadings are for the Shannon and Simpson indices (which we found earlier to be correlated, ρ=-0.95), total area (county area), maximum patch area, proportion of agriculture, which seem to point in the general direction of landscape (county) size. In PC2, the strongest loadings are for land cover richness, total edge, average proportion of (a given class) in the landscape, average perimeter-to-area ratio, which seem to point to a measure of patchiness.

The distribution of the data along those first two eigenvectors is shown in Figure 2.4, and the arrows indicate how the original variables map to PC1 and PC2.

Figure 2.4: Biplot for principal components 1 and 2

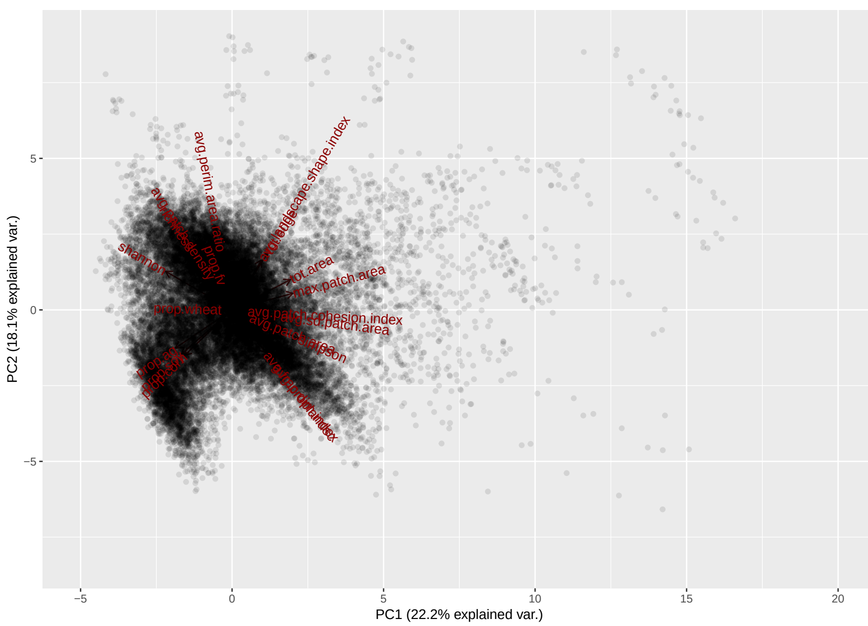

Notes: Plot shows the data reprojected along the first two eigenvectors, which together explain about 45% of the variance (see Figure 2.3).

2.4 Results

2.4.1 Replication of previous studies

We first replicate results in Meehan et al. (2011); Larsen (2013); Larsen et al. (2015) whose approach is akin to yet unlike ours. The results are consigned in Table 2.1, along with variations of our own.

	Meehan 2011	Larsen 2013	Variant 1	Larsen 2015	Variant 2	Variant 3	Variant 4
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
(Intercept)	12.22^{***}						
	(0.87)						
croplandShare	0.26^{***}	0.12	0.02	0.02			
	(0.01)	(0.08)	(0.08)	(0.08)			
prop.corn		155.07^{***}	44.01^{*}	44.01	4346.53	1.65^{***}	30.62
		(28.69)	(21.70)	(26.72)	(7633.90)	(0.46)	(971.81)
prop.soy		132.57^{***}	75.97^{***}	75.97^{***}	-16034.33	1.59^{**}	281.81
		(25.79)	(14.67)	(13.41)	(8326.96)	(0.50)	(820.33)
prop.fv		129.28	189.98^{***}	189.98^{***}	69133.70^{**}	42.51^{***}	7081.44^{**}
		(108.11)	(47.17)	(44.86)	(26479.09)	(1.58)	(2706.79)
prop.ag					-325.57	-0.69^{**}	-27.19
					(4193.32)	(0.25)	(412.45)
avg.patch.area					0.01	0.00	-0.00
					(0.03)	(0.00)	(0.00)
Dep. var	% treated	% treated	% treated	% treated	kg	kg/acre	kg/acre
Sample	Midwest	Midwest	US	US	Midwest	Midwest	US
	2007	1987-2007	1987-2007	1987-2007	2008-2017	2008-2017	2008-2017
Num. obs.	604	1216	5742	5742	5598	5598	27107
R^2 (full model)	0.35	0.89	0.89	0.89	0.76	0.71	0.15
R^2 (proj model)	0.35	0.06	0.05	0.05	0.01	0.13	0.00
Dep. var. mean	25.9	22.0	24.5	24.5	7326	0.09	23.0

Notes: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. Table shows regression results for replications and variants of specifications used in Meehan et al. (2011) (column (1)), Larsen (2013) (column (2)), idem for the entire U.S. (column (3)), idem clustering at the agricultural district level (Larsen et al. (2015), column (4)), using the quantity of insecticides (in kilograms) as opposed to the % of harvested acres treated with insecticides (column (5)), the quantity of insecticides per acre (column (6)), with the full sample of U.S. counties (column (7)). All specifications but (1) include county and year fixed effects, and weather controls.

Table 2.1: Replications

Variables representing the prevalence of non-agricultural habitat is equated to complexity in those approaches (except in (1), the only use of a metric addressing landscape structure), and their importance and sign varies across specifications. The only result that survives all variations on the specifications and choice of the sample and dependent variables is the relationship between the proportion of orchards and vegetables in the county and the use of insecticides: the more land fruits and vegetables occupy in a county, the more insecticides used.

2.4.2 Regressions on untransformed metrics

The results of the estimation of Equation (2.1) with raw FRAGSTATS metrics as $\text{Landscape\_metric}_{it}$ are presented in Table 2.2 (alternative estimations are reported in appendix Table B.1–Table B.3, with little change).

	Insecticides (1)	Herbicides (2)	Fungicides (3)
Richness	632.63	-7111.74^{***}	-788.31
	(1258.75)	(1578.53)	(1908.56)
Shannon	-3036.27	$4799.30^{\cdot}$	-2956.71
	(2277.69)	(2856.34)	(3453.52)
Avg. patch area	-256.68	181.28	-186.16
	(588.86)	(738.46)	(892.85)
Avg. landscape shape index	-2206.18	-12648.91^{***}	-5583.71^{**}
	(1421.18)	(1782.23)	(2154.84)
Avg. perimeter-to-area ratio	757.58	385.59	369.46
	(473.02)	(593.19)	(717.21)
Avg. patch cohesion index	1012.76^{*}	1704.52^{**}	$1307.83^{\cdot}$
	(478.64)	(600.24)	(725.73)
Max. patch area	-7376.53^{*}	-15723.05^{***}	-15378.29^{**}
	(3512.65)	(4405.04)	(5326.00)
Prop. agriculture	114.54	34908.98^{***}	7492.53^{*}
	(2360.19)	(2959.81)	(3578.61)
Prop. wheat	-3563.30^{*}	-38911.62^{***}	-11575.15^{***}
	(1698.39)	(2129.87)	(2575.16)
Prop. fruits and vegetables	8756.64^{***}	1231.45	4805.08^{**}
	(1109.11)	(1390.88)	(1681.67)
Num. obs.	27,109	27,109	27,109
R^2 (full model)	0.97	0.91	0.96
R^2 (proj model)	0.01	0.03	0.01
Dep. var. mean	28,432	101,300	43,113

Notes: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$, $\cdot p < 0.1$. Table shows results for regression of untransformed landscape metrics at the county level against pesticide – insecticide (column (1)), herbicide (column (2)), fungicide (column (3)) – use (kg) for all U.S. counties over 2008-2017. The models all include county and year fixed effects, and weather controls. For comparability, independent variables have been standardized (centered, unit standard deviation).

Table 2.2: Pesticide use models on selected landscape metrics

The column of interest for the purposes of insect population control is column (1) of Table 2.2, where the dependent variable is insecticide use. Two structural metrics are significantly associated with insecticide use.

An increase in maximum patch area is associated with a decrease in insecticide use. There are two possible interpretations for this. Admitting that insecticide use proxies for insect abundance, this means that large patches are detrimental to insect pests, whether because the proximity of other habitat types benefits them,[8] or because their natural enemies are more efficient in large patches (e.g. see Kareiva (1987) – he finds that ladybird beetles exerce better control on aphid populations in non-fragmented landscapes, one explanation being that ladybirds need to aggregate around aphid clusters to control them, which is easier in continuous patches). Should that be the case, that would warrant further investigation into the ecology of insect pests and their natural enemies. On the other hand, taking pesticide use for what it is, rather than for a proxy, this could also just be a symptom of higher efficiency in insecticide (i.e. more sparing) use on larger fields. That would also be an interesting fact,

[8] E.g. hedgerows have been reported to serve as refuge to some insect populations.

but with drastically different implications.

Average patch cohesion index (see B.3.1) is positively associated with insecticide use. In that case the farm management explanation seems less likely. Interpreted as a proxy for insect populations, then, this would indicate that insect pests are better able to proliferate in fragmented landscapes where habitat patches are aggregated in space (as opposed to scattered), which would be consistent with metapopulation theory. Should that be further ascertained, this would mean that there is a benefit to scattering fields of the same crop across the landscape, as far apart from each other as possible, said otherwise, aim for maximal interspersion.

Unsurprisingly, insecticide use is positively associated with the proportion of land devoted to fruits and vegetables in a county (pesticide-intensive, high value-added crops), and is negatively associated with the proportion of wheat.

Columns (2)-(3) of Table 2.2 indicate that the same management or ecological mechanisms are at work for herbicides (weeds) and fungicides (fungi): the signs are identical and most of the variables of importance in (1) are conserved in (2)-(3) (but note the relationship between richness and herbicides – possibly a management effect).

2.4.3 Regressions on rotated landscape metrics

The results of the estimation of Equation (2.1) with the first five principal components of the FRAGSTATS metrics as Landscape_metric$_{it}$ (PC1, PC2, PC3, PC4, PC5) are presented in Table 2.3 (see alternative estimations in Table B.4–Table B.6, with little change).

	Insecticides	Herbicides	Fungicides
	(1)	(2)	(3)
PC1	1655.34	-21746.61^{***}	-2872.69
	(1552.25)	(1954.73)	(2351.81)
PC2	1311.10	-13299.74^{***}	-1973.44
	(1039.31)	(1308.79)	(1574.65)
PC3	-555.37	-7909.05^{***}	-1170.66
	(1341.82)	(1689.74)	(2032.99)
PC4	318.22	-600.00	285.79
	(541.08)	(681.38)	(819.80)
PC5	-1136.60	6048.81^{***}	-150.33
	(735.32)	(925.99)	(1114.09)
Num. obs.	27109	27109	27109
R^2 (full model)	0.97	0.91	0.96
R^2 (proj model)	0.00	0.02	0.01
Dep. var. mean	28432.2	101300.2	43112.75

Notes: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$, $\cdot p < 0.1$. Table shows results for regression of rotated landscape metrics at the county level against pesticide – insecticide (column (1)), herbicide (column (2)), fungicide (column (3)) – use (kg) for all U.S. counties over 2008-2017. The models all include county and year fixed effects, and weather controls.

Table 2.3: Pesticide use models on rotated landscape metrics

To our surprise, none of the first five principal components exhibits any statistical association with insecticide use (insect population size). On the other hand, PCA is not a guided procedure; it best summarizes the data along orthogonal axes in an agnostic way, and the dimensions of landscape composition and structure that matter for insecticide use/pressure might not be those that are the most salient. Considering the full set of principal components extracted (19), we see some PCs being associated with insecticide use (Table B.7). Note that the coefficients obtained are very unstable to the addition of new covariates, except those that are significant.[9]

On the contrary, herbicide use (column (2) of Table 2.3) is exceedingly well associated with the main principal components of our landscape structure and composition data. Similarly to what has been noted above in 2.4.2, it might be because the landscape metrics chosen (FRAGSTATS), applied on the particular land use data at hand (and its spatial and thematic resolutions), is particularly well suited to capture dimensions relevant to weed management by farmers, or dimensions of population dynamics in weeds.

[9]Observed with running the regressions with PC 1–17, not shown.

2.5 Conclusions, and implications for future research

Overall, the results presented in section 2.4 above support the existence of landscape effects on pest pressure, both aspatial (e.g. proportion of vegetable and fruit crops[10]) and spatial (patch cohesion). To our knowledge, this study is the first to consider the effects of landscape complexity on pest pressure, and taking full notice of the spatial dimension of landscape complexity.

Caution is needed in the interpretation of the results, and further research is needed both to ascertain the direction of the causality, explore the underlying mechanisms, and disentangle proxy (pesticide use) from objet of interest (insect populations).

Nonetheless should these findings be confirmed, they open the door to the evaluation of the tradeoff between the operational advantages of landscape simplification and its ecological consequences.

Future research will in particular seek to address the four following limitations to the present study: causality, land cover data adequacy, pesticides as a proxy, short panel.

Indeed the relationships exposed in 2.4 derive from a panel study; while the effects of common shocks and individual county characteristics are taken care of by year and county fixed effects, other unobservables could be the cause of the relationships observed. In particular, a threat to the landscape-causes-pest-population narrative underlying this study is that landscape configuration and composition on the one hand, and pesticide use, on the other hand, may be generated by the same process of agricultural production industrialisation. Indeed, the same economies of scale that motivate the consolidation of parcels into larger, square plots, and the focus on monocultures, could also motivate farmers to apply consistently more (or less) pesticides than otherwise, possibly because it is safer and cheaper to apply them regardless of infestation status (or conversely, smaller fields might be operated less efficiently and sparingly). A solution would be to find locations and events that lead to the realisation of one but not the other dimensions of agricultural industrialisation. The anomalous persistence in places of the metes-and-bounds land demarcation regime alongside the rectangular system could for instance be utilized: it leads the coexistence of a more rugged structure of land plots and surrounding land ordained in a regular checkerboard patter of square fields (Libecap and Lueck, 2011). Thus this setting provides the variation of patch- and landscape-level complexity while keeping other aspects of agricultural intensification unchanged.

The resolution of the land cover data used to measure landscape complexity could hamper the detection of effects of complexity on insect populations (even if they truly exist) in two ways. First, the spatial resolution of the CDL (30 m x 30 m) is fairly coarse compared to the scale at which patch size, edge or fragmentation are typically observed (e.g., 1 m in Kareiva (1987); consider also the width of a hedgerow compared to a 30 m-wide pixel). If the landscape features that matter to insect populations and their natural enemies' operate at a scale below 30 m, they cannot be detected. Second, the resolution of the thematic classification might be irrelevant to the problem at hand. The CDL classifies land cover into 135 classes. This is arguably too much. While some insects are specialists of a single plant (land cover), others are generalists and might consider contiguous patches of wheat, barley, rye as one single large patch. Using the native resolution might thus be ill-suited to study the complexity that matters to insect pest species. Compounding the previous limitations, classification error leads to measurement error, which could be substantial when measuring shapes (e.g. consider

[10]Though as pointed out above, that could also be a pure farm management effect.

Figure B.5). A partial workaround could involve the reclassification of the CDL into fewer land cover classes.

Using pesticides as a proxy for pest pressure introduces obvious concerns over the human factor in pesticide application. The solution adopted by Meehan et al. (2011) was to complement the analysis with direct measurement of aphid pressure (the spatial extent of which was limited, and the data collection period seems to be over at the moment). Finding other sources of insect population data would be ideal, leveraging differences between insecticide compounds could be an interesting workaround.

Finally, the short time frame at hand (2008-2017) does not provide much temporal variation (the delineation of fields is largely fixed). The analysis of older satellite images or aerial photographs could enable the extension of the series further into the past. Displacing the analysis to a place that underwent more change recently (e.g. parts of Europe in the past few decades) and where spatially-explicit land cover maps, and pest data are available, would be another solution.

Appendix B

Appendix to Chapter 2

B.1 Additional figures

B.1.1 Correlation matrices between FRAGSTATS metrics

Figure B.1: Correlation matrix between metrics at the county level

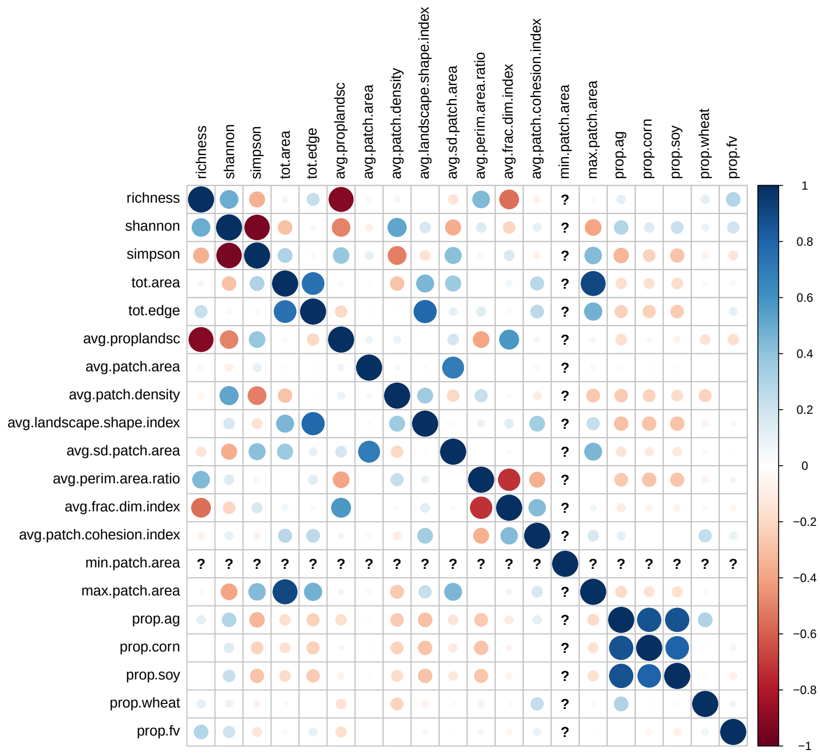

Notes: Matrix shows correlation between landscape metrics aggregated over class (land cover type) at the county level. Full sample ($N = 31,070$). Brighter colors and larger discs indicate stronger correlations (blue, positive, red, negative). Minimum patch area (min.patch.area) is constant because of the resolution of the CDL ($900\,\mathrm{m}^2$).

Figure B.2: Correlation matrix between metrics at the class

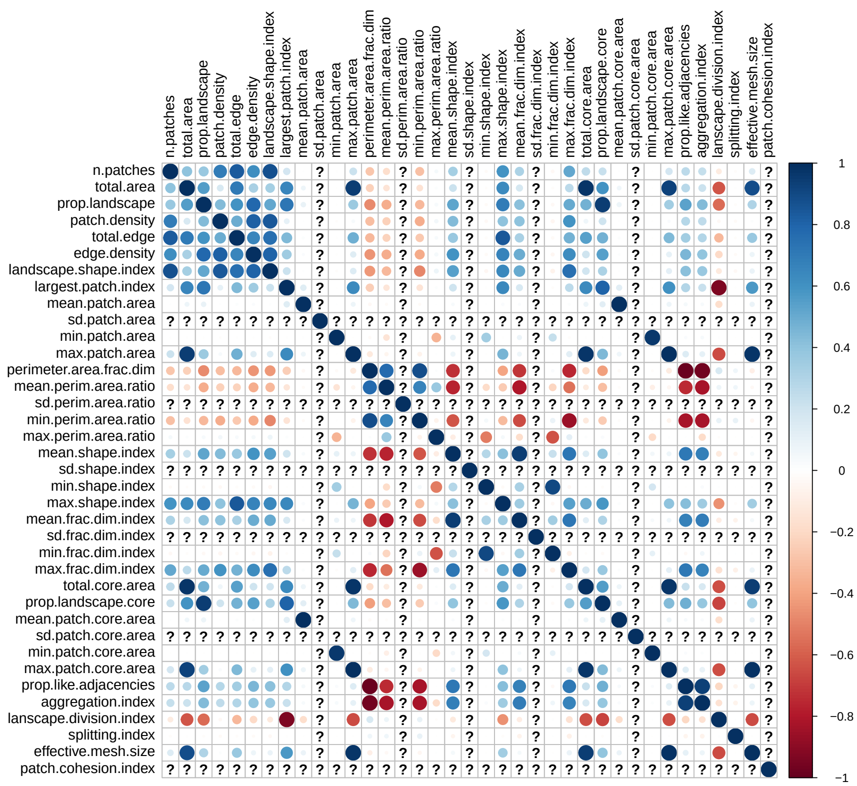

Notes: Matrix shows correlation between landscape metrics calculated at the county level. Full sample ($N = 812,876$). Brighter colors and larger discs indicate stronger correlations (blue, positive, red, negative).

B.1.2 Relationships between modestly correlated FRAGSTATS metrics

Figure B.3: The Shannon index and the average patch density (ρ=0.52)

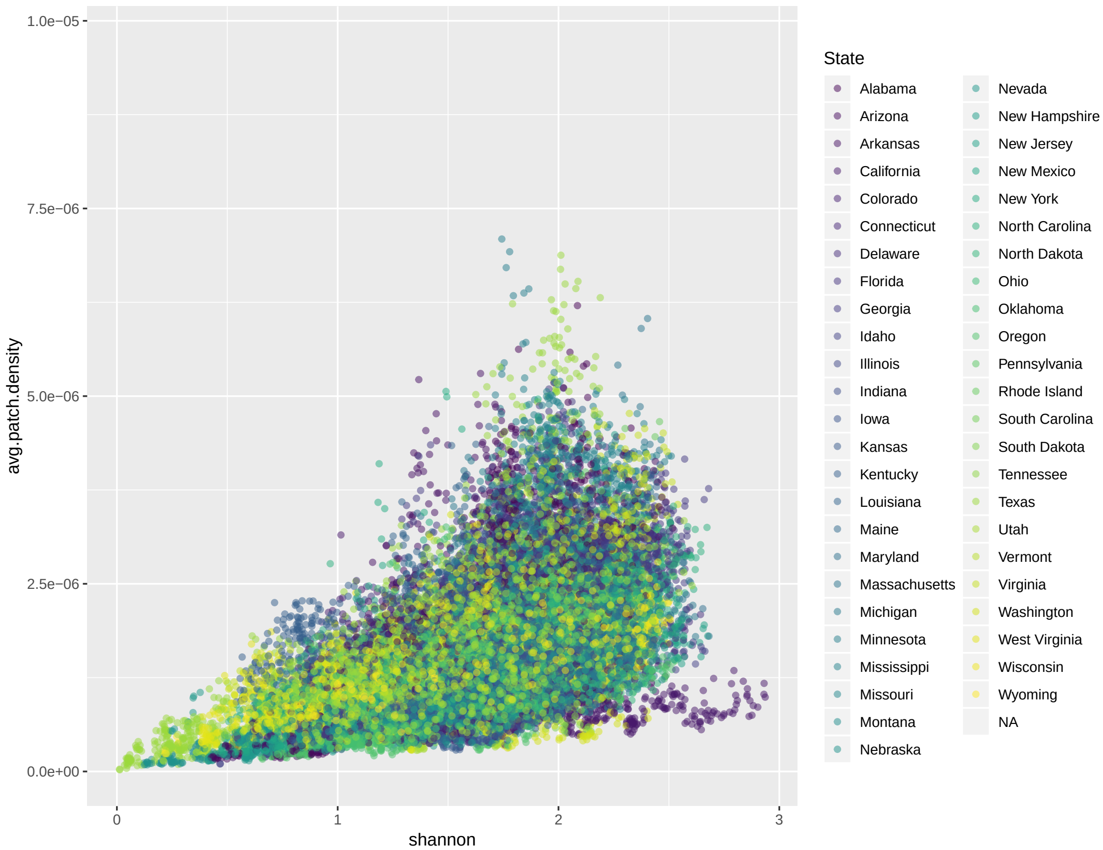

Notes: Plot shows county-year values of the Shannon index of land cover diversity plotted against average patch density and colored by state. Average patch density was omitted from the analyses.

Figure B.4: Landcover richness and the average fractal dimension index (ρ=-0.56)

Notes: Plot shows county-year values of landcover richness plotted against average fractal dimension index and colored by year. Average fractal dimension index was omitted from the analyses.

B.2 Additional tables

B.2.1 Untransformed models, alternative specifications

	Insecticides (1)	Herbicides (2)	Fungicides (3)
Richness	647.30	-8005.74^{***}	-484.38
	(1200.37)	(1506.64)	(1836.30)
Shannon	-4886.99^{*}	4229.98	-5445.61
	(2181.07)	(2737.57)	(3336.57)
Avg. patch area	-90.48	-68.42	-79.61
	(300.30)	(376.93)	(459.40)
Avg. landscape shape index	-3255.19^{*}	-16393.69^{***}	-7281.99^{***}
	(1343.66)	(1686.50)	(2055.51)
Avg. perimeter-to-area ratio	242.02	42.90	-371.37
	(453.20)	(568.83)	(693.30)
Avg. patch cohesion index	991.72^{*}	2499.75^{***}	1931.64^{**}
	(458.48)	(575.46)	(701.38)
Max. patch area	-6882.73^{*}	-14284.39^{***}	-10273.69^{*}
	(3220.21)	(4041.85)	(4926.23)
Prop. agriculture	1365.28	43154.32^{***}	9880.78^{**}
	(2310.35)	(2899.84)	(3534.35)
Prop. wheat	-4355.05^{**}	-38855.60^{***}	-11030.20^{***}
	(1670.38)	(2096.58)	(2555.32)
Prop. fruits and vegetables	14257.47^{***}	2852.87^{*}	11061.49^{***}
	(1053.60)	(1322.42)	(1611.78)
Num. obs.	30,639	30,639	30,639
R^2 (full model)	0.97	0.91	0.96
R^2 (proj model)	0.01	0.02	0.00
Dep. var. mean	27,755	101,090	42,436

Notes: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$, $\cdot p < 0.1$. Table shows results for regression of untransformed landscape metrics at the county level against pesticide – insecticide (column (1)), herbicide (column (2)), fungicide (column (3)) – use (kg) for all U.S. counties over 2008-2017. The models all include county and year fixed effects, no weather controls. For comparability, independent variables have been standardized (centered and unit standard deviation).

Table B.1: Selected landscape metrics: no weather controls

	Insecticides	Herbicides	Fungicides
	(1)	(2)	(3)
Richness	632.63	-7111.74^{***}	-788.31
	(4778.78)	(2263.40)	(4995.97)
shannon	-3036.27	4799.30	-2956.71
	(6180.77)	(6106.40)	(5556.51)
Avg. patch area	-256.68	181.28	-186.16
	(62.77)	(123.98)	(150.84)
Avg. landscape shape index	-2206.18	-12648.91^{***}	-5583.71^{**}
	(2023.07)	(3294.28)	(3497.06)
Avg. perimeter-to-area ratio	757.58	385.59	369.46
	(1114.73)	(705.96)	(1031.78)
Avg. patch cohesion index	1012.76	$1704.52^{\cdot}$	1307.83
	(937.14)	(872.55)	(1216.01)
Max. patch area	-7376.53	-15723.05^{*}	-15378.29
	(5706.55)	(7352.56)	(9328.87)
Prop. agriculture	114.54	34908.98^{**}	7492.53
	(3511.04)	(11253.40)	(5802.86)
Prop. wheat	-3563.30	-38911.62^{***}	-11575.15^{*}
	(2893.06)	(8399.00)	(5334.46)
Prop. fruits and vegetables	8756.64	1231.45	4805.08
	(5456.17)	(4711.23)	(5283.17)
Num. obs.	27,109	27,109	27,109
R^2 (full model)	0.97	0.91	0.96
R^2 (proj model)	0.01	0.03	0.01
Dep. var. mean	28,432	101,300	43,113

Notes: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$, $\cdot p < 0.1$. Table shows results for regression of untransformed landscape metrics at the county level against pesticide – insecticide (column (1)), herbicide (column (2)), fungicide (column (3)) – use (kg) for all U.S. counties over 2008-2017. The models all include county and year fixed effects, and weather controls. Standard errors are clustered at the state level (n=48). For comparability, independent variables have been standardized (centered, unit standard deviation).

Table B.2: Selected landscape metrics models: clustering at the state level

	Insecticides	Herbicides	Fungicides
	(1)	(2)	(3)
Richness	632.63	-7111.74^{*}	-788.31
	(4099.24)	(2802.47)	(4490.42)
Shannon	-3036.27	4799.30	-2956.71
	(5455.48)	(4782.36)	(5200.55)
Avg. patch area	-256.68^{***}	181.28	-186.16
	(69.74)	(121.25)	(127.42)
Avg. landscape shape index	-2206.18	-12648.91^{***}	-5583.71^{*}
	(1826.53)	(3404.94)	(2145.94)
Avg. perimeter-to-area ratio	757.58	385.59	369.46
	(935.55)	(963.59)	(992.53)
Avg. patch cohesion index	1012.76	1704.52^{*}	1307.83
	(825.37)	(746.13)	(1198.35)
Max. patch area	-7376.53	-15723.05^{*}	-15378.29^{*}
	(5494.91)	(7468.09)	(6567.98)
Prop. agriculture	114.54	34908.98^{***}	$7492.53^{\cdot}$
	(2760.05)	(5453.85)	(3840.98)
Prop. wheat	-3563.30	-38911.62^{***}	-11575.15^{*}
	(2140.68)	(6032.80)	(5325.61)
Prop. fruits and vegetables	$8756.64^{\cdot}$	1231.45	4805.08
	(4951.11)	(3596.96)	(4580.52)
Num. obs.	27,109	27,109	27,109
R^2 (full model)	0.97	0.91	0.96
R^2 (proj model)	0.01	0.03	0.01
Dep. var. mean	28,432	101,300	43,113

Notes: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$, $\cdot p < 0.1$. Table shows results for regression of untransformed landscape metrics at the county level against pesticide – insecticide (column (1)), herbicide (column (2)), fungicide (column (3)) – use (kg) for all U.S. counties over 2008-2017. The models all include county and year fixed effects, and weather controls. Standard errors are clustered at the agricultural district level (n=87). For comparability, independent variables have been standardized (centered and unit standard deviation).

Table B.3: Selected landscape metrics models: clustering at the agricultural district level

B.2.2 Rotated models, alternative specifications

	Insecticides (1)	Herbicides (2)	Fungicides (3)
PC1	1146.53	-25114.95^{***}	-2313.60
	(1446.69)	(1818.87)	(2208.20)
PC2	339.88	-15469.78^{***}	$-2734.09\cdot$
	(991.85)	(1247.02)	(1513.94)
PC3	$-2135.97\cdot$	-10264.37^{***}	-4416.35^{*}
	(1222.82)	(1537.40)	(1866.48)
PC4	757.85	-744.47	869.08
	(479.77)	(603.20)	(732.31)
PC5	$-920.25\cdot$	5046.26^{***}	-481.04
	(525.41)	(660.58)	(801.98)
Num. obs.	30639	30639	30639
R^2 (full model)	0.97	0.91	0.96
R^2 (proj model)	0.00	0.01	0.00
Dep. var. mean	27754.95	101089.6	42435.55

Notes: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$, $\cdot p < 0.1$. Table shows results for regression of rotated landscape metrics at the county level against pesticide – insecticide (column (1)), herbicide (column (2)), fungicide (column (3)) – use (kg) for all U.S. counties over 2008-2017. The models all include county and year fixed effects, but no weather controls.

Table B.4: Rotated models: no weather controls

	Insecticides (1)	Herbicides (2)	Fungicides (3)
PC1	1655.34	-21746.61^{***}	-2872.69
	(2583.83)	(6149.08)	(4188.73)
PC2	1311.10	-13299.74^{**}	-1973.44
	(1346.77)	(4089.58)	(2610.00)
PC3	-555.37	-7909.05^{**}	-1170.66
	(2238.05)	(2290.33)	(2790.46)
PC4	318.22	-600.00	285.79
	(695.18)	(1264.17)	(922.85)
PC5	-1136.60	$6048.81\cdot$	-150.33
	(736.27)	(3069.38)	(1324.12)
Num. obs.	27109	27109	27109
R^2 (full model)	0.97	0.91	0.96
R^2 (proj model)	0.00	0.02	0.01
Dep. var. mean	28432.2	101300.2	43112.75

Notes: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$, $\cdot p < 0.1$. Table shows results for regression of rotated landscape metrics at the county level against pesticide – insecticide (column (1)), herbicide (column (2)), fungicide (column (3)) – use (kg) for all U.S. counties over 2008-2017. The models all include county and year fixed effects, and weather controls. Standard errors are clustered at the state level (n=48).

Table B.5: Rotated models: clustering at the state level

	Insecticides	Herbicides	Fungicides
	(1)	(2)	(3)
PC1	1655.34	-21746.61^{***}	-2872.69
	(2251.47)	(4439.90)	(3017.62)
PC2	1311.10	-13299.74^{***}	-1973.44
	(1122.53)	(3253.65)	(1850.99)
PC3	-555.37	-7909.05^{**}	-1170.66
	(1777.51)	(2520.65)	(2557.46)
PC4	318.22	-600.00	285.79
	(557.19)	(935.86)	(746.31)
PC5	$-1136.60^{\cdot}$	6048.81^{**}	-150.33
	(682.34)	(2257.26)	(900.45)
Num. obs.	27109	27109	27109
R^2 (full model)	0.97	0.91	0.96
R^2 (proj model)	0.00	0.02	0.01
Dep. var. mean	28432.2	101300.2	43112.75

Notes: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$, $\cdot p < 0.1$. Table shows results for regression of rotated landscape metrics at the county level against pesticide – insecticide (column (1)), herbicide (column (2)), fungicide (column (3)) – use (kg) for all U.S. counties over 2008-2017. The models all include county and year fixed effects, and weather controls. Standard errors are clustered at the agricultural district level (n=87).

Table B.6: Rotated models: clustering at the agricultural district level

	Insecticides	Insecticides
	(1)	(2)
PC1	1655.34	−4519.33
	(1552.25)	(24095.41)
PC2	1311.10	−3658.54
	(1039.31)	(13792.76)
PC3	−555.37	7410.66
	(1341.82)	(29380.09)
PC4	318.22	634.83
	(541.08)	(2116.19)
PC5	−1136.60	1309.92
	(735.32)	(2640.64)
PC6		−900.38
		(13752.64)
PC7		-9667.95^{***}
		(2319.26)
PC8		1399.23
		(5443.97)
PC9		3637.75
		(18496.68)
PC10		−2828.77
		(7040.03)
PC11		8195.22
		(11931.12)
PC12		2121.68
		(12037.19)
PC13		$-7271.26\cdot$
		(3753.95)
PC14		−289.68
		(2619.01)
PC15		-18544.31^{**}
		(6563.85)
PC16		3223.71
		(4698.96)
PC17		-13128.16^{*}
		(6299.00)
PC18		−9210.53
		(13367.04)
PC19		948.86
		(57526.78)
Num. obs.	27109	27109
R^2 (full model)	0.97	0.96
R^2 (proj model)	0.00	0.01
Dep. var. mean	28432.2	28432.2

Notes: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$, $\cdot p < 0.1$. Table shows results for regression of rotated landscape metrics at the county level against insecticide use (kg) for all U.S. counties over 2008-2017, with only the first 5 PCs (column (1)) or all 19 (column (2)). The models all include county and year fixed effects, and weather controls.

Table B.7: Pesticide use models on rotated landscape metrics

B.3 Other additional material

B.3.1 Most common landscape complexity indices

Relying heavily on Aulong et al. (2005) and McGarigal et al. (2002).

Richness (R) corresponds to the number of distinct land covers found in a landscape (county).

The **Shannon index** is a measure of diversity that takes into account not only the number of items (species, land covers) but also the evenness of the community (e.g. does one species/land cover dominate, or are they all present in equal abundance?): $H' = -\sum_{i=1}^{R} p_i \ln p_i$ with p_i the proportion occupied by land cover type i in the landscape (county).

The **Simpson index** (Simpson, 1949) is a measure of diversity that also takes into both richness and the evenness (it measures "the concentration of the classification," i.e. the probability that two elements picked at random in a population belong to the same type): $\lambda = \sum_{i=1}^{R} p_i^2$ with p_i the proportion occupied by land cover type i in the landscape (county).

Patch area measures the number of contiguous pixels of a given class (times resolution, hence area). Here we use summary statistics of the distributions of patch area (across classes) at the county level, the average and the maximum.

Landscape shape index measure the total length of edge for patches of a given class, but (contrary to "total edge") adjusts for the size of the landscape.

The **perimeter-to-area ratio** of a patch provides information on the shape of habitat patches (or fields, land cover patches). Its minimum is $2/r$, that of a disc of radius r, and is maximised in narrow and long habitat strips.

The **patch cohesion index** is a measure of physical connectivity between habitat patches of a same type (same land cover). It increases when patches are more aggregated in space.

B.3.2 Classification errors

Figure B.5: Illustration of classification errors

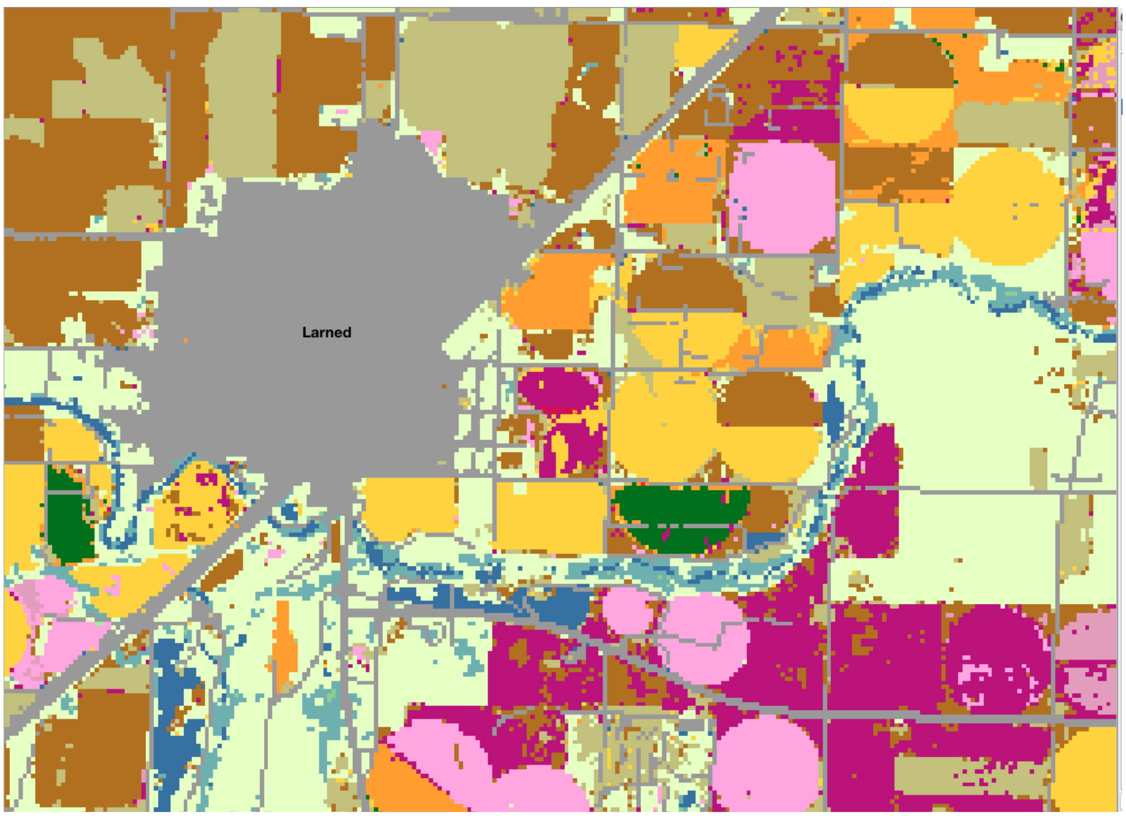

Notes: CDL visualisation (2016) in Pawnee county, KS. 1 pixel = 30 m. Of note here are the dots inside otherwise homogenous patches. While additive metrics (total land cover area, etc.) should be little disturbed by these misclassifications, they create edge, modify shapes in ways that might not be innocuous (a circle is a very simple shape, a donut is a more complex one); in addition, the errors in one class affect the others (e.g. the unreliable classification into pasture (Lark et al., 2017) could alter metrics pertaining to other, more reliably identified land covers.

Chapter 3

Payment for ecosystem services and the preservation of forest cover in Ecuador

"... for it so falls out
That what we have we prize not to the worth
Whiles we enjoy it, but being lacked and lost,
Why, then we rack the value, then we find the
The virtue that possession would not show us
While it was ours."

William Shakespeare, *Much Ado about Nothing*, VI.1 (*ca.* 1600)

Abstract

Ten years ago, in an effort to combat rampant deforestation, the payment for ecosystem service (PES) scheme Socio Bosque was launched in Ecuador: voluntary 20-year contracts between individual landowners and the government mandating bi-yearly payments conditioned to the non-deforestation of a mutually agreed-upon patch of intact forest. Here I evaluate the impact of the scheme on deforestation trends using remotely sensed forest cover, and show that the program fell short of expectations, as there was no measurable difference in deforestation in areas affected by Socio Bosque; I provide suggestive evidence that both targeting and enforcement of the program may be responsible for that outcome. Given that conservation money is in limited supply, both at the national and international level, in Ecuador as in the rest of the world, this disappointing finding is an occasion to question the design, if not the overall merit, of PES schemes, current and future, to ensure their success. This research highlights that particular attention to targeting, attractivity, monitoring and enforcement is paramount, and could still help turn things around for the second half of the program's life.

Keywords: deforestation, land set-aside, payment for ecosystem services.
JEL Classification: Q5, Q56, Q58, Q28.

Introduction

Habitat loss remains the most important threat to biodiversity (International Union for the Conservation of Nature, 2015), and land-use change accounts for just under a quarter of anthropogenic greenhouse gas emissions, in which deforestation is the major contributor (IPCC, 2014; Baccini et al., 2017); while some countries have operated their so-called forest transition (Mather, 1992; Perz and Skole, 2003), deforestation and forest degradation remain rampant in most of the developing tropics, putting the global public goods that are biodiversity and the climate at risk, in addition to creating environmental problems locally (e.g. on the water cycle, Aragão (2012)).

Payments for ecosystem services (PES) consist in incentivizing conservation actions from mostly private individuals, specifically conditioning a cash transfer to the voluntary execution of a specified set of practices. They have been an increasingly popular form of environmental policy, in particular in efforts to prevent deforestation (). Questions regarding their efficacy at obtaining results in general, and halting deforestation in particular, have been present since their emergence as a policy option (Wunder, 2005). For while their voluntary nature makes them appealing, it is also seen as PES's main weakness, as it may by design thwart the policy's additionality (effects after accounting for the baseline): wouldn't the actions have been executed absent the policy, the forest been preserved even without the payments? This concern is particularly vivid when payments are not very substantial compared to standards of living, which is often the case given the limited budget of an NGO, international organization, or ministry of the environment. Enforcement may become a supplementary concern: the myriad of contracts and practices may be impractical to track frequently. Despite being relatively easy to identify compared to other conservation practices, preservation of forest cover may still be hard to monitor (e.g. field visits in rough terrain and vast swaths of land), even with the advent of remote sensing technologies (clouds, resolution, training). As a result, success of such policies has been hard to measure, and their use is still highly favored while their efficacy has yet to be proven.

Ecuador launched its PES scheme in 2008, Socio Bosque, to incentivize small land-holders to keep forest on their land intact, and thus counter the nibbling down of forest patches, under the premise that relieving a cash constraint may reduce the urge to cut down the forest. Requirements to enter Socio Bosque were not tied to deforestation risk (though priority scores were attributed to applications that took deforestation risk into account) but instead principally on the intactness of forest cover, and payments were a function of the surface enrolled (and not to be deforested for 20 years), at a per-hectare rate decreasing in the enrolled surface (MAE, 2012). Over the course of 2008-2015, 2,525 contracts were signed, and over a million hectares were enrolled throughout the country (see Figure 3.1).

Was the scheme actually effective at reducing deforestation? Ten years into the policy, this is the question addressed here. First, observing covariates of Socio Bosque adoption at the parish (*parroquia*) level, I provide suggestive evidence that enrolment was biased towards areas that were at a low risk of deforestation. Second, using remotely sensed forest cover, I propose a quantification of the impact the policy had on deforestation trends within and outside areas affected by the program, and in addition observe and quantify breaches in compliance.

3.1 Background

The Socio Bosque program was launched nationwide in Ecuador in 2008 in an effort to combat deforestation, under the supervision of the Ministry of the Environment (Ministerio del Ambiente). The country currently contains 15.6 Mha of forest cover, and has been experiencing deforestation at around 40 kha per year over the past decade. These forests not only store carbon, but also host wildlife and plant species that make Ecuador one of the megadiverse countries in the world. The program aims at maintaining intact forest stands lying on private land, and at risk of deforestation; it implicitly identifies poverty or credit constrains as risk factors by making rural poverty reduction a secondary goal of the policy.[1] To achieve that, 20-year contracts are offered and signed between the Ministry of the Environment and voluntary individual landowners, establishing on the one hand, that the landowner will protect from deforestation and degradation a preconcerted tract of their land for at least 20 years, in exchange, on the other hand, of payments corresponding to the surface enrolled (biannual instalments in May and October). The maximum rate was $30/ha per year (decreasing in enrolled surface).

Given budget constraints and concerns over transparency, it was also important to minimise transaction costs and complexity[2] and therefore to devise simple rules of eligibility and a simple payment structure. Another challenge to these land set-aside programs, especially to those that, like Socio Bosque, aim at alleviating rural poverty, is that they require not only land ownership but also official land titles. This concern is reflected in the design of the scheme, as the marginal rate of payment decreases with the surface enrolled[3] and legal assistance is provided to applicants to help them ascertain their land titles.

[1] Also by including level of poverty as one of the three criteria for "prioritisation" – although eventually this ranking criterion was never enforced.

[2] As explicitly stated in the document establishing the program, *Manual Operativo: Unificado Proyecto Socio Bosque*, MAE (2012).

[3] In addition, special provisions apply to very small landowners (<20 ha): a single-tier payment scale, more advantageous than the first tier of the regular scale, applies. See Figure C.1 for a partial picture of land ownership in Ecuador.

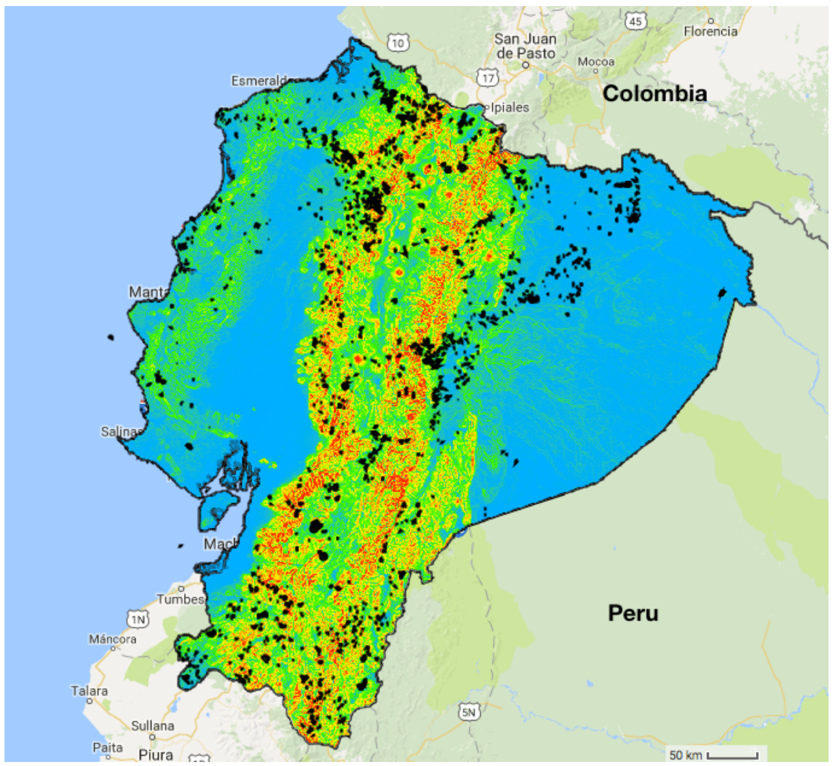

Figure 3.1: Distribution of slope gradients and individual Socio Bosque plots within Ecuador
Note: Black dots represent individual Socio Bosque plots, overlaying slope in degrees (from 0° in blue to 30.5° in red).

3.2 Methods

3.2.1 Data

I obtained Socio Bosque contract boundaries from the National System of Information of Ecuador (http://sni.gob.ec/coberturas), and forest cover loss from the Hansen Global Forest Change data set (v1.5, 2000-2017) Hansen et al. (2013), retrieved from and processed on the Google Earth Engine platform. Accessibility is from (Weiss et al., 2018), accessed and processed on the Earth Engine. Administrative boundaries were from the Database of Global Administrative Areas, GADM (gadm.org).

Table 3.1: Summary statistics on Socio Bosque

Individual contracts	
Nb. contracts	2,525
Nb. beneficiaries	2,157
Total area (ha)	164,636
Avg. contract area (ha)	65
Med. contract area (ha)	38
Max. contract area (ha)	3,364
Min. contract area (ha)	0.3
Avg. area per beneficiary (ha)	76
Med. area per beneficiary (ha)	41

"Contract" refers to entries defined by a unique contract number (unit of analysis), "beneficiary" refers to a unique first and last name combination (up to 9 contracts).

3.2.2 Analyses: adoption

I use two approaches to assess whether enrolment into Socio Bosque was spatially biased towards marginal, little deforestation-prone, land. The first approach relies on linear regressions of deforestation covariates on enrolment, at the parish level (comparing levels of adoption between all 988 parishes), and within parish (i.e. comparing Socio Bosque plot(s) characteristics with their average value in the same parish).

3.2.3 Analyses: impact

To get at the impact of the policy, it is obviously not sufficient to sum the hectares enrolled and claim they have been saved, nor is it even to compare forest cover loss before and after the roll-out of the program, for deforestation could be slowing down (or accelerating) in Ecuador as a whole, thus masking the effects of the policy itself. I used a difference-in-differences approach to compare forest cover loss between (1) parishes with and without Socio Bosque, (2) Socio Bosque plots and the remainder of the parish. The difference-in-differences method enables to compare treated, non-treated, pre-, post-policy forest cover loss, and thus estimate its effect, under the assumption that treated and non-treated units follow parallel trends (which is verified, see Figure 3.2). I estimated:

$$y_{it} = \sum_{t \neq 2007} \beta_\tau 1\{t = \tau\} 1\{i \in \mathbb{T}\} + \lambda_i + \lambda_t + \varepsilon_{it} \tag{3.1}$$

Table 3.2: Summary statistics for parishes with and without Socio Bosque enrolment

	SB parishes	Non-SB parishes	Test for difference
A. Parish characteristics			
population	8708.592	19037.66	1.8718**
ln(population)	8.228911	8.544173	3.8496***
accessibility (med)	140.2529	92.56836	-3.6108***
accessibility (avg)	159.3193	100.079	-4.5038***
eligible area (ha)	11672.46	4788.544	-5.5758***
communal contracts	8023771	1.81e+07	1.0994
slope (med)	16.57248	12.22017	-9.2585***
slope (avg)	17.45698	13.37632	-9.2340***
cash crops (annuals)	9762015	1.42e+07	2.3074**
cash crops (trees)	1.11e+07	2.02e+07	1.8583**
pasture	6.69e+07	3.58e+07	-4.9084***
B. Measures of forest cover loss in the years preceding the launch of Socio Bosque			
loss (ha)	52.2595	26.69808	-10.2919***
ln(loss)	2.632126	1.574022	-21.8370***
IHS(loss)	3.322268	2.210392	-26.0454***
loss (%)	.1630154	.1515194	-1.5331*

Note: in bold variables with significant differences. Test for difference gives the t-statistic for difference in means, stars indicate significance level ($^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$) of the unilateral test. Panel B spans years 2001 through 2007 ($N_B = 6,916$, while $N_A = 988$). Accessibility is in minutes to the nearest city, slope is in degrees, *loss (%)* is in percentage of parish area (both in hectares), *IHS(loss)* is the inverse hyperbolic sine of forest cover loss (in hectares).

With y_{it} the log forest cover loss in parish i and year t (2001-2017), $\mathbb{T}$ the set of parishes participating in Socio Bosque, β_t the coefficients of interest, λ_i, λ_t, parish and year fixed-effects, respectively, errors clustered at the parish level. Robustness checks take 2008 as base year instead of 2007 (as uptake in 2008 was very small compared to that of subsequent years), use natural instead of logged values, or compare non-enrolled parishes with only the areas of participating parishes that are indeed under Socio Bosque contract (which amounts to comparing the red and blue series in Figure C.2), with similar results and identical conclusions.

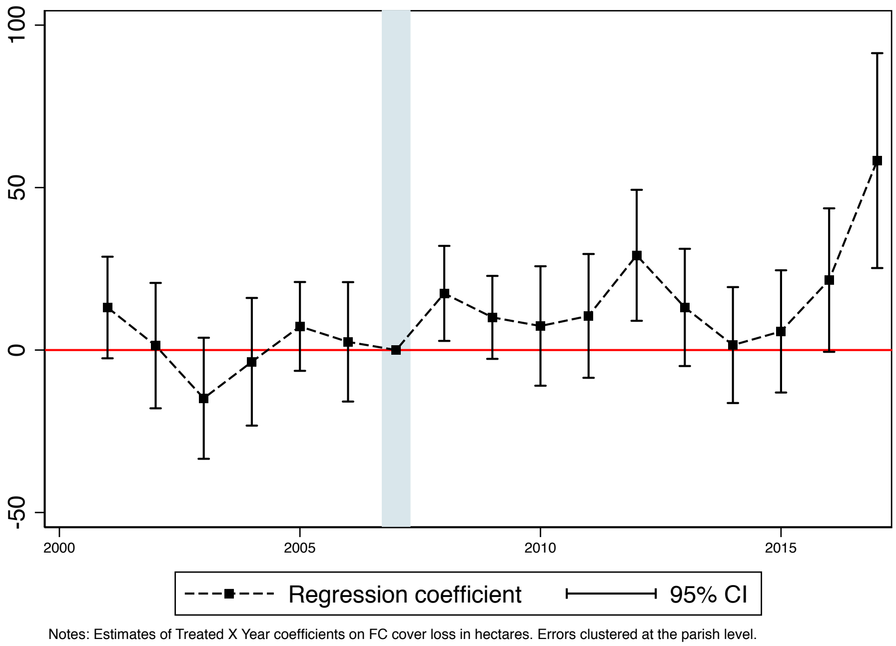

Figure 3.2: Impact of Socio Bosque on forest cover loss

Note: Difference-in-differences coefficients and their 95% confidence intervals (β_ts of Equation 3.1), with y_{it} forest cover loss in hectares, hence the interpretability of β_ts as the effect of Socio Bosque on forest cover *loss* in a given year and compared to non-participating parishes, in hectares; note that *none* is significantly negative (i.e. forest cover loss *never* decreased in SB parishes), and the point estimate does not go beyond -20, i.e. 15 hectares *at best* were preserved in a given year, which is dwarfed by the size of the typical contract (65 ha on average, median of 38 ha), let alone that of a parish. On the other hand, towards the end of the series, estimates, while small, become significantly positive, indicating a larger forest lost in SB parishes. The vertical line flags the reference year, 2007, right before the program started being implemented.

3.3 Results

3.3.1 Enrolment

As can be anticipated at the inspection of Figure 3.1, enrolment was biased towards steep and marginal land, that is less well suited for agriculture and other economic development. Indeed, Table 3.2 shows that parishes with enrolment in Socio Bosque are typically: less populated, less well-connected (it takes longer on average to reach a city), steeper, larger and better endowed in Socio Bosque-eligible land (i.e. forested, but not under another form of protection such as a natural reserve), less engaged in growing cash crops (bananas, coffee, etc.) and pasture.

It is somewhat surprising then, given deforestation drivers previously described in the literature (Geist and Lambin, 2002; Yackulic et al., 2011), that those same marginal parishes should have *larger* deforested surfaces for the pre-Socio Bosque period 2001-2007 (refer to Panel B. of Table 3.2).

Accordingly, self-selection into Socio Bosque of the land-owners least likely to engage in clear-cutting, and of the land least likely to be deforested, is likely to hamper the intended benefits of the program.

3.3.2 Impact

To get at the impact of the policy, it is obviously not sufficient to sum the hectares enrolled and claim they have been saved, nor is it even to compare forest cover loss before and after the roll-out of the program, for deforestation could be slowing down (or accelerating) in Ecuador as a whole, thus masking the effects of the policy itself. The difference-in-differences method enables to compare treated, non-treated, pre-, post-policy forest cover loss, and thus estimate its effect, under the assumption that treated and non-treated units follow parallel trends (which is verified, see Figure 3.2: before 2007, the coefficients are not different from zero).

Figure C.2 shows the evolution of forest cover loss in parishes with some (green dots), and with no (red dots, "non-treated") enrolment (left axis), as well as forest cover loss *within* Socio Bosque plots (right axis), 2001-2017. A striking observation is that all three series seem to be following the same temporal dynamics. A more important insight is that enrolled plots were not deforestation-free before Socio Bosque started (despite intact cover being a precondition for eligibility), and were not deforestation-free after, either.

Going past observations, the difference-in-differences analysis (see Figure 3.2) reveals that deforestation trends were not significantly different in treated and non-treated areas before the launch of Socio Bosque, but that, if anything, forest cover loss increased in "treated" parishes.

3.4 Discussion

This finding of no or moderately negative impact of the program on forest cover is disturbing; yet several plausible mechanisms, directly engendered by the Socio Bosque program, might lead to it. These are not mutually exclusive. First, the design of the program itself could explain the discrepancy between the objectives and the present findings. Indeed, Socio Bosque increases incomes and relieves liquidity constraints for the participants; as (Alix-Garcia et al., 2013) have shown with the Opportunidades program in Mexico, such transfers can raise the consumption of land-intensive goods, and thus can lead to increased deforestation (Socio Bosque is also meant as a poverty-alleviation program). Besides, no feature of Socio Bosque aims at curtailing leakage, despite it being a well-known and paramount issue in environmental policy, whether it pertains to greenhouse gas emissions (Stern, 2007) or forest conservation (Meyfroidt and Lambin, 2009; Alix-Garcia et al., 2012) – there is nothing indeed to prevent land-owners from enrolling one fraction of their forested land and cutting down a portion of the rest (possibly with the added help of the cash provided by the Socio Bosque program). Second, as Figure C.2 shows, levels of deforestation *inside* Socio Bosque plots (dotted line) were non-zero throughout the duration of the program, and followed fairly closely the parish-level trends (solid and dashed lines); thus another explanation for the null or even negative results is that the program was poorly enforced, and potentially even facilitated by the additional cash available (following the same logic as explained above). Finally, the payments for the program were interrupted in a staggered fashion over the course of 2014-15 as the government ran into a budgetary crisis; one interpretation of Figure 3.2 is then that the parishes that used to receive cash from the government, and to count on it to balance their household budgets, suddenly faced an unexpected liquidity constraint and a breach in contract on the part of the government, and resorted to cutting down trees as a result.

3.5 Conclusions

Policies for habitat preservation are increasingly favoring contractual schemes whereby conservation actions undertaken by private individuals are voluntary but incentivized by payments. This is where their appeal lies, as opposed to schemes relying on restriction of use or exclusion. It is also their weakness, but whether that is a fatal one is still an open debate. The present study analyzes such a scheme in Ecuador, and shows that not only enrolment was biased towards marginal land that was at a lower risk of deforestation anyway, but also that compliance fell short. As a result, deforestation dynamics were not altered by the program. This implies that targeting and enforcement should be a priority when considering designing a payment for ecosystem service scheme.

Appendix C

Appendix to Chapter 3

C.1 Additional figures

Figure C.1: Land tenure in Ecuador

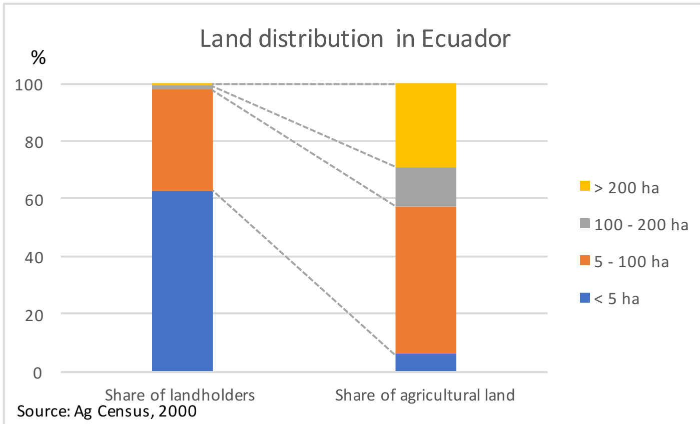

Notes: Own figure. Data from the Ecuadorian Ministry of Agriculture, 2000 agricultural census. This serves as an *indication* of the skewed structure of land ownership in Ecuador, but note that it only concerns agricultural land (which can't, by design, be enrolled under Socio Bosque).

Figure C.2: Deforestation trends in Ecuador.

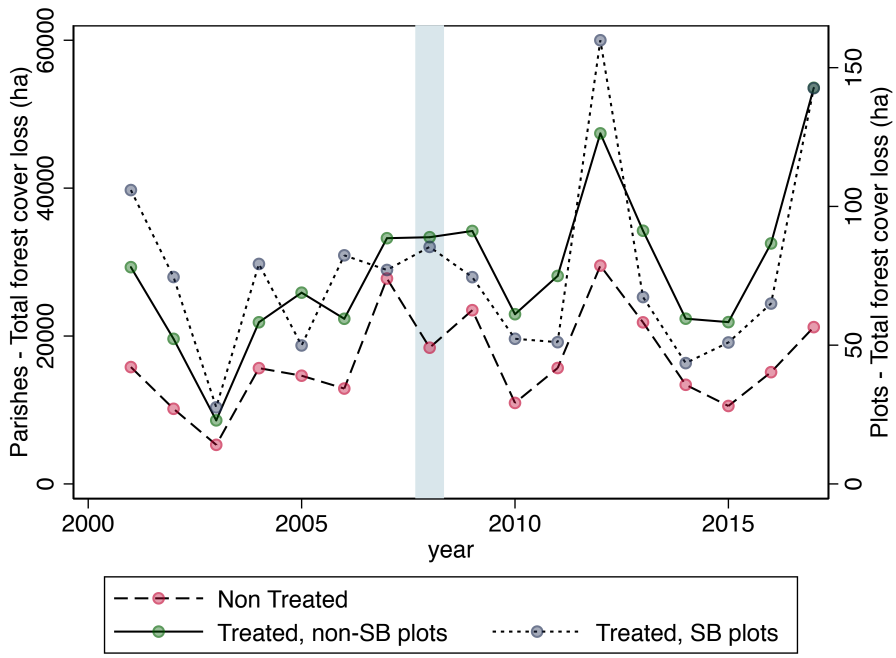

Notes: Figure displays deforestation trends in Socio Bosque plots, treated parishes, non-treated parishes, in hectares per year. The vertical line in 2008 corresponds to the year Socio Bosque was launched. Note separate scales for parishes (left; dashed and solid lines) and plots (right; dotted line).

Chapter 4

Asylum applications respond to temperature fluctuations

with Wolfram Schlenker

published as: Missirian, A., & Schlenker, W. (2017). Asylum applications respond to temperature fluctuations. *Science, 358*(6370), 1610–1614.

> “[...] *66 is the path of a people in flight, refugees from dust and shrinking ownership, from the desert’s slow northward invasion, from the twisting winds that that howl up out of Texas, from the floods that bring no richness to the land and steal what little richness is there. From all of these the people are in flight, and they come into 66 from the tributary side roads, from the wagon tracks and the rutted country roads. 66 is the mother road, the road of flight.*”
>
> John Steinbeck, *The Grapes of Wrath* (1939)

Abstract

International negotiations on climate change, along with recent upsurges in migration across the Mediterranean Sea, have highlighted the need to better understand the possible effects of climate change on human migration – in particular, across national borders. Here we examine how, in the recent past (2000-2014), weather variations in 103 source countries translated into asylum applications to the European Union, which averaged 351,000 per year in our sample. We find that temperatures that deviated from the moderate optimum (∼20°C) increased asylum applications in a nonlinear fashion, which implies an accelerated increase under continued future warming. Holding everything else constant, asylum applications by the end of the century are predicted to increase, on average, by 28% (98,000 additional asylum applications per year) under representative concentration pathway (RCP) scenario 4.5 and by 188% (660,000 additional applications per year) under RCP 8.5 for the 21 climate models in the NASA Earth Exchange Global Daily Downscaled Projections (NEX-GDDP).

Keywords: climate change, agriculture, temperature shocks, migration.
JEL Classification: Q5, Q54, F22, Q15.

The European Union (EU) has seen an unprecedented wave of immigration in 2015 (Frontex, 2016) as part of a larger surge in migration across the Mediterranean Sea that began in 2014. Many of the migrants flee war-torn countries such as Syria, Afghanistan, or Iraq, and there is an active debate as to whether a change in climatic conditions has contributed to, and will amplify, such migration flows. For example, a 2015 study has shown that the unrest in Syria was preceded by a record drought that led to lower agricultural yields and forced farmers to migrate to urban areas (Kelley et al., 2015). Although that study does not attribute the Syrian conflict to the drought, the authors argue that it added another stressor. These arguments have gained traction outside the academic literature. For instance, the Pentagon calls climate change a "threat multiplier" (Obama, 2015). However instead of looking at individual countries, we take a step back and investigate the role of weather shocks in global distress-driven migration to the EU in 2000–2014; i.e., preceding the recent crisis. Asylum applications to the EU from the 103 source countries in our sample totaled 1.5 million in 2015; that is, more than 4 times the average in our sample. Previous studies had found a relationship between weather variations and migration (Bohra-Mishra et al., 2014; Cai et al., 2016; Gray and Wise, 2016), but ours is the first to focus on distress-driven migration (as measured by asylum applications) on a global scale. Two centuries ago, the "year without a summer" (1816), following the volcanic eruption of Mount Tambora in Indonesia, saw massive crop failures throughout the Northern Hemisphere, caused by the aerosol-obscured atmosphere and unseasonal climate. It triggered sizeable migrations as peasants deserted their fruitless farms (Luterbacher and Pfister, 2015). Here we provide quantified evidence of a similar phenomenon taking place in the present day, whereby weather shocks on agricultural regions in 103 countries around the globe directly influence emigration, now toward the EU.

The relationship of international migration decisions to economic situation in both the source and destination country has been extensively documented. Migration's response to income or wealth corresponds in an inverted U shape: Positive income shocks in the home country enable individuals to overcome liquidity constraints and finance migration costs (Clemens, 2014). Richer households are not liquidity-constrained and show a negative migration-income relationship as improving conditions at home make it less desirable to leave (Bazzi, 2017). (See supplementary text sections D.1 and D.2 for a more detailed review and discussion (D).) Migration barriers have been described as one of the biggest distortions in the global economy (Clemens, 2011).

Causes of migration are not limited to the desire for better economic opportunities: humans flee persecution and war. We investigate how exogenous weather fluctuations affect one facet of migration: asylum applications, which equal roughly Embedded Image of the overall migration flows over our sample frame. Our sample included the 103 non-Organisation for Economic Co-operation and Development (OECD) source countries that reported asylum applications to the EU in each year between 2000 and 2014. It covered, on average, 351,000 asylum applications per year, the majority (140,000) coming from the 31 Asian countries, including Afghanistan and Iraq, each of which supplied ∼25,000 applicants. The 46 African and 11 non-EU countries in Europe accounted for 100,000 applicants each, whereas 16 countries in the Americas accounted for the rest (tables D.7 to D.9). For example, 55,943 people from Serbia applied for asylum in the EU in 2000. Applications from all source countries average 378,000 per year—that is, our sample covered 93% of all applications to the EU. Recent research (McGuirk and Burke, 2017) suggests that, in agricultural production areas, there should be a negative relationship between economic conditions and conflict, which then translates into asylum applications.

Our baseline regression links annual asylum applications from each source country outside the OECD to any EU member state. We use a panel analysis with source-country and year fixed effects, which is equivalent to a joint demeaning of all variables and accounting for common annual shocks. In other words, we link anomalies in log applications to weather anomalies once common annual shocks are absorbed (e.g., the global financial crisis in 2008). Our specification examines whether hotter-than-normal temperatures will increase or decrease asylum applications from a given source country. Because our dependent variable is in logs, we estimate relative impacts, which is preferable as the number of applications differs greatly among source countries in absolute terms. We allow the effect to vary by the average weather variable: Hotter-than-usual temperatures can reduce asylum applications for cold countries and increase them for hot countries. Our model includes both average temperature and precipitation. The coefficients and standard errors are given in table D.1.

We find a statistically significant relationship between fluctuations in asylum applications and weather anomalies: Applications are lowest for average temperatures around 20°C and increase if the weather is too cold or too hot. We choose to focus here on the EU because it receives the largest share of asylum application and, despite having a high rejection rate, remains a major provider of international protection (Missirian and Schlenker, 2017); other target ensembles are considered in the sensitivity checks. Colder countries in Europe outside the EU are predicted to account for fewer asylum applications in a warming world, whereas hotter countries, especially in Asia and Africa, are expected to see sizeable increases in a warming world (tables D.7 to D.9).

The coefficients on temperature are displayed in Figure 4.1. We show a quadratic response function (dashed brown line), as well as flexible restricted cubic splines (solid brown line). Both use the contemporaneous average temperature in the source country, averaged over the maize growing area and season. These models correspond to columns (1a) and (3a) of table D.1, respectively. Each line gives the point estimate and is normalized so that the minimum of the response function is zero. We find a highly significant relationship ($P < 0.01$ for joint significance) between logged asylum applications and average temperature over the maize growing area and season for the 103 source countries in our sample. If we average the weather on the basis of population in a grid cell (table D.2), the P value becomes 0.14 and the temperature variables are no longer significant, which suggests that weather shocks over the agricultural area are the crucial channel. The use of different weather data sets yields comparable results for seasonal averages (table D.3). Including data on political conflicts as controls (table D.4) produces important predictors of asylum applications, but the estimated relationship with temperature only slightly weakens, suggesting that they either pick up other forms of aggression or persecution because our conflict measures are limited to certain continents and actors or that the conflict data has measurement error.

The average temperature for which asylum applications are lowest is 21.4°C for the quadratic model and 19.9°C for the spline model. These values correspond to the optimal temperature range for agriculture (Schlenker and Roberts, 2009). Countries that are currently warmer than the optimal temperature would thus be predicted to produce an increase in asylum applications under a warmer climate. The range of observed average temperatures over our 15-year panel is depicted as green horizontal lines in Figure 4.1, and green "x" symbols denote the average over all 15 years. (A map of the current average climate over the maize growing season is provided in Fig. D.2, whereas Fig. D.3

Figure 4.1: Response of asylum applications to the EU with respect to the annual average temperature over the maize growing season

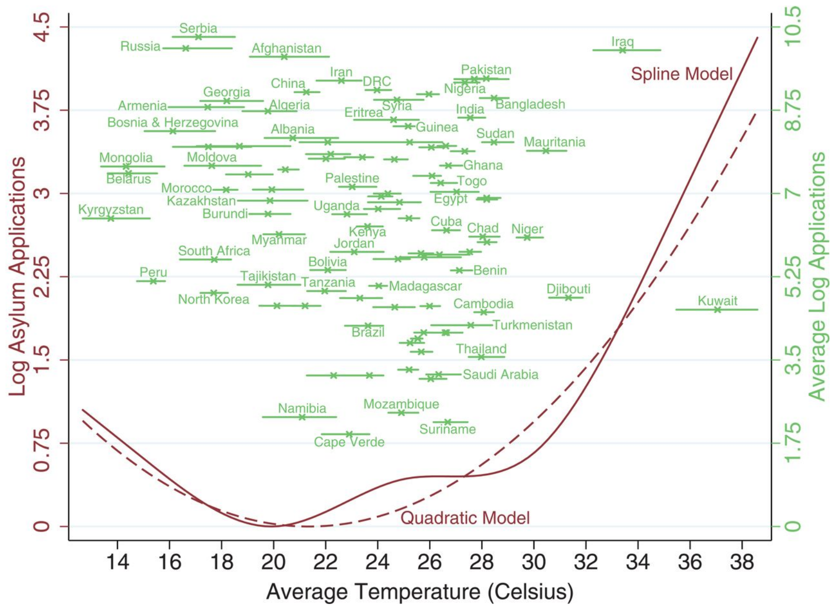

Notes: The quadratic response function is shown as a dashed brown line, whereas the restricted cubic spline is shown as a solid brown line (knots at 15°, 20°, 25°, 30°, and 35°C). Standard errors for the coefficients are given in table D.1. Because the models are in logs, the left y axis indicates the relative impact of changing temperatures on asylum applications. Each model controls for a quadratic function in season-total precipitation, as well as source-country and year fixed effects. The mean of the 15 annual average temperatures and log asylum applications (right y axis) for each source country are denoted by green "x" symbols. Because the models use weather anomalies in the identification, the green lines display the variation in annual average temperature in each country, ranging from the lowest to the highest observed value in the 15-year period.

shows the average temperatures over all months of the year and all grid cells in a country.) Because the regression is in log points, a y value of 1 implies an increase of 100 log points, or a $e^1 = 2.72$-fold increase in the number of applications.

Although the quadratic specification lends itself easily to interpretation of the regression coefficients and allows for some nonlinearity, it remains restrictive by assuming symmetry around the optimum. The more flexible model using restricted cubic splines with five evenly spaced knots between 15° and 35°C (supplementary text section D.2.2) enables us to relax this symmetry assumption, as well as the forced linearity in the marginal impact. The discovered relationship is slightly asymmetric, which suggests that temperatures above the optimum level are more harmful than those below this level.

Total precipitation, on the other hand, is not an important predictor of migration, consistent with previous research on conflict that indicates that temperature, as opposed to precipitation, is a stronger predictor of conflict (Burke et al., 2009). Moreover, the relative changes in temperature under future climate change scenarios translate into larger changes in yields than do precipitation changes (Schlenker and Lobell, 2010). When we exclude precipitation from the regression (column (2a) of table D.1), we obtain similar results.

Having established a consistent and robust U-shaped relationship between the weather in a source country and asylum applications—that is, temperatures that are too low or too high will lead to higher asylum applications—we now turn to simulations of how these applications will be altered under global climate change. We present both the response to hypothetical uniform temperature increases ranging from 1° to 5°C, as well as the predicted changes under the 21 global climate models in the NEX-GDDP (NASA Earth Exchange Global Daily Downscaled Projections) CMIP5 (Coupled Model Intercomparison Project phase 5) archive that estimate spatially heterogeneous warming scenarios. The U-shaped migration-temperature relationship suggests that colder source countries will experience a reduction in asylum applications to the EU under warmer temperatures, whereas warmer countries will see an increase thereof, as they lie to the right of the temperature that minimizes applications.

When we simulate the aggregate effect of a uniform 1° to 5°C warming, we compare the extrapolated number of asylum applications to the predicted applications under the historic observed weather in 2000–2014. Figure 4.2 accordingly illustrates the predicted changes in total applications to the EU by using the restricted cubic spline model (solid brown line in Figure 4.1 or column (3a) from table D.1), as well as the probability that they increase compared with the baseline. We prefer the cubic spline specification because the model is better at capturing nonlinearities within the observed range of temperatures and is more conservative in out-of-sample predictions, as the relationship is assumed to be linear above the highest knot whereas the quadratic model assumes increasing rates. Truncating the response function and forcing it to become flat outside the observed range does not substantially affect the overall results (Figs. D.6 and D.7). The results are comparable if we use the quadratic specification, as shown in the first row of Figure 4.3. The likelihood of an increase in applications is shown as a blue line in Figure 4.2 (right y axis) and ranges from 85% under +1°C warming to 99% under both +4° and +5°C warming. The change in the volume of applications is highly nonlinear: A 1°C warming results in a relative modest 6% increase in applications, but a 5°C warming leads to a 175% increase. The predicted mean change is positive in all models and under all warming scenarios.

Asylum applications are also forecasted using temperature data from the 21 climate models in the NEX-GDDP archive (supplementary text section D.3.4) for the medium term (2030–2059) as well as the end of the century (2070–2099). We compare them to an analogous 30-year historic baseline (1976–2005) in the NEX-GDDP data. The mean estimate and the confidence intervals are constructed using the same sample from the parameter distribution as for the uniform temperature increases and are shown in Figs. D.8 and D.9. Asylum applications are predicted to increase, on average, by 28% under representative concentration pathway (RCP) scenario 4.5 and 188% under RCP 8.5 by the end of the century, but there is large heterogeneity among countries (tables D.7 to D.9), with some colder countries seeing declines as they become warmer and closer to the optimal temperature.

These predictions are ceteris paribus. On the one hand, they might overstate the responsiveness, as the model uses historic weather shocks to identify the relationship while we apply it to a permanent warming scenario in which countries can engage in adaptive responses (e.g., shifting the growing season). On the other hand, these predictions might also be underestimates, as historic weather shocks are small enough in size that they likely do not capture disruptive events (such as major civil unrest) in case of continuous warming.

There are several likely mechanisms behind the sensitivity of fluctuations in asylum applications to temperature anomalies. First, there is a strong nonlinear relationship between agricultural yields and temperature (Schlenker and Roberts, 2009; Lobell et al., 2011; Tack et al., 2015). Moderate temperatures (i.e., in the lower 20°C range) over the growing season are ideal, with both hot and cold

Figure 4.2: Predicted changes to asylum applications under uniform climate change scenarios

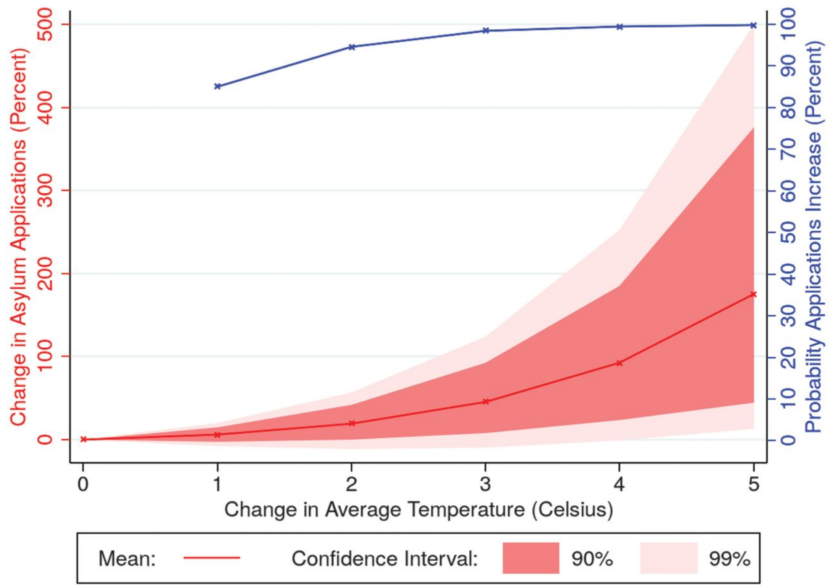

Notes: We used 1000 samples drawn from the joint distribution of the model parameters (solid brown line in Figure 4.1) to repeatedly predict the change in the percentage of total asylum applications filed in the EU. The solid red line shows the predicted change in percent, whereas the shaded areas illustrate the 90% and 99% confidence intervals. The blue line (right y axis) indicates the probability that asylum applications increase.

temperatures reducing yields. Second, gross domestic product (GDP) growth rates have been found to be very sensitive to temperature, even on the nonagricultural components of the GDP and even in industrialized countries (Dell et al., 2012; Burke et al., 2015). For both yields and GDP, being too hot is worse than being too cold. An improvement in agricultural output or GDP will help producers and workers (if they are paid their marginal product, which increases) in those regions, reduce the incentive to join criminal activity, and lead to less conflict (McGuirk and Burke, 2017). Third, the same sign in the relationship with weather is found for aggressive behavior, which increases with temperatures that have been shown to reduce output (Reifman et al., 1991). This relationship holds consistently across several spatial and temporal scales (Hsiang et al., 2013). The first two mechanisms are a priori conducive to increased distress-driven migration in the event that weather deviates from the moderate optimum, and all three predict that increases in very warm locations are likely to be associated with higher asylum applications.

The results of our baseline regression suggest that suboptimal weather (namely, temperatures that are too cold or too hot) increases asylum applications to the EU. One might wonder whether these additional applications are caused by heightened persecution or just by changing economic conditions, with both being credible intermediates in a causal chain linking weather anomalies and demand for asylum. To clarify this issue, we examine the numbers of accepted applications per year between a source country and the EU. We do find that the weather-induced spikes in applications translate into roughly three times higher acceptance rates in the following two years (see supplementary text section D.2.4 for more detail), suggesting that destination countries classify the additional cases as more deserving than the average applicant and see them as refugees and not economic migrants.

Figure 4.3: Sensitivity checks to various modeling assumptions

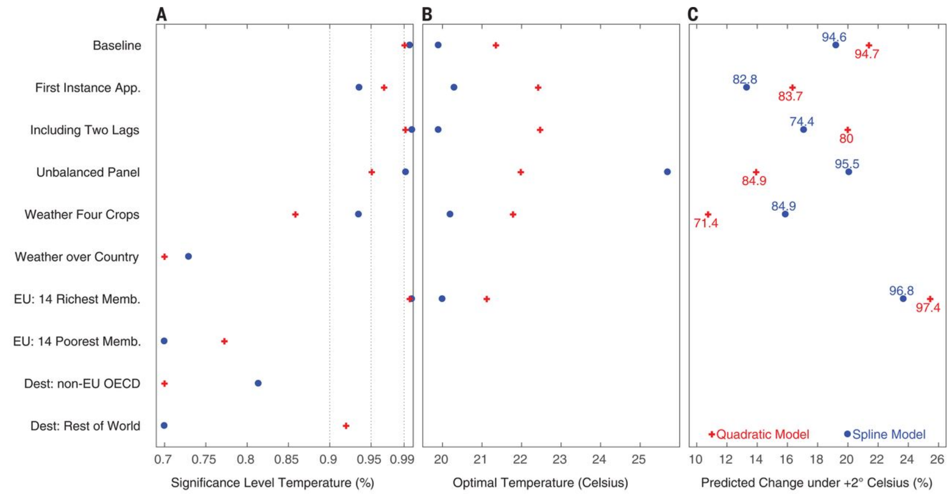

Notes: Each row represents a separate model. (**A**) P value for joint significance of the temperature variables (values below 0.7 are shown as 0.7). (**B** and **C**) For cases in which there is a significant link between temperature and asylum applications (average significance level above 0.9 for the quadratic and spline model), (**B**) shows the temperature that minimizes asylum applications, and (**C**) illustrates the predicted change in applicants (x axis) as well as the probability that a +2°C uniform warming will lead to an increase in applications to the EU (number next to the estimate).

We conduct several sensitivity checks of our baseline results that are summarized in Figure 4.3 to rule out the possibility of our model picking up spurious correlations in the data. This figure includes three panels: Panel A shows the joint significance for all temperature variables. In cases for which there is a significant link between temperatures and asylum applications, panel B depicts the temperature that minimizes asylum applications, and panel C illustrates the predicted change in asylum applications under uniform +2°C warming, as well as the probability that applications will increase. Results displayed in panels B and C are generally robust.

The first row in Figure 4.3 presents our baseline results for comparison. Results for the quadratic model are shown in red, whereas results for the model using restricted cubic splines are shown in blue (all models account for a quadratic function in season-total precipitation, but because they were not significant, the coefficients are not used when we evaluate climate impacts).

The second row limits asylum applications to case openings (rather than second instances, appeals, etc.). In principle, both first and subsequent applications could be influenced by weather, and the results are comparable when we limit the data to first instances. In our baseline regression, we simply added all instances.

The third row includes two lags of all weather variables. We show the combined results of the contemporaneous term as well as the two lags. Lagged variables might capture two opposite effects: (i) delayed increase of applications, accounting for cases when people apply the following year(s), as they might not be able or willing to flee right away, and (ii) forward displacement, whereby weather shocks induce the applications of individuals who were contemplating to leave in the next year(s) anyway, in which case the contemporaneous effect should be counterbalanced by coefficients of the opposite sign in the following years. Either these two effects are of small magnitude or they balance each other out, as the model with two lags produces similar predictions.

The fourth row uses a model that includes all source countries, even if they do not report applications for all of the 15 years, with little change (the optimal temperature in the spline model in panel B becomes a bit larger, but this has little effect on the predicted climate impacts in panel C).

The fifth row derives the weather not just over the maize growing area and season but also over the growing areas and seasons of the other three large staple commodities: wheat, rice, and soybeans. These four weather variables are then averaged using area weights. The temperature coefficients are no longer significant for the quadratic model, but the remaining results are robust to these changes.

The sixth row averages the weather over the entire country and year and no longer produces significant results. This suggests that the results are driven by events that happen in the agricultural (rural) areas during the time the crops are grown and is in line with the predictions that a significant result should mainly be observed for agricultural production areas (supplementary text section D.2.1). We are not able to distinguish whether these events affect agriculture directly or other sectors in the same (rural) area at the same time of the year.

The seventh and eighth rows break the destination into two subgroups: the 14 richest and 14 poorest member states of the EU. Results are only significant for applications to the 14 richest member countries that absorb most of the applicants.

The last two rows examine applications to countries outside the EU, which we break into two subgroups on the basis of whether or not they belong to the OECD. There is no significant relationship with weather in the source country for the former, but these countries accept a small share of asylum applications (Missirian and Schlenker, 2017). The rest of the world (non-EU, non-OECD) predominantly accepts applications from neighboring countries, but we observe only a marginally significant effect of the weather in the source country for the spline model.

Finally, we also tested for heterogeneity among source countries but generally did not find significantly different sensitivities to weather for various subgroups. The results, in which we split the sample into roughly equal halves based on (i) the corruption index of the source country, (ii) the latitude of the source country, (iii) the population of the source country, (iv) the share of the source countries' workforce that is employed in agriculture, or (v) the distance between the source country and the EU, are shown in tables D.5 and D.6.

In summary, we link annual asylum applications received by the EU member states to average temperature over the maize growing area and season in the source country and find a nonlinear relationship, especially for those applications filed into the richer EU member states. Moderate temperatures around 20°C minimize asylum applications. Both colder and hotter temperatures increase migration flows. Extrapolating those results, an increase in temperatures in source countries is predicted to lead to an increase in asylum applications to the EU as well, following a highly nonlinear response function. Our findings support the assessment that climate change, especially continued warming, will add another "threat multiplier" that induces people to seek refuge abroad. Weather impacts in low-income source countries will not be confined to those countries or regions but will instead likely spill over into developed countries through increased refugee flows.

Appendix D

Appendix to Chapter 4

D.1 Literature Review

D.1.1 The Effect of Weather and Climate on Migration

There is an extensive literature on how environmental factors influence migration decisions. Recent concerns over greenhouse gas emissions and the associated change in climatic condition, for example aggravated droughts or sea level rise, have brought the effect of environmental conditions on migration back to the forefront (Piguet and Laczko, 2014, pp. 4-5). Various temporal scales (short-term versus long-term migration), spatial scopes (within country or international), climatic factors (temperature and precipitations), as well as channels (income and conflict) have been investigated. Most studies find that migration responds to temperature variations, but the results are mixed: Cai et al. (2016) show that migration towards mostly OECD countries is increasing in population-weighted temperatures in the source country, and that the effect is driven by those source countries that have a predominantly agricultural economy. Bohra-Mishra et al. (2014) find a nonlinear relationship with temperature, and identify an optimal temperature above which temperature increases lead to an increase in migration, while the opposite is true below the threshold. They argue that in Indonesia the optimal temperature is 25°C, mirroring optimal temperatures found in agriculture. According to Bohra-Mishra et al. (2014) this highlights the importance of the income channel in the temperature-migration relationship (a negative income shock leads to a higher probability of permanent emigration from the affected Indonesian province), and incidentally corroborates "the importance of the agricultural linkage" argued by Cai et al. (2016). Similarly, Cattaneo and Peri (2016) find that warmer temperatures increase the trend to emigrate to urban areas or other countries in middle income countries but not in poor countries, where households are liquidity constrained. The study assumes a linear response to temperature trends (except for a sensitivity check where they estimate two linear response functions, one for hot and one for cold countries). On the other hand, Gray and Wise (2016), for instance, find no relationship between temperature anomalies and emigration in two of the five countries they study (Senegal and Nigeria), highlighting the non-triviality of the effects of environmental factors on migration. Finally, while most studies find little or no effect of precipitation on migration, Henry et al. (2004) find that in Burkina Faso, precipitation is strongly related to emigration, but in intricate ways: the likelihood to emigrate, as well as the distance and the duration vary depending on how wet the region is and on the magnitude of the water stress, and across socio-demographic groups. For example, men tend to migrate abroad much more when precipitations fall within historical ranges than when they are just below average.

D.1.2 The Effect of Conflict on Migration

A possible channel how temperature might impact distress-driven migration is conflict. Diverging results have been published on the matter (e.g. see Adger et al. (2014); Oppenheimer et al. (2014) for an overview), most studies conclude that climate variability cause, or are at least associated, with violence. Opposing analyses have been refuted on statistical grounds (Hsiang et al., 2015). For instance, Hendrix and Salehyan (2012) show increase in social conflict to be associated with rainfall shocks (positive and negative). Iyigun et al. (2017) find a relationship between *long-term* temperature change and violence, documenting the increased violence caused by the cooling that occurred during the Little Ice Age. The article also provides evidence that it is possible for intensification effects (consecutive "bad" events amplify) to dominate over adaptation effects (response to changing conditions become muted through adaptation) in the case of long-run climate change. In other words, the persistent change in climate makes things worse for people each year, and adaptation does not occur fast or significantly enough, possibly because of the damaged social structures and limited resources imposed by the climate changes. Going beyond individual studies, Hsiang et al. (2013) propose a meta-analysis of the quantitative evidence at hand, and find that "deviations from normal precipitation and mild temperatures systematically increase the risk of conflict, often substantially", and further, that "each 1-SD change in climate toward warmer temperatures or more extreme rainfall increases the frequency of interpersonal violence by 4% and intergroup conflict by 14% (median estimates)."

D.2 Model

D.2.1 Economic Model of Migration Decisions

Migration flows respond positively to the difference in living standards in destination and source countries, whether they are due to different incomes or amenities (Sinha and Cropper, 2013; Albouy et al., 2016). Individuals i in source country s are contemplating whether to apply for asylum in destination country d. They have an indirect utility function $v(p, y, c)$ that depends on local prices p, income y, and conflict c. Utility is decreasing in the price level ($\frac{\partial v}{\partial p} < 0$) and increasing in income ($\frac{\partial v}{\partial y} > 0$) as one can afford to buy less goods if prices go up, or if one's income is lower, all else being the same. The income distribution $g()$ of the individuals i in the source country ranges over $[\underline{y}, \overline{y}]$. Finally, utility is decreasing in conflict ($\frac{\partial v}{\partial c} < 0$).

We use subscripts s and d to differentiate between source and destination countries. Income varies across indviduals i, while prices and conflict are given by local conditions that are the same for all individuals. People desire to migrate if $v_{id}(p_d, y_{id} - m, c_d) > v_{is}(p_s, y_{is}, c_s)$, but can do so only if they can afford the migration cost m, i.e., $m < y_{is}$. Otherwise liquidity constraints will keep them from realizing this desire. In summary, people that will migrate are given by (the lower bound of the integral is due to the liquidity constraint):

$$\int_{m}^{\overline{y}} \underbrace{1_{[v_{id}(p_d, y_{id}-m, c_d) > v_{is}(p_s, y_{is}, c_s)]}}_{\text{indicator function}} g(y_{is}) dy_{is}$$

Recent empirical research has shown that extremely hot temperatures are detrimental in many ways: they decrease yields (Schlenker and Roberts, 2009), labour supply (Graff Zivin and Neidell, 2014), and climate amenities (Albouy et al., 2016), while increasing mortality and energy cost (Deschênes and Greenstone, 2011). Let w_s measure extreme heat, so $\frac{\partial y_{is}}{\partial w_s} < 0$. There are four channels through which weather will impact migration, either directly or indirectly:

1) Effect of weather on economic outcomes (income y_{si}):

   1a) Effect on people that already could afford to migrate: Given that $\frac{\partial v_{is}}{\partial y_{is}} \frac{\partial y_{is}}{\partial w_s} < 0$, an increase in w_s will decrease the utility of staying in the source country, and increase the difference between destination and source countries. This will make migration more desirable. However, incomes in destination countries already tend to be a lot larger than in source (mostly developing) countries, so the indicator $1_{[v_{id}(p_d, y_{id}-m, c_d) > v_{is}(p_s, y_{is}, c_s)]}$ might not be impacted for the set of people that were able to afford to move in the first place. Earlier research on migration finds this channel to be of less importance for poorer people.

1b) Effect on who can afford to migrate: Since extreme heat will decrease income $\frac{\partial y_{is}}{\partial w_s} < 0$, the set of people that can afford the moving cost m will contract. This results in a negative association between extreme heat and migration and is consistent with a pro-cyclical response to income shocks, as more heat implies worse economic conditions and in turn less migration.

2) Effect of weather on conflict:

2a) There is a direct effect of extreme heat (higher temperatures) on aggression (Reifman et al., 1991; Hsiang et al., 2013). An increase in heat increases conflict ($\frac{\partial c_s}{\partial w_s} > 0$), decreases the utility in the source country, and increases the gap between destination and source countries. This corresponds to a positive correlation between extreme heat (or temperature more generally) and asylum applications.

2b) The second channel depends on whether an area is a net agricultural exporter, i.e., whether it is producing more than its inhabitants consume. As outlined above, extreme heat will decrease yields. In production areas, a decrease in extreme heat increases yields and profits, as prices are set by the global market and hence don't change in response to local production shocks. Workers will also benefit if they are paid their marginal product, which is now higher. Given that producers and workers in these agricultural areas are doing better, it becomes less desirable to join an armed "factor conflict" (McGuirk and Burke, 2017). Overall, there is a positive correlation between extreme heat and conflict. On the other hand, in urban areas where more is consumed than produced, the increase in agricultural surplus increases the rewards from engaging in food riots or "output conflicts" (McGuirk and Burke, 2017), since weather is highly spatially correlated and hence weather in consumption areas is positively correlated with weather in close-by production areas. This results in a negative correlation between extreme heat and asylum applications.

Previous studies examining economic rationales for emigration (points 1a and 1b above) have found that liquidity constraints (1b) tend to dominate initially when individuals are poor (Bazzi, 2017). Richer constituents might not longer be constraint and show the opposite. For similar reasons, cash transfers to poor individuals increase emigration (Angelucci, 2015; Bryan et al., 2014), both in the current as well as future periods, as they help people to overcome the migration cost hurdle. Earthquakes reduce emigration from El Salvador (Yang et al., 2008) because the aggregate shock they generate on the local economy thwarts potential migrants' reliance on financial assistance in order to pay for the upfront costs of migration. The opposite is true for *return* migration, as liquidity constraints for returning to the home country usually play less of a role and hence point 1a dominates. For example, fewer migrants return when the destination country has positive exchange rate shocks (Yang, 2006) and more migrants return from Australia when their respective home country is doing well (Abarcar, 2017).

Our study does not focus on the economic reasons of migration, but on asylum-seekers that flee and seek international protection abroad (points 2a and 2b above). We focus on weather in agricultural production regions for two reasons. First, the direct effect of weather on aggression, and the indirect effect through incomes (agricultural profits and wages) go in the same direction. By contrast, in urban consumption areas, they go in opposite directions, and the overall effect is hence unknown. Second, "factor conflicts" (McGuirk and Burke, 2017) are more severe than "output conflicts." The former push people off their land while the latter result in food riots that might not induce people to leave a country altogether and seek refugee status. While the literature looking at the economic reasons for migration has found a pro-cyclical response for poor individuals, which would be equivalent to a positive sensitivity of migration to extreme heat, we expect a negative sensitivity for asylum applications, i.e. they should increase with extreme heat.

D.2.2 Econometric Specification

We run a fixed effect model linking asylum applications to the European Union to weather in the source country. Specifically, the model is:

$$y_{st} = \alpha_s + \sum_{\tau=0}^{l} \mathbf{W}_{s(t-\tau)} \boldsymbol{\beta}_\tau + \gamma_t + \epsilon_{st}$$

Log asylum applications y_{st} from source country s to any member of the European Union in year t are regressed on source-country fixed effects α_s to account for baseline differences as well the weather $\mathbf{W}_{st} = [W_{st1}, \ldots, W_{stk}]$ (k weather variables) in source country s in year t. Since we are including source-country fixed effects, our regression only uses the random and exogenous year-to-year variation in weather (anomalies) in the identification but does not rely on baseline difference (e.g., different forms of government might result in different average number of refugees fleeing a country). This is preferable for obtaining causal estimates of the relationship (Dell et al., 2014). Figure D.4 displays histogram of resulting anomalies by year for the 103 countries in our sample for both average temperature over the maize growing area and season (top graph) as well as log applications (bottom graph). Figure D.5 shows the temporal evolution for the four countries with the largest number of asylum applications: Serbia, Iraq, Russia, and Afghanistan. Note the randomness in year-to-year fluctuations in the temperature anomalies (top graph), which is ideal for statistical purposes. The bottom graph shows application anomalies for the four countries with the largest number of applications. While there is variation over time, there are no clear trends.

The baseline specification uses only contemporaneous weather ($l = 0$). For the weather variables $\mathbf{W}_{st}$, we use a quadratic both in the average temperature over the maize growing season and in the season-total precipitation over the maize growing season ($k = 4$ as we get two terms for temperature and precipitation each). To allow for a more flexible response of asylum applications to temperature, we also estimate a model where we replace the quadratic in average temperature with a restricted cubic spline in average temperature with 5 knots (at 15, 20, 25, 30, and 35°C), which results in $k = 6$ weather variables (4 temperature variables and the same two precipitation variables). Some specifications also include two lags $l = 2$ of the weather variables ($\mathbf{W}_{it}, \mathbf{W}_{i(t-1)}, \mathbf{W}_{i(t-2)}$) to allow for delayed effects. We include year fixed effects γ_t as temporal controls to adjust for common shocks like the 2008 financial crisis. Error term ϵ_{it} uses White's correction (robust standard errors).

D.2.3 Simulating Changing Temperatures

We simulate the aggregate impacts from uniform warming in the following way. We predict asylum applications for the historic observed weather in 2000-2014 and add them over all 15 years and all source countries. Next, we add the predicted warming to each historic weather record, e.g., increase the observed temperature in each country and year by 1°C and recalculate the predicted applications for all years and countries before summing them. The change is the ratio of the predicted applications under the warmer weather to the predicted historic applications in our sample. We are summing over all countries, and the effects in countries that historically had higher numbers of asylum applications will hence be weighted more heavily. In order to incorporate the model uncertainty in our estimates, we repeat the above exercise 1000 times, and each time draw from the joint distribution of the parameters. This gives us a range of predicted aggregate impacts.

We also derive predicted climate change impacts for the 21 climate models included in the NEX-GDDP data. We again predict applications both in the future as well as an historic baseline - we choose 1976-2005 baseline as the year 2005 was the last year of historic temperature data used in CMIP5 projections, i.e., different RCP scenarios started in 2006. We again aggregate over all countries when we summarize the results. A break-down of the effect by source country is shown in Tables D.7-D.9.

D.2.4 Anomalies in Asylum Applications and Acceptance

We derive an indirect test to find whether spikes in weather-induced asylum applications are due to actual persecution by checking whether those anomalies translate into higher acceptance rates in the following two years. If asylum applications spike for economic reasons, we should see a decline in the acceptance rate as destination countries deem the application invalid. On the other hand, if the spike is due to increased conflict or persecution, the acceptance rate should increase as destination countries deem these applications valid and accept them. The average acceptance rate of applications into the European Union in 2000-2014 was 11.4 percent.

We compute anomalies in asylum applications due to weather shocks as the predicted change in the number of application from a source country that are explained by deviations in the weather variables from their respective averages $\overline{\mathbf{W}}_s = [\overline{W}_{s1} \ldots, \overline{W}_{sk}]$ for the k temperature variables (k is 2 in case of the quadratic specification, and 4 when we use the restricted cubic splines in average temperature). The application anomaly n_{st} for source country s in year t is hence

$$n_{st} = \underbrace{e^{\widehat{\alpha_s} + \mathbf{W}_{st}\widehat{\boldsymbol{\beta}} + \widehat{\gamma}_t + \frac{1}{2}\widehat{\sigma^2_\epsilon}}}_{\text{app. observed weather}} - \underbrace{e^{\widehat{\alpha_s} + \overline{\mathbf{W}}_s\widehat{\boldsymbol{\beta}} + \widehat{\gamma}_t + \frac{1}{2}\widehat{\sigma^2_\epsilon}}}_{\text{app. mean weather}}$$

Where the parameters $\widehat{\alpha_s}$, $\widehat{\boldsymbol{\beta}}$ and $\widehat{\gamma}_t$ are the coefficients from our baseline regression of log asylum applications on weather (see Section D.2.2), and $\widehat{\sigma^2_\epsilon}$ is the predicted variance of the error terms ϵ from the same regression. In a second stage we then examine decisions (acceptances) d_{st} in the following two years

$$d_{st} = \mu_s + \sum_{\tau=0}^{2} \delta_\tau n_{s(t-\tau)} + \phi_t + \nu_{st}$$

The coefficients $\delta_0, \delta_1, \delta_2$ measure how shocks (anomalies) translate into additional acceptances. The sum of the three coefficients is 0.40 (t-value: 2.91) for the quadratic model and 0.27 (t-value: 4.75) under the spline model. Both are much higher than the baseline average acceptance rate of 0.11, suggesting that these weather-induced shocks to applications are assessed as legitimate by destination authorities at a much higher rate. In other words, this indicates that weather shocks lead to increased conflict and persecution, which lead people to migrate and be recognized as needing international protection through refugee status.

D.3 Data

D.3.1 Asylum Applications

We use freely available data on asylum-seekers by the United Nations High Commissioner on Refugees.[1] The data set includes the number of asylum applications during each year between source and destination countries. We add all applications from a given non-OECD source country to destination countries that are part of the European Union (EU) to obtain the total number of applications coming from that country into the EU in a given year. This data is available for the years 2000-2014. Our baseline specification concentrates on countries that had non-zero applications in each of the 15 years.

D.3.2 Weather Data: University of Delaware (Baseline)

We pair this annual panel of applications by non-OECD source countries to weather data in the source country. The University of Delaware data set (version 4 has data until 2014) gives monthly average temperature and total precipitation on a 0.5x0.5 degree grid.[2] We aggregate this gridded monthly weather data to an annual country level in two ways. First, we simply average all grid cells in a country over the entire year. Second, since a large share of the population in most source countries work in agriculture, and since this sector is depend on weather fluctuations, we construct the weather exposure for the maize growing area, maize being the staple commodity that accounts for the largest share of human's caloric intake and is grown around the world (Roberts and Schlenker, 2013). Figure D.1 shows where maize is grown around the globe. We use the gridded data set of Sacks et al. (2010) to construct the fraction of each climate grid in a country that grows maize. Moreover, the data set gives the start and end dates of the maize growing season. Since the weather data is monthly, we define the growing season to start on the first of the month of the average planting date and to end on the last of the month of the average harvest date. If the crop is grown more than once a year, we focus on the first season. The average temperature is then averaged over all grids in a country, where the weights equal the maize growing area in a cell. These are time-invariant weights. This gives us the weather averaged over the time and area where maize is grown in a country. We first calculate nonlinear transformations (the square or a restricted cubic spline of the average temperature over season) for each climate grid, and then the average, since the average of a nonlinear function is different than the nonlinear function of the average.

In a sensitivity check we also calculate the weather over the growing area of the three other big staple commodities besides maize: rice, wheat, and soybeans. Together with maize they account for 75% of the caloric intake of humans, either direct or indirect through the use as feedstock for animals

[1] Source: http://popstats.unhcr.org/en/asylum_seekers

[2] https://www.esrl.noaa.gov/psd/data/gridded/data.UDel_AirT_Precip.html

(Cassman, 1999). The weather for a country will be different for different crops, as various commodities are grown in different parts of a country as well as during different months. As a result, the average temperature for rice in a country is generally higher than the average temperature for wheat in that same country. Finally, we use the weighted average of the weather averaged over the four commodities, where we use the growing area for each crop in Sacks et al. (2010) as weights.

D.3.3 Weather Data: Berkeley Earth

For a sensitivity check we also downloaded the daily weather data from Berkeley Earth for the years 2000-2013.[3] The grid is slightly coarser at 1 degree latitude and longitude, and the temporal coverage ends in 2013 instead of 2014. The aggregation procedure is analogous to that applied to the University of Delaware data except that we now use the exact days of the growing season (not the start of the month when planting occurs, or the end of the month when harvest occurs that we used for the monthly data).

D.3.4 Climate Projections: NEX-GDDP

We download the global daily downscaled projections from NASA's Earth Exchange (NEX-GDDP).[4] It includes daily data for 21 climate models under the RCP4.5 and RCP8.5 scenarios on a 0.25 grid scale in both latitude and longitude. We use this to derive both future predicted temperature scenarios (medium term 2030-59 as well as long term 2070-2099) as well as a historic baseline (1976-2005). We use the same aggregation procedure as for the baseline data in Section D.3.2, i.e., still define the growing season from the first of the month of the planting date to the last day of the month of harvest and then aggregate all grids in a country using the cropland area to make the climate projections consistent with the data used in our baseline econometric analysis.

D.3.5 Population Weights

We also obtained a gridded data set on the population counts from NASA's Socioeconomic Data Applications Center:[5] Gridded Population of the World (GPW), v4 for the year 2000. To keep the time frame comparable, we first derive the average growing season of maize among all grids in a country and then average the weather over this time period, but instead of using the maize growing area as weight, we use the population count in each grid cell as weight.

D.3.6 Additional Data

We conduct several sensitivity checks to see whether there is heterogeneity in the response for different sets of source countries. We generally split the source countries into roughly equal halves.

[3] http://berkeleyearth.org/data/

[4] https://nex.nasa.gov/nex/projects/1356/

[5] http://sedac.ciesin.columbia.edu/data/set/gpw-v4-population-density/data-download

Geographic Data

We use the country shapefiles by ESRI to derive the distance between a destination country and the European Union, calculated as the minimum distance between the coordinates of the source country and the coordinates of the 28 EU member states. A common land border would imply a distance of 0. A distance of 2700 km roughly splits our sample in half between close and far-away source countries.

Similarly, we define countries as tropical if any of the latitude coordinates of their country shapefile are within 23.5 degrees of the Equator.

Socio-economic Data: CIA

We use data from the CIA World Factbook[6] on population in the source country as well as GDP per capita among the 28 destination countries in the European Union.

We separate source countries based on whether they have a population of less or more than 15 million (this measure is time-invariant) and see whether the results differ.

The 28 EU member states are split into the countries with the highest/lowest GDP per capita in sensitivity checks. The 14 member states with the highest GDP per capita according to the CIA Factbook are: Luxembourg (102900), Ireland (54300), Netherlands (49300), Sweden (48000), Austria (47500), Germany (47400), Denmark (45800), Belgium (44100), France (41400), Finland (41200), United Kingdom (41200), Italy (35800), Spain (35200), Malta (34700). The 14 member states with the lowest GDP per capita according to the CIA Factbook are: Bulgaria (18400), Romania (20600), Croatia (21300), Latvia (24500), Greece (25600), Hungary (26000), Poland (26400), Portugal (27800), Lithuania (28000), Estonia (28700), Slovakia (29500), Slovenia (30900), Cyprus (31000), Czech Republic (31500).

Socio-economic Data: World Bank

We downloaded data from the World Bank[7] on the share of the population that is working in agriculture. We average this annual time series for our sample frame 2000-2014 to obtain time-invariant average shares by country.

Corruption Data

Data from Transparency International[8] ranks countries based on how corrupt they are. We use the latest country index of 2015 as it covers more countries than what is available for 2000.

Conflict Data

Data from the UCDP/PRIO Armed Conflict Dataset version 4-2016[9] was used to construct the number of minor (between 25 and 999 battle-related deaths per year) as well as major conflicts (at least 1000 battle-related deaths). This data set include conflicts where at least one party is a government of a

[6] https://www.cia.gov/library/publications/

[7] http://data.worldbank.org/indicator/SL.AGR.EMPL.ZS

[8] http://www.transparency.org/

[9] http://ucdp.uu.se/downloads/

state. We assigned a conflict to a country if it was named among the list of countries on side A or B (not counting secondary countries involvement that provide support). A government can be involved in more than one conflict, and hence this variable can take on integer values.

Data from the Armed Conflict Location and Event Data Project[10] lists political violence and protest events for mainly African countries. We sum all fatalities for each country and year in thousands.

D.4 Additional Results

D.4.1 Weather Data: Sensitivity Check using Population Weights

Table D.2 contrasts the results when we include weather variables that are averaged over the maize growing area and season in a country in columns (a), weather variables that are averaged over the average maize growing season in a country using population weights in columns (b), as well as both weather measures in column (c). It should be noted that the temperature measures over the agricultural area in highly significant (p-value < 0.01 in columns (a)), while the population weighted ones in column (b) are not (p-values of 0.14 and 0.15 for the quadratic and spline model respectively). If we include both weather measures in column (c), p-values decrease, but the maize-area weighted temperature coefficients have lower p-values than the ones for the population-weight ones.

D.4.2 Weather Data: Sensitivity Check using Berkeley Earth

Results are given in Table D.3. The first three columns (1a)-(1c) show the results for the model using a quadratic specification in the average temperature, while the last three columns (2a)-(2c) show the results for a model using restricted cubic splines in average temperature. Columns (a) give the baseline results when we use temperature and precipitation data from the University of Delaware. Columns (b) replicate the analysis using the same weather data source, but only estimate it for the countries and years where we have data available from both data sources. Fewer countries are covered in Berkeley Earth's grid and the data ends one year earlier (2013 versus 2014). Columns (c) use the same countries and years as columns (b) but replace the monthly temperature data from the University of Delaware with the daily data from Berkeley Earth. We keep the precipitation data fixed from the University of Delaware to focus on one change at a time.

For the quadratic specification in columns (1a)-(1c), the results are very consistent. For the spline model, columns (2a) and (2b) give comparable results, but the temperature coefficients are no longer significant in column (2c). We believe this is likely due to the coarser grid of the Berkeley Earth data set. For non-linear models, it is important to first apply the nonlinear transformation before averaging instead of first averaging and then applying the nonlinear transformation.

We also experimented by first conducting a nonlinear transformation of the daily temperatures before aggregating to the season, but the results were not significant as we likely have a large amount of attenuation bias due to the required interpolation, i.e., the same reason why column (2c) was no longer significant.

[10] http://www.acleddata.com/data/

D.4.3 Including Conflict Data

Table D.4 includes various measures of conflict in the estimation as additional controls. While these data have their own drawback as conflicts are clearly endogenous, one might wonder whether the inclusion of these controls at least have the expected sign (and hence confirm the validity of the asylum application data) and whether accounting for conflicts makes our results on weather go away (if the main channel is through conflicts, which are now directly measured). Conflict variables are all highly statistically significant: more conflicts in the source country increases the number of asylum applications from that country. Major conflicts have a larger impact than minor conflicts.

Note that our baseline estimates remain robust when we include data on conflicts. This does not necessarily imply that we are not measuring the effects of conflict. First, some of the conflict measures (fatalities in ACLED) are only available for African countries in our sample spanning four continents. The coefficients on weather might hence be identified by the remaining countries (see Tables D.7-D.9 where we allow the coefficients to vary by continent and obtain similar results). By the same token, the other conflict data measures conflicts where at least one party is a government, and hence excludes many violent clashes of non-governmental groups. Second, the conflict data might simply have lots of measurement issues, which would explain why our data based on observable weather is still identified.

D.4.4 Truncation of Asylum-Temperature Response Function

One potential worry is that our estimated response function extends beyond the historically observed temperatures after warming or relies heavily on observations at the upper or lower end to define the nonlinear response function. We address this by truncating the response function such that it is forced to flatten out in Figure D.6 at various thresholds ranging between 0°C to 2°C above or below the historic extrema, +2°C being our default warming scenario for the sensitivity checks in the main paper. Note how the estimated quadratic response function is nearly unchanged, except for the parts at the boundaries that by assumption are forced to flatten out.

The effect of this truncation on predicted climate impact is shown in Figure D.7. The x-axis gives the truncation around the boundaries, while the y-axis gives the predicted climate change impacts on asylum applications for uniform climate change scenarios ranging from +1°C to +5°C. As shown by the flat green line, the predicted impacts are virtually unchanged for our baseline +2°C as function of the width of the truncation. Only for significant warming (e.g., +5°C) does the truncation reduce the estimated impacts, which is not surprising, as continued warming shifts countries to a part of the response function that is assumed to be non-responsive, i.e., where warming has no effect. Our main finding that asylum applications will increase with climate change is reduced in magnitude but not reversed.

D.4.5 Predicted Changes in Asylum Applications under NEX-GDDP

We simulate the effects of future climate change again using 1000 bootstrap runs of the parameters estimated using the University of Delaware data for the years 2000-2014. All parameter estimates except for temperature coefficients are lumped together as an source-country fixed effect. We calculate the predicted asylum applications for a 30-year future interval as well as the baseline 1976-2005 using temperature of the NEX-GDDP data set. The reason why we predict applications for a comparable

30-year baseline (1976-2005) is that the weather data (Delaware and NEX-GDDP) are on different geographic grids and we evaluate the nonlinear transformation (square or spline) before aggregating over all grid cells. Part of the difference might hence be due to aggregation bias (extremes are less frequently observed if the temperature data is averaged over a larger area). Comparing outcomes under future temperature outcomes in the NEX-GDDP data to historic baselines on the same grids avoids such differences.

Changes in the medium term are shown in Figure D.8, while changes in the long-term are shown in Figure D.9. Each figure shows the predicted changes in percent for the 21 climate models as well as the model mean averaged over all 21 models. The distribution within each climate model is solely due to the parameter uncertainty of the statistical model linking weather to asylum applications as the NEX-GDDP data gives only one ensemble member for each model. The black vertical line give the mean impact for each model as well as model average: a +28% increase under RCP 4.5 and a 188% increase under RCP 8.5 by the end of he century (2070-2099). The respective probabilities that asylum applications will increase are 92% and 99%.

D.4.6 Heterogeneity of Asylum-Temperature Response Function

Tables D.5 and D.6 allow the response function for the quadratic and spline models, respectively, to differ by observable differences across countries. Column (1) replicates the baseline estimate. We split the data based on observable characteristics in roughly equal halves: Column (2a) looks at the subgroup of countries that are considered more corrupt by Transparency International, (2b) at countries that are within 23.5 degrees of the Equator, (2c) at countries that have a population of at least 15 million, column (2d) at countries where at least 42% of the population work in agriculture, and column (2e) at countries that are at least 2700 km from the European Union. The p-value for whether the temperature coefficients are different for the subgroup are generally not significant.

Columns (4a)-(4b) in Tables D.7 - D.9 allow the response function for the quadratic model to differ by the continent of origin (Africa, Asia, Europe, and America). The standard errors are generally larger as we are estimating more parameters, but the point estimates are comparable with a slightly higher overall impact (an increase of 34% under RCP4.5 by the end of the century compared to +29% if we pool countries from the four continents in one regression in columns (3a)-(3b) in Table D.9). We therefore use the pooled estimate in our baseline regression as it is more conservative.

D.5 Additional figures

Figure D.1: Maize Growing Area

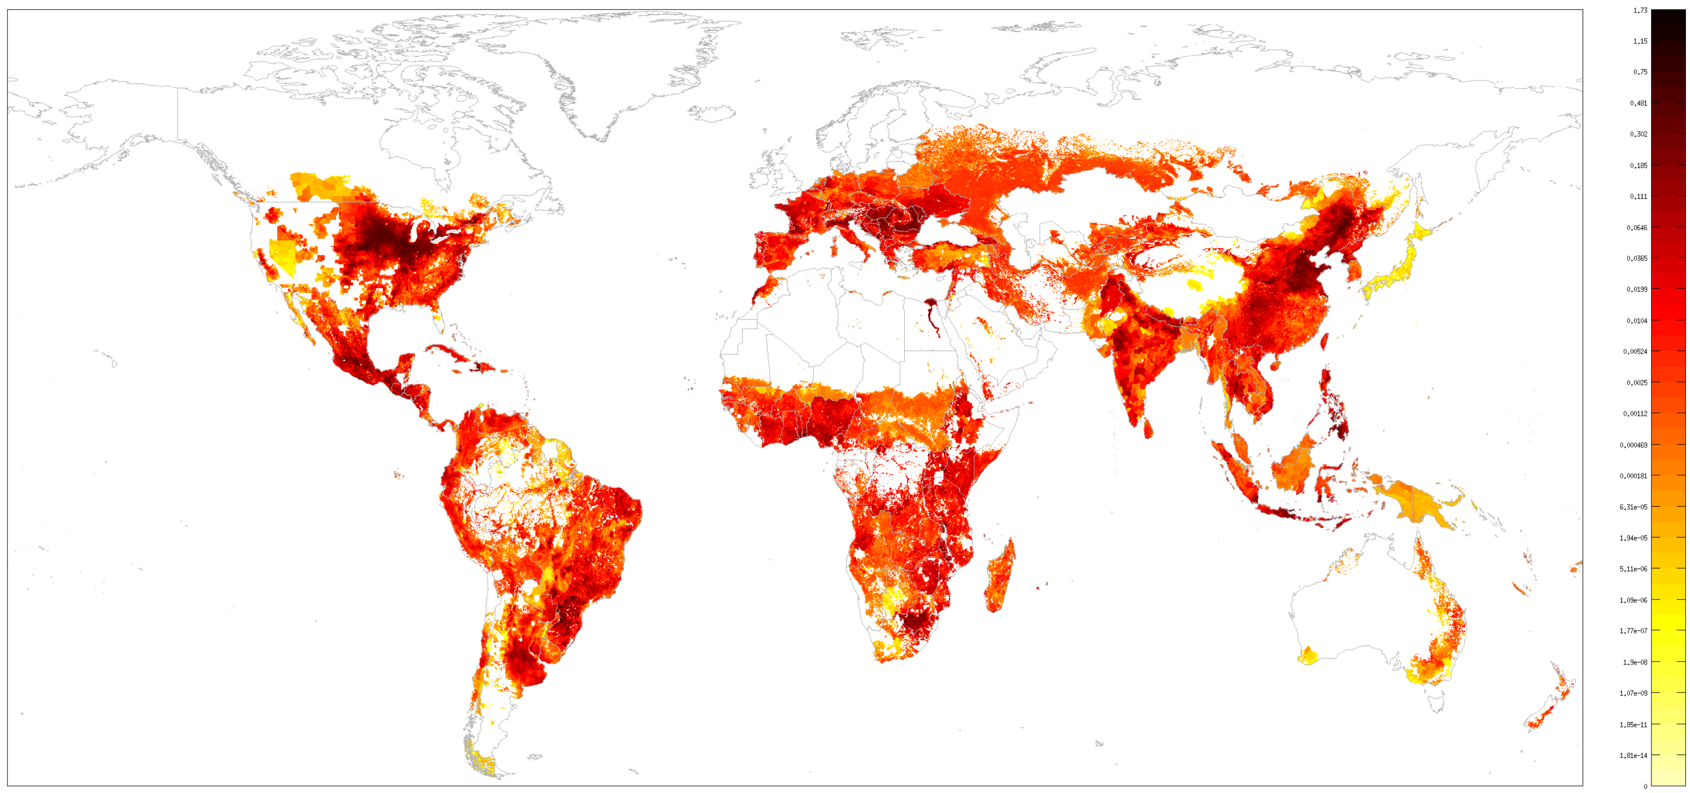

Notes: Figure displays the fraction of each grid cell in (Sacks et al., 2010) used to grow maize (note the nonlinear scale on the right). Numbers greater than 1 indicate double cropping.

Figure D.2: Average Temperature Over Maize Growing Area and Season

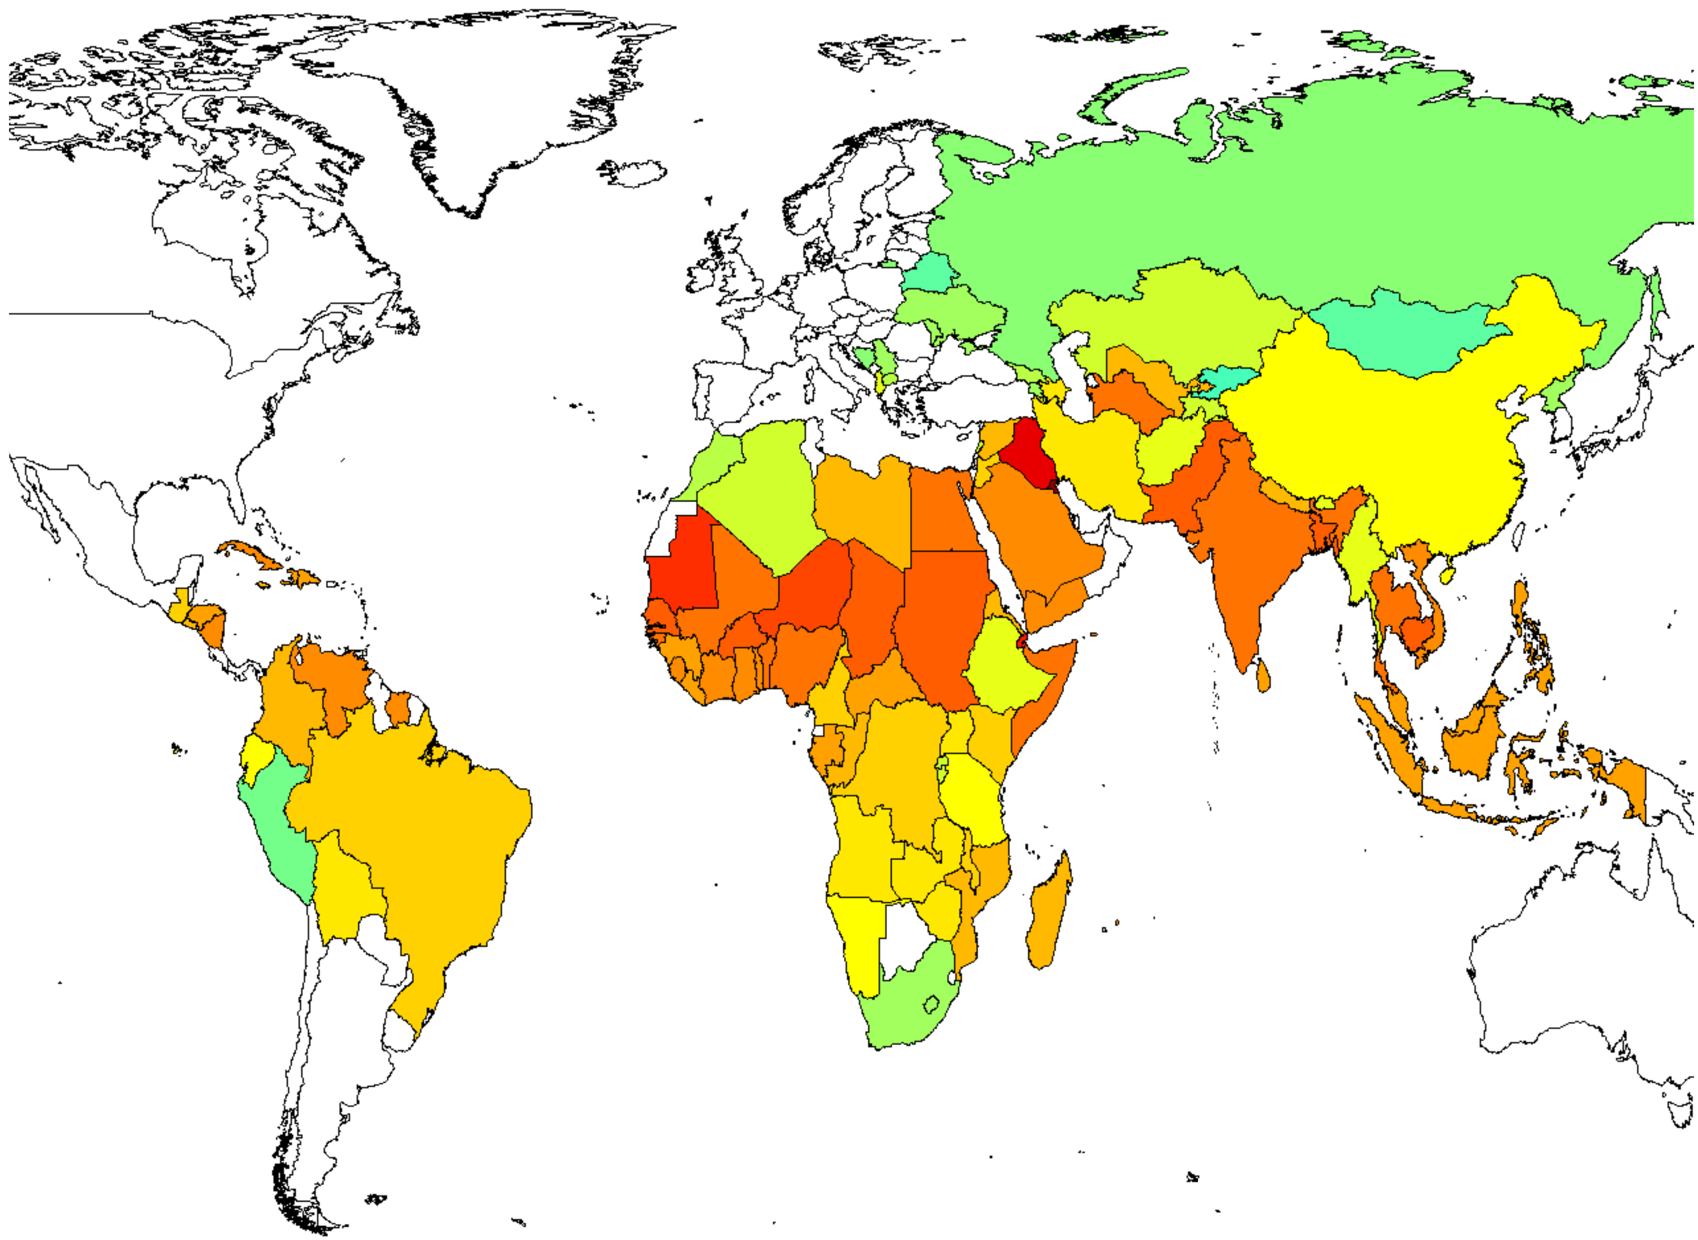

Notes: Figure displays the average temperature in a country. Monthly temperature data from the University of Delaware was weighted by the amount of maize growing area in each grid cell (see Figure D.1) for the months that maize is grown.

Figure D.3: Average Temperature Over Country

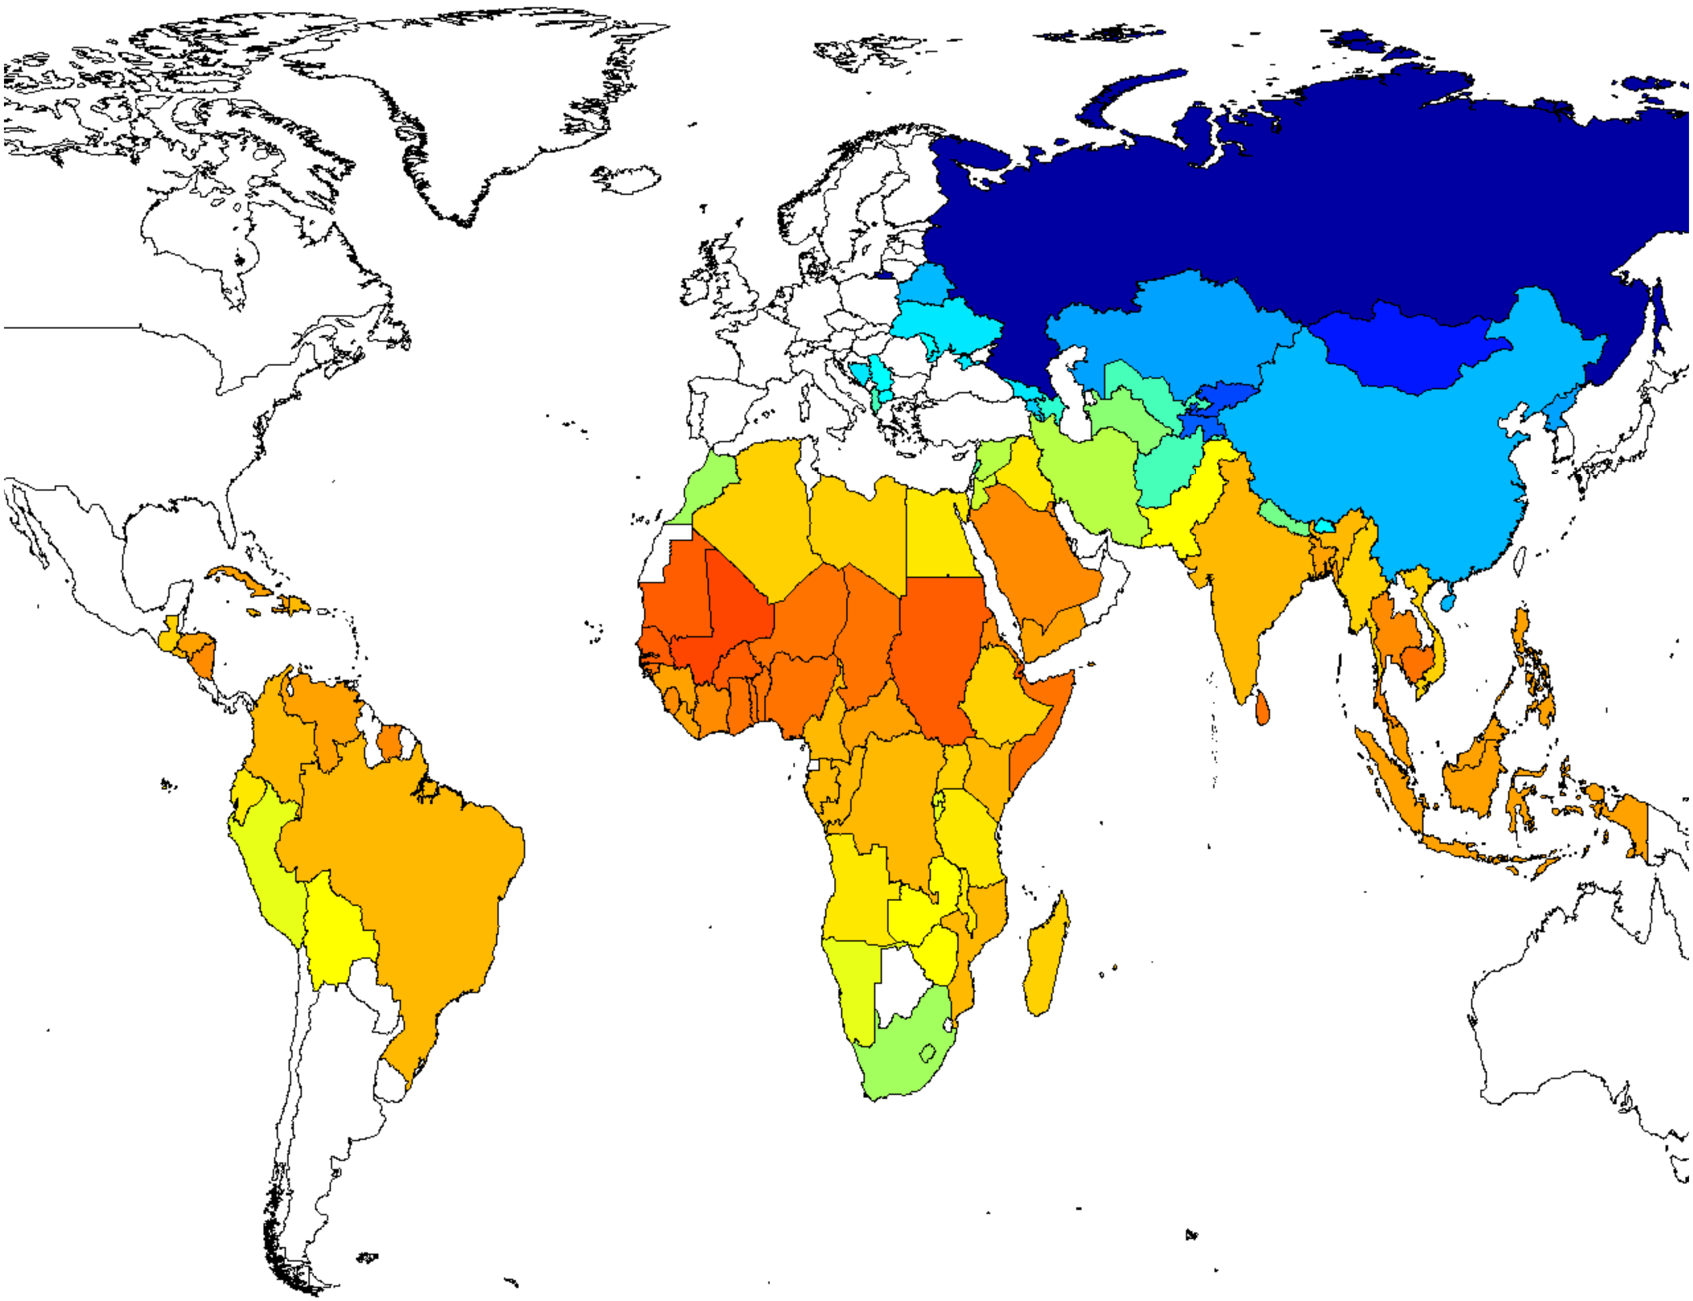

Notes: Figure displays the average temperature in a country. It averages monthly temperature from the University of Delaware Climate Data over all grid cells in a country for the entire year.

Figure D.4: Histogram of Temperature and Application Anomalies

Panel A: Temperature Anomalies

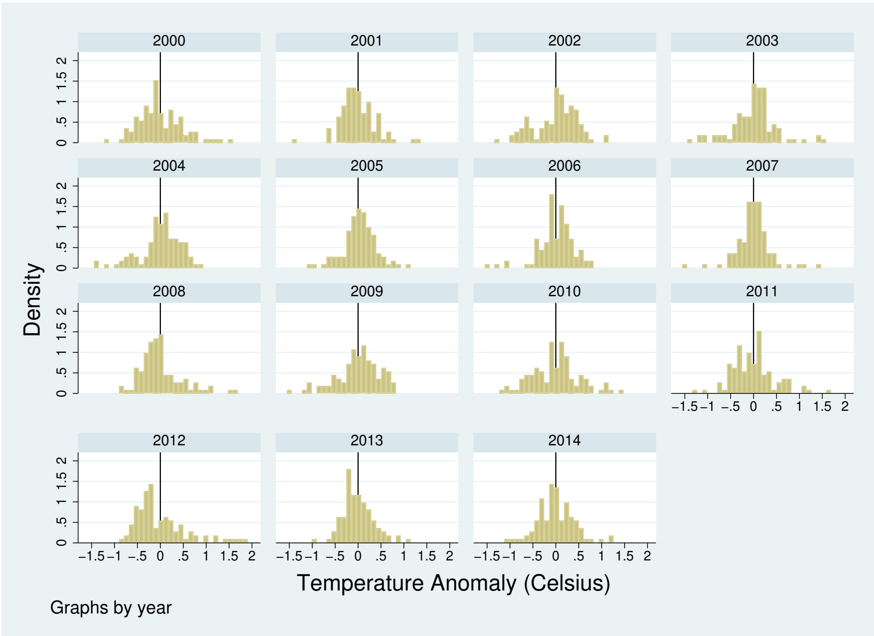

Panel B: Application Anomalies

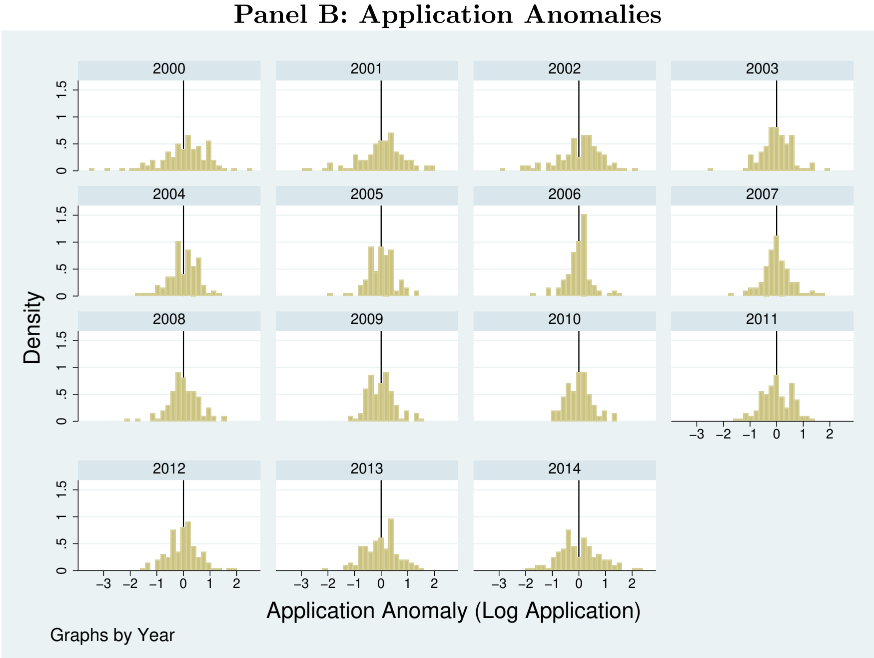

Notes: Graphs display the distribution of anomalies after removing country and year fixed effects, i.e., the variation that is used in our model to estimate the coefficients of interest. We use the 103 countries in our baseline model that are listed in Tables D.7-D.9. The top graph shows temperature anomalies (Celsius) while the bottom graph shows application anomalies (log application). The top graph excludes one observation (-3) for clearer exposition.

Figure D.5: Temperature and Application Anomalies for Top 4 Countries

Panel A: Temperature Anomalies

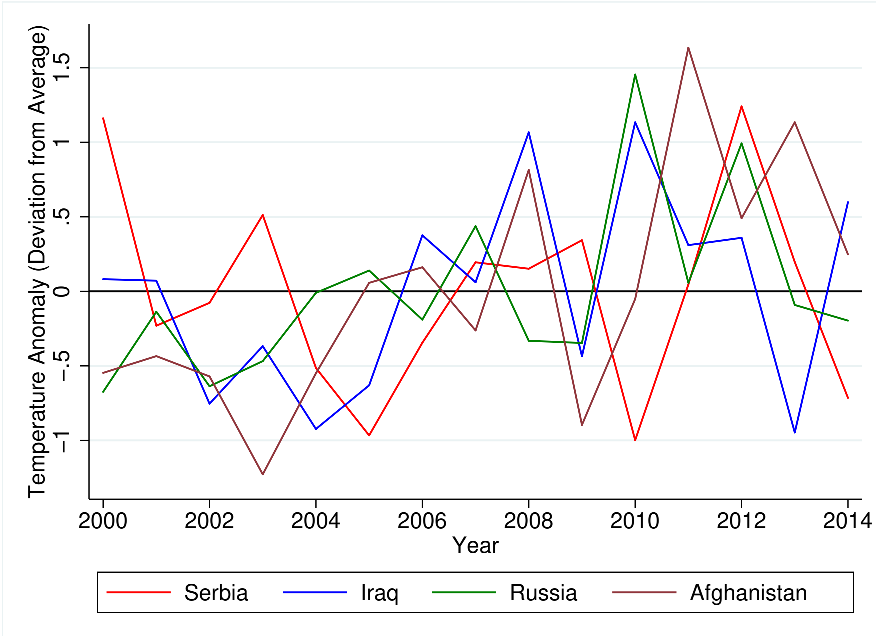

Panel B: Application Anomalies

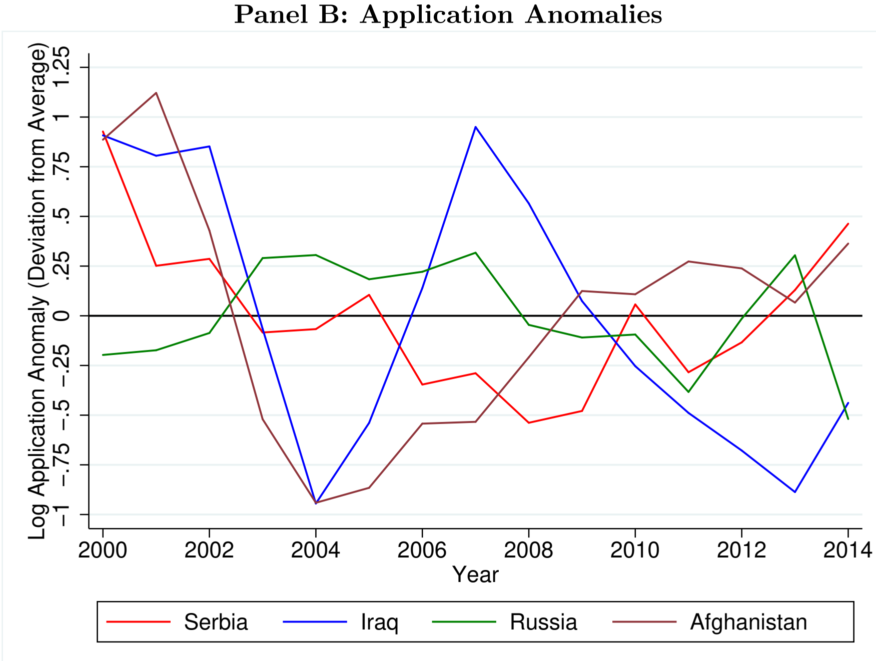

Notes: Graphs display anomalies after removing country and year fixed effects, i.e., the variation that is used in our model to estimate the coefficients of interest. The anomalies are plotted for the four countries with the largest average number of asylum applications in our sample. The top graph shows temperature anomalies (Celsius) while the bottom graph shows application anomalies (log application). The top graph excludes one observation (-3) for clearer exposition.

Figure D.6: Truncated Quadratic Response Function

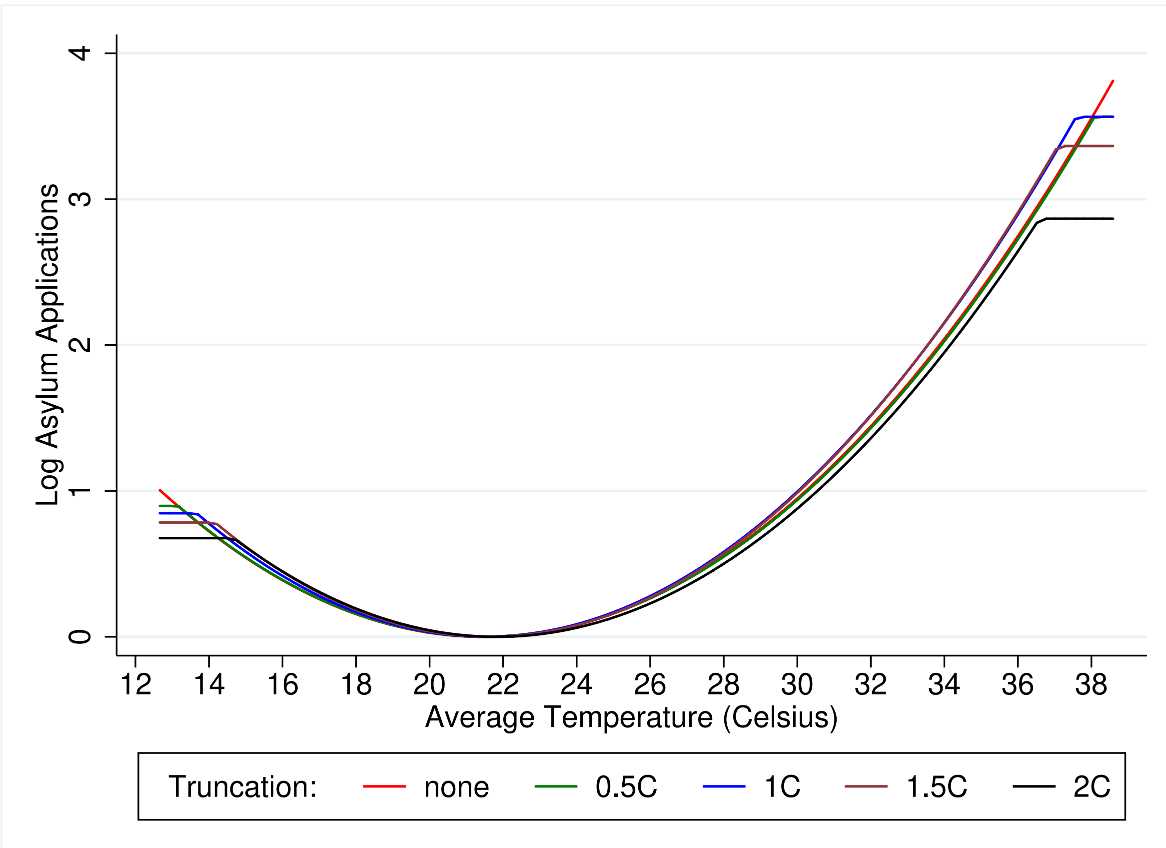

Notes: Graph displays the response functions between average weather over the maize growing season when we truncate the extremes: A truncation of x°C implies that the response function becomes non-responsive (flat) if temperatures exceed a threshold of x°C below the historic maximum or fall below a threshold of x°C above the historic minimum. The truncation is varied between 0°C and 2°C in 0.5°C intervals.

Figure D.7: Predicted Changes under 2°C Warming for Truncated Response Function

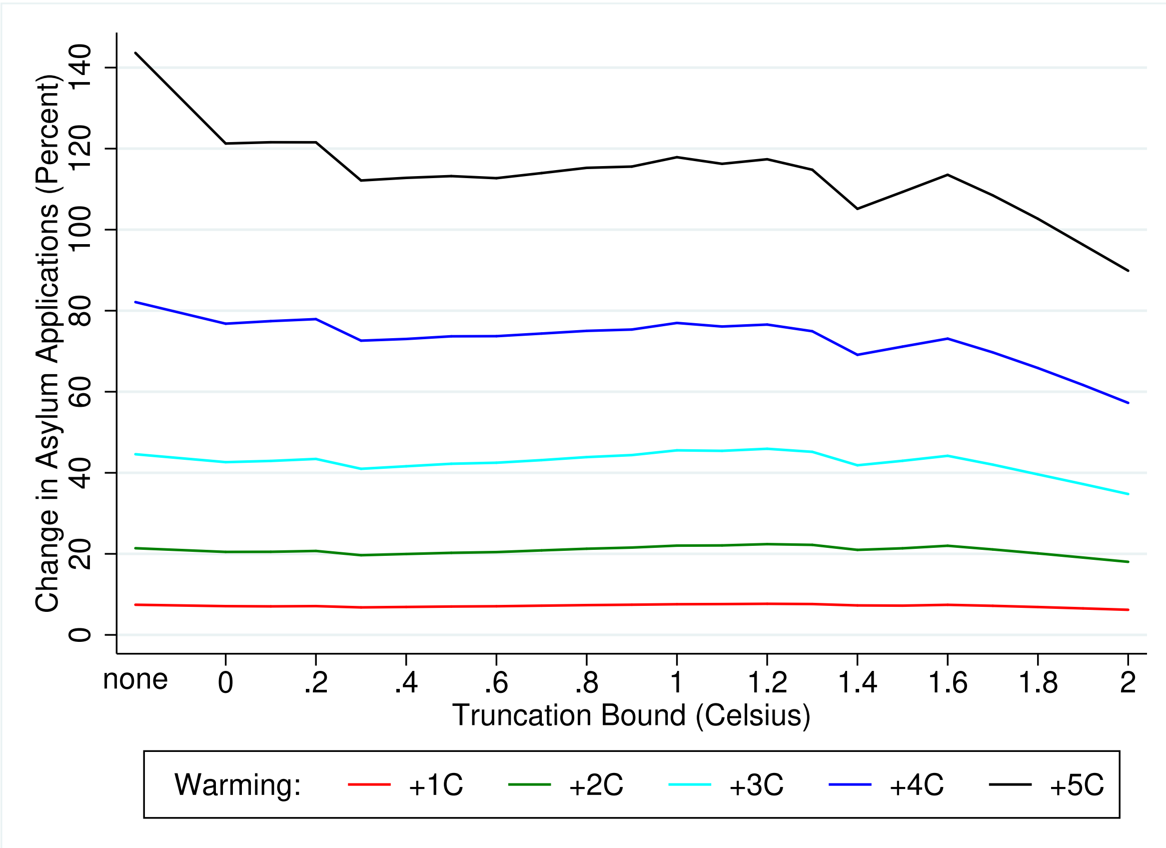

Notes: Graph displays the predicted change in asylum applications as a function of truncation at the boundaries. The leftmost point of each line displays predicted changes without any truncation, i.e., the function is evaluated after a 1, 2, 3, 4 or 5°C warming, even if these new temperatures are warmer than anything that was observed in the past. A truncation of 0°C implies that temperatures are not allowed to be warmer than the highest observed temperature record. A truncation of x°C implies that the response is flat if temperatures exceed a threshold of x°C below the historic maximum or fall below a threshold of x°C above the historic minimum. The truncation is varied between 0°C and 2°C in 0.1°C interval steps.

Figure D.8: Predicted Changes in Asylum Applications in Medium Term (2030-2059)

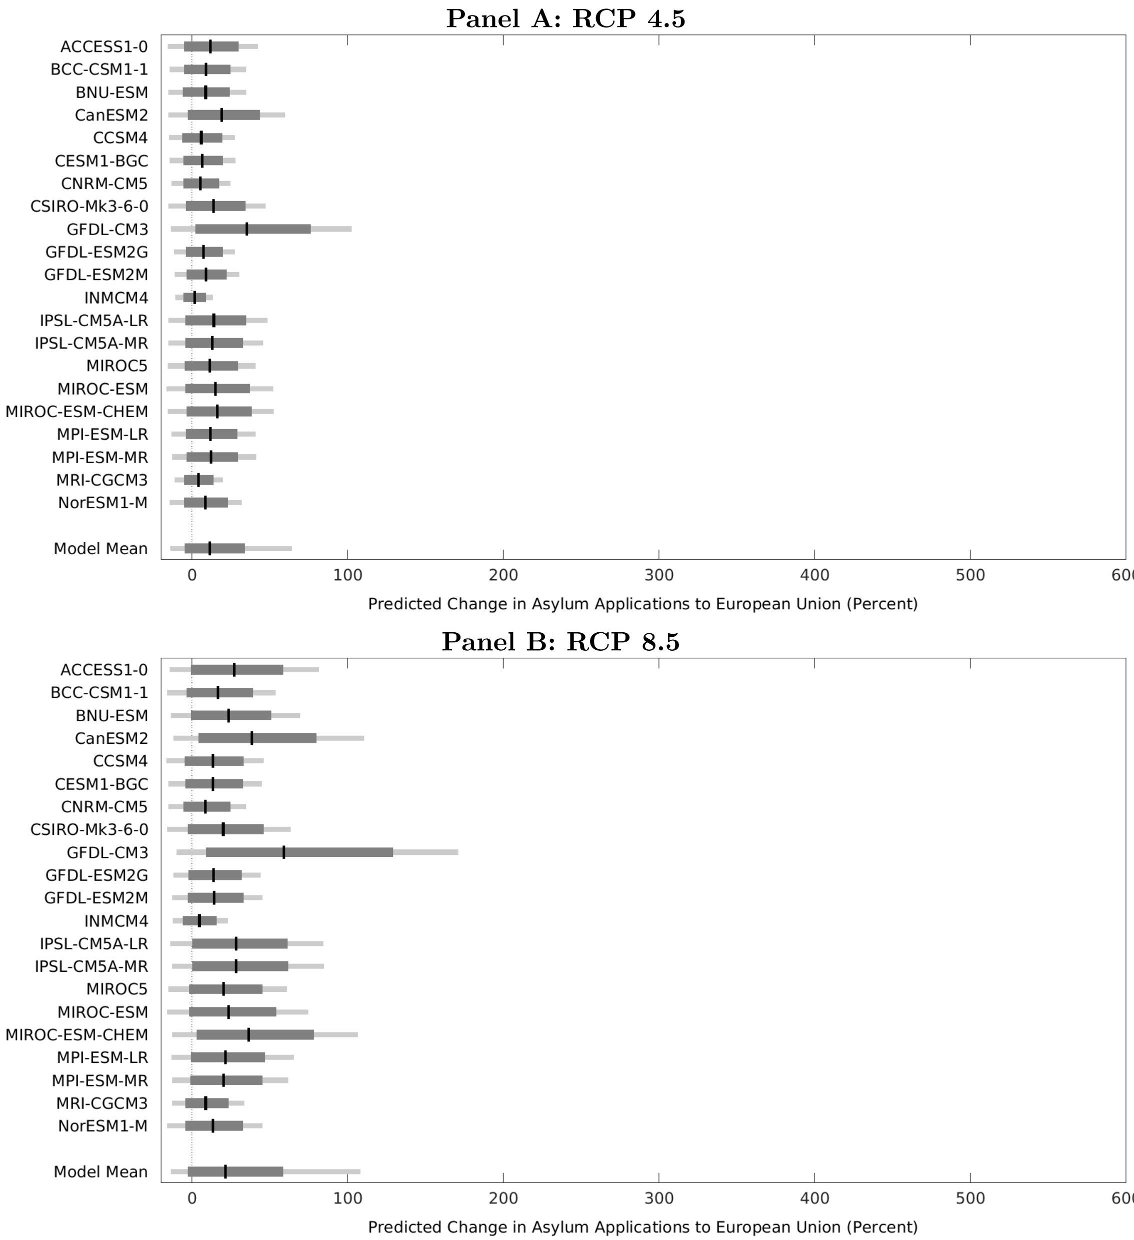

Notes: Graphs display the predicted change in asylum applications in the medium term: it is the change in predicted asylum applications using the NEX-GDDP climate projections in 2030-2059 (averaged over the 30 years) compared to the predicted applications in the baseline period 1976-2005 (averaged over the 30 years). The top graph shows the results under the RCP4.5 scenario, the bottom graph under RCP8.5 scenario for the 21 models in NEX-GDDP as well as the model average over all 21 models. The solid vertical black line gives the average impact when bootstrapping the model parameter uncertainty under the spline model, while the dark grey bars give the 90% confidence interval and the light grey bars the 99% confidence interval.

Figure D.9: Predicted Changes in Asylum Applications in Long Term (2070-2099)

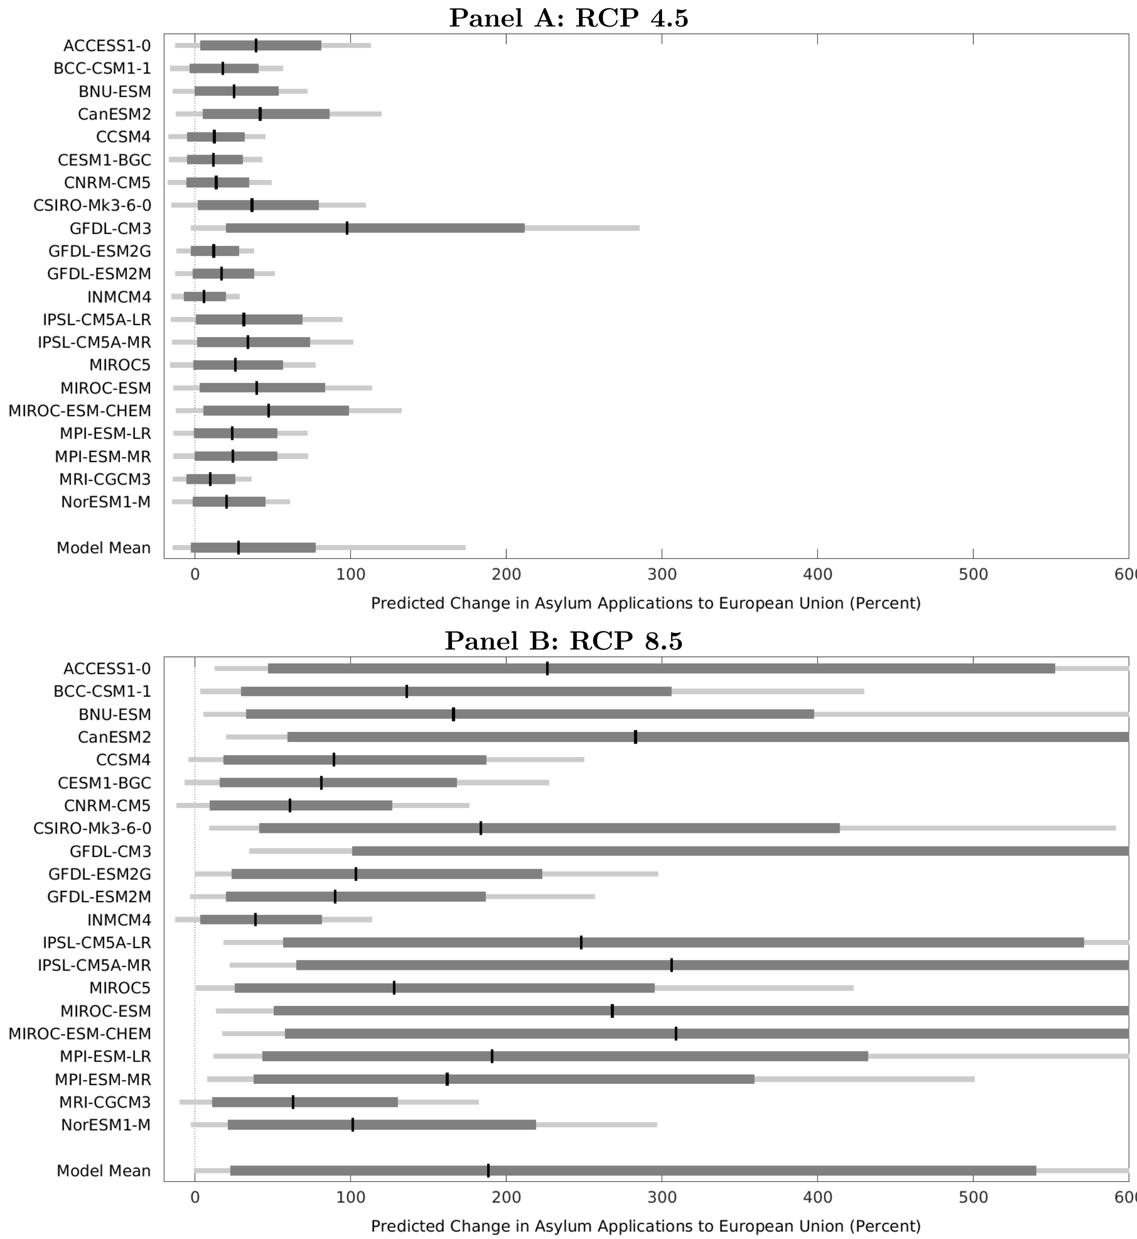

Notes: Graphs display the predicted change in asylum applications in the medium term: it is the change in predicted asylum applications using the NEX-GDDP climate projections in 2070-2099 (averaged over the 30 years) compared to the predicted applications in the baseline period 1976-2005 (averaged over the 30 years). The top graph shows the results under the RCP4.5 scenario, the bottom graph under RCP8.5 scenario for the 21 models in NEX-GDDP as well as the model average over all 21 models. The solid vertical black line gives the average impact when bootstrapping the model parameter uncertainty under the spline model, while the dark grey bars give the 90% confidence interval and the light grey bars the 99% confidence interval (truncated at +600%).

D.6 Additional tables

Table D.1: Asylum Applications to the EU as a Function of Weather

	Quadratic in Average Temperature				Spline	
	(1a)	(1b)	(2a)	(2b)	(3a)	(3b)
Average Temperature - 1st term	-0.539***	-0.706*	-0.556***	-0.799**	-0.187	-0.356
	(0.187)	(0.368)	(0.193)	(0.381)	(0.114)	(0.300)
Average Temperature - 2nd term	0.013***	0.016**	0.012***	0.016**	1.034	1.973
	(0.004)	(0.007)	(0.004)	(0.007)	(0.728)	(1.562)
Average Temperature - 3rd term					-2.877	-5.541
					(2.213)	(4.626)
Average Temperature - 4th term					4.243	7.069
					(2.876)	(5.872)
Precipitation	0.244	1.132			0.220	1.137
	(0.530)	(1.325)			(0.526)	(1.319)
Precipitation Squared	0.104	0.038			0.103	0.027
	(0.177)	(0.350)			(0.175)	(0.352)
Observations	1545	1545	1545	1545	1545	1545
Countries	103	103	103	103	103	103
P-value Temperature	.00937	.00831	.018	.0284	.00295	.00021
Optimal Temperature	21.4	22.5	22.7	25.8	19.9	19.9
Lags Included	0	2	0	2	0	2

Notes: Table regresses annual log asylum applications to EU countries on a quadratic or restricted cubic spline (with 5 evenly spaced knots between 15 and 35°C) in the average temperature and total precipitation during the maize growing season and area in the source country. All columns include source-country and year fixed effects. Columns (b) furthermore include two lags (the reported coefficients are the sum of the contemporaneous and lagged effects). Stars indicate significance: ***, **, and * stand for significance at the 1%, 5%, and 10% level, respectively. The footer also gives the optimal temperature, i.e., the average temperature in a country when fewest people seek asylum in Europe as well as the p-value for the joint significance of all temperature variables. Finally, it shows both the overall R-squared, i.e., the fraction of the variation in asylum applications that is explained by the full model including fixed effects, as well as the partial R-squared when all variables except for the weather variables are partialled out.

Table D.2: Asylum Applications to the EU: Including Population-Weighted Weather

	Quadratic Model			Spline Model		
	(1a)	(1b)	(1c)	(2a)	(2b)	(2c)
Avg. Temp. (Maize) - 1st term	-0.539***		-0.475*	-0.187		-0.253*
	(0.187)		(0.274)	(0.114)		(0.143)
Avg. Temp. (Maize) - 2nd term	0.013***		0.010	1.034		2.246**
	(0.004)		(0.006)	(0.728)		(0.986)
Avg. Temp. (Maize) - 3rd term				-2.877		-7.538**
				(2.213)		(3.280)
Avg. Temp. (Maize) - 4th term				4.243		10.890**
				(2.876)		(4.655)
Precipitation (Maize)	0.244		-0.803	0.220		-0.819
	(0.530)		(0.651)	(0.526)		(0.644)
Precip. (Maize) Squared	0.104		0.611**	0.103		0.606**
	(0.177)		(0.243)	(0.175)		(0.240)
Avg. Temp. (Pop) - 1st term		-0.359*	-0.079		-0.066	0.059
		(0.200)	(0.239)		(0.100)	(0.132)
Avg. Temp. (Pop) - 2nd term		0.009*	0.003		-0.009	-1.355
		(0.005)	(0.006)		(0.721)	(0.932)
Avg. Temp. (Pop) - 3rd term					0.483	5.145
					(2.384)	(3.295)
Avg. Temp. (Pop) - 4th term					0.145	-7.170
					(3.441)	(5.006)
Precipitation (Pop)		0.001*	0.001**		0.001*	0.001*
		(0.000)	(0.001)		(0.000)	(0.001)
Precip. (Pop) Squared		-0.000	-0.000***		-0.000	-0.000***
		(0.000)	(0.000)		(0.000)	(0.000)
Observations	1545	1545	1545	1545	1545	1545
Countries	103	103	103	103	103	103
P-val. Temp (Maize)	.00937		.227	.00295		.119
P-val. Temp (Pop)		.14	.846		.15	.468
Optimal Temp (Maize)	21.4		23.6	19.9		18.9
Optimal Temp (Pop)		19.7	12.8		24.5	25.0

Notes: Table regresses annual log asylum applications to EU countries on a quadratic or restricted cubic spline (with 5 evenly spaced knots between 15 and 35°C) in the average temperature and total precipitation during the maize growing season in the source country. All columns include source-country and year fixed effects. Columns (a) replicate the baseline estimates. Columns (b) includes the same weather measures that are now weighted by the population in each grid instead of the maize growing area in a grid. We average weather over the average maize season of all grids in a country to keep the temporal span underlying the weather average comparable, but simply change how grids within a country are weighted. Columns (c) include both the maize and population weighted weather. Stars indicate significance: ***, **, and * stand for significance at the 1%, 5%, and 10% level, respectively. The footer also gives the optimal temperature, i.e., the average temperature in a country when fewest people seek asylum in Europe as well as the p-value for the joint significance of all temperature variables.

Table D.3: Asylum Applications to the EU: Sensitivity to Weather Data Set

	Quadratic Model			Spline Model		
	(1a)	(1b)	(1c)	(2a)	(2b)	(2c)
Average Temperature - 1st term	-0.539***	-0.513**	-0.557**	-0.187	-0.178	0.028
	(0.187)	(0.210)	(0.212)	(0.114)	(0.121)	(0.029)
Average Temperature - 2nd term	0.013***	0.012**	0.011**	1.034	1.150	-0.540
	(0.004)	(0.005)	(0.004)	(0.728)	(0.743)	(0.606)
Average Temperature - 3rd term				-2.877	-3.577	1.554
				(2.213)	(2.215)	(2.005)
Average Temperature - 4th term				4.243	5.710**	-1.955
				(2.876)	(2.799)	(3.056)
Precipitation	0.244	0.699	0.555	0.220	0.656	0.634
	(0.530)	(0.615)	(0.622)	(0.526)	(0.610)	(0.628)
Precipitation Squared	0.104	-0.154	-0.121	0.103	-0.150	-0.139
	(0.177)	(0.168)	(0.168)	(0.175)	(0.166)	(0.170)
Observations	1545	1260	1260	1545	1260	1260
Countries	103	90	90	103	90	90
P-value Temperature	.00937	.0336	.0345	.00295	.00108	.547
Optimal Temperature	21.4	21.4	24.3	19.9	19.5	40.0
Weather Data Source	Delaware	Delaware	Berkeley	Delaware	Delaware	Berkeley

Notes: Table regresses annual log asylum applications to EU countries on a quadratic or restricted cubic spline (with 5 evenly spaced knots between 15 and 35°C) in the average temperature and total precipitation during the maize growing season in the source country. All columns include source-country and year fixed effects. Columns (a) use the baseline weather data from the University of Delaware. Columns (b) uses the same weather data as columns (a) but restricts the set of countries and years to match what is available in Berkeley Earth. Columns (c) use the same countries and years (2000-2013) as columns (b) but replaces the average monthly temperature data with daily temperature data from the Berkeley Earth. (Note that the precipitation data is still based on the monthly data from the University of Delaware to keep other things constant). Stars indicate significance: ***, **, and * stand for significance at the 1%, 5%, and 10% level, respectively. The footer also gives the optimal temperature, i.e., the average temperature in a country when fewest people seek asylum in Europe as well as the p-value for the joint significance of all temperature variables.

Table D.4: Asylum Applications to the EU: Including Measures of Conflict

	Quadratic Model				Spline Model			
	(1a)	(1b)	(2a)	(2b)	(3a)	(3b)	(4a)	(4b)
Avg. Temp. - 1st term	-0.539***	-0.505**	-0.556***	-0.520**	-0.187	-0.188	-0.208*	-0.207*
	(0.187)	(0.193)	(0.193)	(0.199)	(0.114)	(0.116)	(0.113)	(0.115)
Avg. Temp. - 2nd term	0.013***	0.012***	0.012***	0.011**	1.034	1.160	1.056	1.182
	(0.004)	(0.004)	(0.004)	(0.004)	(0.728)	(0.724)	(0.733)	(0.728)
Avg. Temp. - 3rd term					-2.877	-3.448	-3.122	-3.689*
					(2.213)	(2.200)	(2.239)	(2.222)
Avg. Temp. - 4th term					4.243	5.215*	4.889	5.849**
					(2.876)	(2.860)	(2.956)	(2.934)
Precipitation	0.244	0.216			0.220	0.189		
	(0.530)	(0.525)			(0.526)	(0.520)		
Precipitation Squared	0.104	0.107			0.103	0.105		
	(0.177)	(0.177)			(0.175)	(0.173)		
Minor conflict (Uppsala)		0.146**		0.145**		0.144**		0.143**
		(0.064)		(0.064)		(0.064)		(0.064)
Major conflict (Uppsala)		0.413***		0.423***		0.425***		0.437***
		(0.134)		(0.136)		(0.137)		(0.138)
Fatalties (ACLED)		0.158**		0.157**		0.161**		0.161**
		(0.071)		(0.072)		(0.070)		(0.071)
Fatalties Squared		-0.012**		-0.012**		-0.012**		-0.012**
		(0.006)		(0.006)		(0.006)		(0.006)
Observations	1545	1545	1545	1545	1545	1545	1545	1545
Countries	103	103	103	103	103	103	103	103
P-value Temperature	.00937	.0193	.018	.0361	.00295	.00122	.00177	.000725
Optimal Temperature	21.4	21.3	22.7	22.7	19.9	19.7	20.1	19.8

Notes: Table regresses annual log asylum applications to EU countries on a quadratic or restricted cubic spline (with 5 evenly spaced knots between 15 and 35°C) in the average temperature and total precipitation during the maize growing season in the source country. All columns include source-country and year fixed effects. Columns (a) replicate the baseline estimates. Columns (b) include the number of minor and major conflicts a country is involved in (Uppsala Conflict Data Program) as well as a quadratic in the fatalities from political violence in thousands (Armed Conflict Location and Event Data Project). Stars indicate significance: ***, **, and * stand for significance at the 1%, 5%, and 10% level, respectively. The footer also gives the optimal temperature, i.e., the average temperature in a country when fewest people seek asylum in Europe as well as the p-value for the joint significance of all temperature variables.

Table D.5: Heterogeneity of Asylum Applications to the EU - Quadratic Model

	(1)	(2a)	(2b)	(2c)	(2d)	(2e)
Average Temperature	-0.539***	-0.608***	-0.561***	-0.521**	-0.611***	-0.586***
	(0.187)	(0.195)	(0.174)	(0.212)	(0.229)	(0.200)
Avg. Temperature Squared	0.013***	0.015***	0.014***	0.013***	0.016***	0.015***
	(0.004)	(0.004)	(0.004)	(0.005)	(0.004)	(0.004)
Avg Temp x I_{subgroup}		-0.343	-0.575	-0.483	0.060	-0.092
		(0.257)	(0.540)	(0.343)	(0.236)	(0.310)
Avg Temp Squared x I_{subgroup}		0.007	0.011	0.011	-0.002	0.001
		(0.006)	(0.011)	(0.007)	(0.006)	(0.008)
Precipitation	0.244	0.552	0.300	0.976	1.505	0.527
	(0.530)	(0.691)	(0.556)	(0.642)	(1.030)	(0.499)
Precipitation Squared	0.104	0.124	0.024	-0.039	-0.183	-0.169
	(0.177)	(0.190)	(0.138)	(0.204)	(0.387)	(0.169)
Precipitation x I_{subgroup}		-0.204	0.096	-0.728	-0.821	0.285
		(0.758)	(0.782)	(0.684)	(0.555)	(0.747)
Prec. Squared x I_{subgroup}		0.192	0.158	0.227	0.349**	0.125
		(0.241)	(0.258)	(0.178)	(0.173)	(0.208)
Observations	1545	1545	1545	1545	1545	1545
Observations Subgroup		765	855	765	780	765
Countries	103.0	103.0	103.0	103.0	103.0	103.0
Optimal Temperature	21.4	19.8	19.6	20.4	19.2	20.0
Optimal Temp. Subgroup		21	22.5	21.2	20.5	21.8
P-value Temp. Common	.00937	.0000186	.000248	.0176	.0000432	.000165
P-value Temp. Subgroup		.411	.555	.316	.74	.761
Subset		Corrupt	Latitude	Population	Ag Labour	Distance

Notes: Column (1) of Table replicates the baseline quadratic specification (columns (1a) of Table D.1). Columns (2a)-(2e) add an interaction term of the weather variables with a dummy that is one if an source county belongs to a specific subgroup. Column (2a) looks at the subgroup of countries that are more corrupt by Transparency International, (2b) at countries that are within 23.5 degree of the equator, (2c) at countries that have a population of at least 15 million, column (2d) at countries where at least 42% of the population work in agriculture, and column (2e) at countries that are at least 2700km from the European Union. Stars indicate significance: ***, **, and * stand for significance at the 1%, 5%, and 10% level, respectively. The footer also gives the optimal temperature, i.e., the average temperature in a country when fewest people seek asylum in Europe as well as the p-value for the joint significance of all temperature variables. It also includes a separate test when the temperature variables are interacted with the subset dummy to see whether they are different.

Table D.6: Heterogeneity of Asylum Applications to the EU - Spline Model

	(1)	(2a)	(2b)	(2c)	(2d)	(2e)
Avg. Temp Spline - 1st term	-0.187	-0.207	-0.197	-0.131	-0.038	-0.343**
	(0.114)	(0.141)	(0.121)	(0.130)	(0.159)	(0.145)
Avg. Temp Spline - 2nd term	1.034	1.440*	1.245	0.994	-0.620	2.481***
	(0.728)	(0.859)	(0.850)	(0.827)	(0.884)	(0.765)
Avg. Temp Spline - 3rd term	-2.877	-4.144	-3.008	-3.249	3.940	-7.002***
	(2.213)	(2.759)	(2.795)	(2.587)	(2.985)	(2.291)
Avg. Temp Spline - 4th term	4.243	5.970	3.368	5.822*	-6.714	8.393***
	(2.876)	(3.889)	(3.990)	(3.399)	(4.419)	(3.053)
Avg Temp x $I_{subgroup}$ - 1st term		-0.079	-0.684	-0.286	-0.233	0.071
		(0.216)	(0.433)	(0.257)	(0.154)	(0.103)
Avg Temp x $I_{subgroup}$ - 2nd term		-0.100	2.535	1.304	1.812	-1.372
		(1.221)	(2.035)	(1.610)	(1.148)	(0.885)
Avg Temp x $I_{subgroup}$ - 3rd term		0.806	-6.390	-2.754	-6.457*	5.348
		(3.493)	(5.577)	(4.710)	(3.388)	(3.596)
Avg Temp x $I_{subgroup}$ - 4th term		-1.154	7.049	1.413	11.063**	-13.502
		(4.101)	(6.912)	(5.489)	(4.480)	(10.679)
Precipitation	0.220	0.523	0.265	0.963	1.460	0.400
	(0.526)	(0.691)	(0.565)	(0.633)	(1.042)	(0.497)
Precipitation Squared	0.103	0.115	0.028	-0.047	-0.141	-0.131
	(0.175)	(0.196)	(0.138)	(0.199)	(0.396)	(0.168)
Precipitation x $I_{subgroup}$		-0.230	0.072	-0.706	-0.756	0.216
		(0.762)	(0.778)	(0.694)	(0.549)	(0.766)
Prec. Squared x $I_{subgroup}$		0.199	0.163	0.226	0.323*	0.153
		(0.242)	(0.257)	(0.181)	(0.168)	(0.216)
Observations	1545	1545	1545	1545	1545	1545
Observations Subgroup		765	855	765	780	765
Countries	103.0	103.0	103.0	103.0	103.0	103.0
Optimal Temperature	19.9	19.4	19.6	19.2	23.6	19.3
Optimal Temp. Subgroup		20.4	20.7	19.9	20.6	40
P-value Temp. Common	.00295	8.55e-24	.0054	8.29e-21	.00396	.000108
P-value Temp. Subgroup		.445	.455	.6	.0569	.61
Subset		Corrupt	Latitude	Population	Ag Labour	Distance

Notes: Column (1) of Table replicates the baseline spline specification (columns (3a) of Table D.1). Columns (2a)-(2e) add an interaction term of the weather variables with a dummy that is one if an source county belongs to a specific subgroup. Column (2a) looks at the subgroup of countries that are more corrupt by Transparency International, (2b) at countries that are within 23.5 degree of the equator, (2c) at countries that have a population of at least 15 million, column (2d) at countries where at least 42% of the population work in agriculture, and column (2e) at countries that are at least 2700km from the European Union. Stars indicate significance: ***, **, and * stand for significance at the 1%, 5%, and 10% level, respectively. The footer also gives the optimal temperature, i.e., the average temperature in a country when fewest people seek asylum in Europe as well as the p-value for the joint significance of all temperature variables. It also includes a separate test when the temperature variables are interacted with the subset dummy to see whether they are different.

Table D.7: Historic Applications and Predicted Changes - Africa

	Appl.		**Spline**		**Quadratic**		**Quadratic**
	(1)	(2a)	(2b)	(3a)	(3b)	(4a)	(4b)
Algeria	6766	-3%	(-40%,53%)	-10%	(-36%,20%)	-22%	(-75%,68%)
Angola	3144	24%	(-25%,97%)	13%	(-10%,52%)	-4%	(-46%,58%)
Burundi	902	16%	(-25%,73%)	4%	(-15%,29%)	-9%	(-50%,45%)
Cameroon	2473	19%	(-15%,67%)	20%	(-3%,58%)	3%	(-34%,59%)
Cape Verde	10	11%	(-20%,53%)	13%	(-5%,45%)	-1%	(-36%,45%)
Central African Republic	390	9%	(-35%,76%)	41%	(5%,109%)	21%	(-28%,106%)
Chad	487	48%	(-8%,173%)	59%	(4%,174%)	38%	(-25%,183%)
Comoros	572	4%	(-25%,42%)	22%	(1%,56%)	8%	(-23%,54%)
Congo (Brazzaville)	2504	4%	(-33%,56%)	29%	(2%,83%)	12%	(-27%,77%)
Democratic Republic of the Congo	10344	7%	(-32%,60%)	30%	(0%,84%)	11%	(-31%,80%)
Benin	241	14%	(-28%,89%)	42%	(6%,117%)	24%	(-23%,116%)
Ethiopia	1868	6%	(-22%,44%)	3%	(-17%,30%)	-11%	(-54%,48%)
Eritrea	7603	30%	(-8%,107%)	39%	(2%,117%)	18%	(-30%,114%)
Djibouti	147	69%	(7%,192%)	66%	(11%,168%)	45%	(-26%,187%)
Gabon	49	5%	(-31%,55%)	29%	(3%,80%)	13%	(-24%,76%)
Gambia	1778	24%	(-22%,109%)	48%	(8%,132%)	30%	(-23%,132%)
Ghana	2186	9%	(-34%,78%)	41%	(5%,109%)	22%	(-24%,107%)
Guinea	4783	8%	(-31%,63%)	30%	(1%,86%)	12%	(-29%,82%)
Ivory Coast	3215	8%	(-36%,78%)	41%	(5%,115%)	21%	(-25%,111%)
Kenya	571	13%	(-17%,53%)	23%	(-2%,62%)	6%	(-32%,62%)
Liberia	1035	3%	(-31%,46%)	26%	(2%,68%)	11%	(-23%,64%)
Libya	1289	14%	(-19%,65%)	28%	(-2%,82%)	8%	(-35%,78%)
Madagascar	222	10%	(-20%,49%)	15%	(-5%,45%)	0%	(-35%,50%)
Malawi	196	15%	(-24%,68%)	23%	(-4%,66%)	4%	(-38%,70%)
Mali	3851	22%	(-27%,118%)	52%	(4%,156%)	30%	(-27%,153%)
Mauritania	2923	96%	(2%,357%)	78%	(1%,253%)	57%	(-27%,275%)
Mauritius	73	7%	(-20%,40%)	15%	(-1%,43%)	3%	(-26%,43%)
Morocco	1580	-12%	(-39%,20%)	-13%	(-35%,10%)	-22%	(-71%,51%)
Mozambique	12	9%	(-28%,61%)	30%	(1%,77%)	11%	(-30%,77%)
Namibia	16	19%	(-28%,83%)	14%	(-12%,56%)	-4%	(-51%,68%)
Niger	466	102%	(13%,322%)	78%	(12%,228%)	57%	(-27%,261%)
Nigeria	11851	20%	(-23%,98%)	45%	(7%,123%)	26%	(-25%,122%)
Guinea-Bissau	370	12%	(-31%,82%)	42%	(6%,111%)	24%	(-24%,109%)
Rwanda	1280	1%	(-28%,40%)	-6%	(-25%,15%)	-17%	(-61%,43%)
Senegal	1531	22%	(-24%,110%)	48%	(7%,136%)	29%	(-23%,136%)
Sierra Leone	2414	3%	(-34%,52%)	30%	(3%,76%)	13%	(-25%,73%)
Somalia	13061	23%	(-22%,97%)	47%	(8%,117%)	29%	(-23%,120%)
South Africa	290	6%	(-34%,63%)	-6%	(-29%,21%)	-18%	(-67%,58%)
Zimbabwe	3001	18%	(-25%,77%)	19%	(-6%,63%)	1%	(-41%,66%)
Sudan	3397	58%	(-3%,190%)	66%	(4%,187%)	44%	(-27%,200%)
Togo	1558	8%	(-33%,73%)	38%	(5%,104%)	20%	(-24%,101%)
Uganda	753	17%	(-24%,71%)	17%	(-6%,55%)	0%	(-40%,59%)
Egypt	1715	35%	(-25%,164%)	69%	(9%,209%)	41%	(-30%,199%)
Tanzania	163	15%	(-22%,64%)	14%	(-7%,45%)	-2%	(-40%,52%)
Burkina Faso	490	33%	(-19%,141%)	57%	(8%,167%)	36%	(-26%,171%)
Zambia	32	17%	(-27%,78%)	25%	(-3%,75%)	5%	(-38%,76%)
Total	103600	18%	(-15%,76%)	34%	(2%,95%)	15%	(-29%,93%)

Notes: Table gives the average number of asylum applications per year for our sample period 2000-2014 in column (1). Columns (a) give the predicted changes in asylum applications in percent by the end of the century (2070-2099) compared to 1976-2005 among the 21 climate models in NEX-GDDP under the RCP4.5 scenario, while columns (b) give the 95% confidence interval. Columns (2a)-(2b) use the spline model pooling all source countries, while columns (3a)-(3b) use the quadratic model pooling all source countries. Finally, columns (4a)-(4b) use again the quadratic model but allows the coefficients of the weather variables (both average temperature and precipitation) to vary by the four continents in our data set.

Table D.8: Historic Applications and Predicted Changes - Asia

	Appl.		**Spline**		**Quadratic**		**Quadratic**
	(1)	(2a)	(2b)	(3a)	(3b)	(4a)	(4b)
Afghanistan	23429	-3%	(-31%,36%)	1%	(-29%,44%)	33%	(-12%,122%)
Bangladesh	8535	39%	(-11%,133%)	54%	(9%,143%)	72%	(14%,201%)
Bhutan	111	2%	(-25%,35%)	5%	(-17%,36%)	26%	(-6%,82%)
Myanmar	546	3%	(-24%,38%)	1%	(-19%,30%)	22%	(-7%,73%)
Cambodia	116	34%	(-11%,121%)	49%	(8%,130%)	65%	(13%,182%)
Sri Lanka	9086	4%	(-27%,42%)	23%	(2%,60%)	37%	(6%,92%)
China	10145	9%	(-27%,57%)	9%	(-16%,48%)	34%	(-4%,106%)
Occupied Palestinian Territory	1386	31%	(-28%,116%)	14%	(-9%,54%)	35%	(-0%,106%)
India	6295	32%	(-9%,109%)	46%	(7%,119%)	64%	(12%,174%)
Indonesia	62	7%	(-23%,46%)	24%	(2%,63%)	37%	(7%,94%)
Iran	12471	34%	(-10%,112%)	27%	(-8%,90%)	60%	(3%,176%)
Iraq	25513	399%	(72%,1482%)	257%	(34%,955%)	322%	(41%,1274%)
Kazakhstan	1034	-0%	(-33%,40%)	-3%	(-31%,33%)	25%	(-14%,98%)
Jordan	341	28%	(-30%,114%)	23%	(-8%,85%)	51%	(2%,163%)
North Korea	176	-2%	(-37%,46%)	-13%	(-37%,14%)	10%	(-20%,56%)
Kuwait	138	490%	(53%,1933%)	352%	(33%,1431%)	411%	(34%,1723%)
Kyrgyzstan	698	-28%	(-62%,9%)	-39%	(-69%,-4%)	-16%	(-52%,31%)
Lebanon	1673	4%	(-30%,51%)	-2%	(-25%,29%)	22%	(-11%,80%)
Malaysia	53	5%	(-30%,55%)	30%	(3%,78%)	45%	(8%,117%)
Mongolia	1981	-31%	(-67%,7%)	-29%	(-58%,1%)	-8%	(-40%,36%)
Nepal	852	13%	(-16%,58%)	21%	(-6%,66%)	44%	(2%,123%)
Pakistan	13636	112%	(21%,350%)	87%	(12%,268%)	122%	(20%,400%)
Philippines	145	8%	(-24%,51%)	27%	(3%,70%)	40%	(7%,105%)
Saudi Arabia	31	23%	(-29%,110%)	54%	(5%,141%)	83%	(13%,224%)
Vietnam	3174	32%	(-9%,107%)	46%	(7%,120%)	65%	(12%,176%)
Syria	17437	23%	(-26%,109%)	47%	(2%,141%)	79%	(10%,237%)
Tajikistan	196	-9%	(-38%,26%)	-4%	(-32%,37%)	28%	(-15%,120%)
Thailand	42	44%	(-5%,148%)	54%	(9%,146%)	71%	(13%,204%)
Turkmenistan	89	33%	(-35%,178%)	79%	(8%,238%)	120%	(17%,378%)
Uzbekistan	943	13%	(-38%,95%)	43%	(-1%,135%)	78%	(8%,236%)
Yemen	344	25%	(-2%,77%)	28%	(-0%,81%)	47%	(5%,136%)
Total	140677	73%	(10%,239%)	62%	(5%,202%)	93%	(13%,297%)

Notes: Table gives the average number of asylum applications per year for our sample period 2000-2014 in column (1). Columns (a) give the predicted changes in asylum applications in percent by the end of the century (2070-2099) compared to 1976-2005 among the 21 climate models in NEX-GDDP under the RCP4.5 scenario, while columns (b) give the 95% confidence interval. Columns (2a)-(2b) use the spline model pooling all source countries, while columns (3a)-(3b) use the quadratic model pooling all source countries. Finally, columns (4a)-(4b) use again the quadratic model but allows the coefficients of the weather variables (both average temperature and precipitation) to vary by the four continents in our data set.

Table D.9: Historic Applications and Predicted Changes - Europe and America

	Appl.		**Spline**		**Quadratic**		**Quadratic**
	(1)	(2a)	(2b)	(3a)	(3b)	(4a)	(4b)
			Panel A: Europe				
Albania	5235	9%	(-35%,72%)	3%	(-26%,49%)	37%	(-54%,292%)
Azerbaijan	3312	13%	(-35%,85%)	31%	(-6%,104%)	166%	(-52%,954%)
Armenia	6809	-18%	(-50%,20%)	-18%	(-47%,18%)	-17%	(-68%,87%)
Bosnia and Herzegovina	5277	-27%	(-59%,11%)	-27%	(-54%,3%)	-36%	(-76%,26%)
Belarus	1788	-33%	(-70%,7%)	-31%	(-60%,-1%)	-43%	(-83%,17%)
Georgia	8152	-1%	(-39%,50%)	-11%	(-40%,24%)	-3%	(-61%,125%)
Moldova	3151	-7%	(-44%,52%)	-13%	(-40%,18%)	-10%	(-61%,91%)
Russia	24118	-21%	(-54%,17%)	-22%	(-50%,9%)	-27%	(-71%,46%)
Serbia	32573	-14%	(-49%,33%)	-17%	(-45%,14%)	-19%	(-65%,68%)
Ukraine	4396	-16%	(-49%,26%)	-17%	(-45%,13%)	-19%	(-65%,63%)
Macedonia	4516	10%	(-42%,100%)	-6%	(-36%,39%)	15%	(-60%,216%)
Total	99326	-14%	(-46%,26%)	-16%	(-44%,16%)	-11%	(-63%,101%)
			Panel B: America				
Bolivia	287	11%	(-16%,62%)	17%	(-10%,68%)	3%	(-57%,127%)
Brazil	86	15%	(-22%,72%)	31%	(1%,89%)	35%	(-55%,247%)
Colombia	1516	12%	(-13%,52%)	21%	(-3%,64%)	15%	(-50%,150%)
Cuba	648	13%	(-26%,77%)	38%	(5%,102%)	60%	(-55%,379%)
Dominican Republic	151	7%	(-23%,46%)	24%	(1%,62%)	25%	(-46%,173%)
Ecuador	222	-5%	(-27%,19%)	4%	(-14%,28%)	-12%	(-50%,38%)
El Salvador	156	10%	(-29%,67%)	33%	(2%,91%)	41%	(-56%,282%)
Guatemala	27	9%	(-18%,45%)	14%	(-9%,51%)	0%	(-52%,98%)
Haiti	3553	5%	(-26%,46%)	24%	(2%,62%)	26%	(-47%,176%)
Honduras	41	14%	(-24%,69%)	24%	(-2%,72%)	19%	(-52%,174%)
Jamaica	413	6%	(-25%,47%)	24%	(2%,65%)	27%	(-46%,181%)
Nicaragua	36	11%	(-26%,65%)	30%	(1%,88%)	36%	(-53%,254%)
Peru	181	-16%	(-41%,9%)	-19%	(-42%,5%)	-37%	(-83%,44%)
Suriname	15	17%	(-34%,127%)	52%	(6%,182%)	124%	(-63%,753%)
Venezuela	87	26%	(-20%,113%)	49%	(6%,136%)	90%	(-62%,565%)
Total	7418	7%	(-22%,48%)	24%	(0%,65%)	28%	(-45%,197%)
			Panel C: Africa+Asia+Europa+America				
Total	351021	28%	(-5%,99%)	29%	(-5%,96%)	34%	(-5%,123%)

Notes: Table gives the average number of asylum applications per year for our sample period 2000-2014 in column (1). Columns (a) give the predicted changes in asylum applications in percent by the end of the century (2070-2099) compared to 1976-2005 among the 21 climate models in NEX-GDDP under the RCP4.5 scenario, while columns (b) give the 95% confidence interval. Columns (2a)-(2b) use the spline model pooling all source countries, while columns (3a)-(3b) use the quadratic model pooling all source countries. Finally, columns (4a)-(4b) use again the quadratic model but allows the coefficients of the weather variables (both average temperature and precipitation) to vary by the four continents in our data set.

Conclusion

Here you can write some introductory remarks about your chapter. I like to give each sentence its own line.

When you need a new paragraph, just skip an extra line.

New Section

By using the asterisk to start a new section, I keep the section from appearing in the table of contents. If you want your sections to be numbered and to appear in the table of contents, remove the asterisk.

Bibliography

Abarcar, P. (2017). The return motivations of legal permanent migrants: Evidence from exchange rate shocks and immigrants in Australia. *Journal of Economic Behavior and Organization*, 144(December):62–77.

Adger, W. N., Puhlin, J. M., Barnett, J., Dabelko, G. D., Hovelsrud, G. K., Levy, M. A., Oswald Spring, Ú., and Vogel, C. H. (2014). Human Security. In Field, C. B., Barros, V., Dokken, D., Mach, K., Mastrandrea, M. D., Bilir, T., Chatterjee, M., Ebi, K., Estrada, Y., Genova, R., Girma, B., Kissel, E., Levy, A., MacCracken, S., Mastrandrea, P., and White, L., editors, *Climate Change 2014: Impacts, Adaptation, and Vulnerability. Part A: Global and Sectoral Aspects. Contribution of Working Group II to the Fifth Assessment Report of the Intergovernmental Panel of Climate Change*, chapter 12, pages 755–791. Cambridge University Press, Cambridge, UK.

Albouy, D., Graf, W., Kellogg, R., and Wolff, H. (2016). Climate Amenities, Climate Change, and American Quality of Life. *Journal of the Association of Environmental and Resource Economists*, 3(1):205–246.

Alix-Garcia, J. M., McIntosh, C., Sims, K. R. E., and Welch, J. R. (2013). The Ecological Footprint of Poverty Alleviation: Evidence from Mexico's Oportunidades Program. *Review of Economics and Statistics*, 95(2):417–435.

Alix-Garcia, J. M., Shapiro, E. N., and Sims, K. R. E. (2012). Forest conservation and slippage: Evidence from Mexico's national payments for ecosystem services program. *Land Economics*, 88(4):613–638.

Almond, D. and Currie, J. (2011). Killing Me Softly: The Fetal Origins Hypothesis. *Journal of Economic Perspectives*, 25(3):153–172.

Anderson, M. L. and Auffhammer, M. (2013). Pounds That Kill: The External Costs of Vehicle Weight. *Review of Economic Studies*, 81(2):535–571.

Angelucci, M. (2015). Migration and Financial Constraints: Evidence from Mexico. *The Review of Economics and Statistics*, 97(1):224–228.

Aragão, L. E. O. C. (2012). The rainforest's water pump. *Nature*, 489:217–218.

Aulong, S., Erdlenbruch, K., and Figuières, C. (2005). Un tour d'horizon des critères d'évaluation de la diversité biologique. *Économie Publique*, 16(1):3–46.

Baccini, A., Walker, W., Carvalho, L., Farina, M., Sulla-Menashe, D., and Houghton, R. A. (2017). Tropical forests are a net carbon source based on aboveground measurements of gain and loss. *Science*, 358(6360):230–234.

Bagla, P. (2010). Hardy cotton-munching pests are latest blow to GM crops. *Science*, 327(5972):1439.

Baker, N. T. and Stone, W. W. (2015). Estimated Annual Agricultural Pesticide Use for Counties of the Conterminous United States, 1992–2009. *Geological Survey Data Series 907*, pages 1–9.

Bazzi, S. (2017). Wealth Heterogeneity and the Income Elasticity of Migration. *American Economic Journal: Applied Economics*, 9(2):219–255.

Bender, D. J., Contreras, T. A., and Fahrig, L. (1998). Habitat Loss and Population Decline: A Meta-Analysis of the Patch Size Effect. *Ecology*, 79(2):517.

Bentley, J. W. (1987). Economic and Ecological Approaches to Land Fragmentation: In Defense of a Much-Maligned Phenomenon. *Annual Review of Anthropology*, 16:31–67.

BenYishay, A. and Mobarak, A. M. (2019). Social Learning and Incentives for Experimentation and Communication. *The Review of Economic Studies*, 86(3):976–1009.

Bohra-Mishra, P., Oppenheimer, M., and Hsiang, S. M. (2014). Nonlinear permanent migration response to climatic variations but minimal response to disasters. *Proceedings of the National Academy of Sciences*, 111(27):9780–9785.

Boltanski, L. and Thévenot, L. (1991). *De la justification. Les économies de la grandeur.* Gallimard, collection NRF Essais, Paris.

Boryan, C., Yang, Z., Mueller, R. M., and Craig, M. (2011). Monitoring US agriculture: the US Department of Agriculture, National Agricultural Statistics Service, Cropland Data Layer Program. *Geocarto International*, 26(5):341–358.

Bryan, G., Chowdhury, S., and Mobarak, A. M. (2014). Underinvestment in a Profitable Technology: The Case of Seasonal Migration in Bangladesh. *Econometrica*, 82(5):1671–1748.

Bulow, J. I., Geanakoplos, J. D., and Klemperer, P. D. (1985). Multimarket Oligopoly: Strategic Substitutes and Complements. *Journal of Political Economy*, 93(3):488–511.

Bunch, T. R., Gervais, J. A., Buhl, K., and Stone, D. (2012). Dicamba Technical Fact Sheet. *National Pesticide Information Center, Oregon State University Extension Services.*

Burke, M., Hsiang, S. M., and Miguel, E. (2015). Global non-linear effect of temperature on economic production. *Nature*, 527:235–239.

Burke, M., Miguel, E., Satyanath, S., Dykema, J. A., and Lobell, D. B. (2009). Warming increases the risk of civil war in Africa. *Proceedings of the National Academy of Sciences of the United States of America*, 106(49):20670–20674.

Burton, N., Shirriff, S. W., and Beckie, H. J. (2014). Response of glyphosate-resistant kochia (Kochia scoparia L. Schrad.) to alternative herbicides. *Canadian Journal of Plant Science*, 94(8):1407–1411.

Cai, R., Feng, S., Oppenheimer, M., and Pytliková, M. (2016). Climate variability and international migration: The importance of the agricultural linkage. *Journal of Environmental Economics and Management*, 79:135–151.

Carlton, D. W. and Perloff, J. M. (1994). Monopolies, Monopsonies, and Dominant Firms. In *Modern Industrial Organization*, pages 132–174. Harper Collins, 2nd edition.

Cassman, K. G. (1999). Ecological intensification of cereal production systems: Yield potential, soil quality, and precision agriculture. *Proceedings of the National Academy of Sciences of the United States of America*, 96(11):5952–5959.

Cattaneo, C. and Peri, G. (2016). The migration response to increasing temperatures. *Journal of Development Economics*, 122:127–146.

Chaplin-Kramer, R., O'Rourke, M. E., Blitzer, E. J., and Kremen, C. (2011). A meta-analysis of crop pest and natural enemy response to landscape complexity. *Ecology Letters*, 14(9):922–932.

Clemens, M. A. (2011). Economics and emigration: Trillion-dollar bills on the sidewalk? *Journal of Economic Perspectives*, 25(3):83–106.

Clemens, M. A. (2014). Does Development Reduce Migration? In Lucas, R. E. B., editor, *International Handbook on Migration and Economic Development*, chapter 6, pages 152–185. Edward Elgar Publishing Ltd, Cheltenham (UK) and Northampton, MA (USA).

Costello, C. and Polasky, S. (2008). Optimal harvesting of stochastic spatial resources. *Journal of Environmental Economics and Management*, 56(1):1–18.

Costello, C., Quérou, N., and Tomini, A. (2017). Private eradication of mobile public bads. *European Economic Review*, 94:23–44.

Cowan, R. and Gunby, P. (1996). Sprayed to Death: Path Dependence, Lock-in and Pest Control Strategies. *The Economic Journal*, 106(436):521.

Creppell, I. (2007). Reviewed Work(s): On Justification: Economies of Worth by Luc Boltanski, Laurent Thévenot. *Perspectives on Politics*, 5(1):144–146.

Cushman, S. A., McGarigal, K., and Neel, M. C. (2008). Parsimony in landscape metrics: Strength, universality, and consistency. *Ecological Indicators*, 8(5):691–703.

Dell, M., Jones, B. F., and Olken, B. A. (2012). Climate Shocks and Economic Growth: Evidence from the Last Half Century. *American Economic Journal: Macroeconomics*, 4(3):66–95.

Dell, M., Jones, B. F., and Olken, B. A. (2014). What Do We Learn from the Weather? The New Climate-Economy Literature. *Journal of Economic Literature*, 52(3):740–798.

Deryugina, T., Heutel, G., Miller, N. H., Molitor, D., and Reif, J. (2019). The Mortality and Medical Costs of Air Pollution: Evidence from Changes in Wind Direction. *NBER Working Paper*, (22796).

Deschênes, O. and Greenstone, M. (2011). Climate Change, Mortality, and Adaptation: Evidence from Annual Fluctuations in Weather in the US. *American Economic Journal: Applied Economics*, 3(4):152–185.

Dias, M., Rocha, R., and Soares, R. R. (2019). Glyphosate Use in Agriculture and Birth Outcomes of Surrounding Populations. CEEP Working Paper N. 1.

Fahrig, L., Baudry, J., Brotons, L., Burel, F., Crist, T. O., Sirami, C., Siriwardena, G., and Martin, J.-L. (2011). Functional landscape heterogeneity and animal biodiversity in agricultural landscapes. *Ecology Letters*, 14:101–112.

Farmer, B. H. (1960). On Not Controlling Subdivision in Paddy-Lands. *Transactions and Papers (Institute of British Geographers)*, 28:225–235.

Frank, E. G. (2017). The Effects of Bat Population Losses on Infant Mortality through Pesticide Use in the U.S. (JMP). pages 1–86.

Frontex (2016). Risk Analysis for 2016. Technical report, Warsaw, Poland.

Geist, H. J. and Lambin, E. F. (2002). Proximate causes and underlying driving forces of tropical deforestation. *BioScience*, 52(2):143–150.

Godard, O. (2004). De la pluralité des ordres – Les problèmes d'environnement et de développement durable à la lumière de la théorie de la justification. *Géographie, économie, société*, 6(2004/3):303–330.

Graff Zivin, J. and Neidell, M. (2014). Temperature and the allocation of time: implications for climate change. *Journal of Labor Economics*, 32(1):1–26.

Gray, C. L. and Wise, E. (2016). Country-specific effects of climate variability on human migration. *Climatic Change*, 135(3):555–568.

Griliches, Z. (1957). Hybrid Corn: An Exploration in the Economics of Technological Change. *Econometrica*, 25(4):501–522.

Guiteras, R., Levinsohn, J., and Mobarak, A. M. (2019). Demand Estimation with Strategic Complementarities: Sanitation in Bangladesh.

Han, W., Yang, Z., Di, L., and Mueller, R. M. (2012). CropScape: A Web service based application for exploring and disseminating US conterminous geospatial cropland data products for decision support. *Computers and Electronics in Agriculture*, 84:111–123.

Hansen, M. C., Potapov, P. V., Moore, R., Hancher, M., Turubanova, S. A., Tyukavina, A., Thau, D., Stehman, S. V., Goetz, S. J., Loveland, T. R., Kommareddy, A., Egorov, A. V., Chini, L., Justice, C. O., and Townshend, J. R. G. (2013). High-Resolution Global Maps of 21st-Century Forest Cover Change. *Science*, 850(November):850–854.

Heal, G., Walker, B. H., Levin, S. A., Arrow, K. J., Dasgupta, P., Daily, G. C., Ehrlich, P. R., Maler, K.-G., Kautsky, N., Lubchenco, J., Schneider, S. H., and Starrett, D. (2004). Genetic diversity and interdependent crop choices in agriculture. *Resource and Energy Economics*, 26(2):175–184.

Hendrix, C. S. and Salehyan, I. (2012). Climate change, rainfall, and social conflict in Africa. *Journal of Peace Research*, 49(1):35–50.

Henry, S., Schoumaker, B., and Beauchemin, C. (2004). The impact of rainfall on the first out-migration: A multi-level event-history analysis in Burkina Faso. *Population and Environment*, 25(5):423–460.

Hsiang, S. M., Burke, M., and Miguel, E. (2013). Quantifying the influence of climate on human conflict. *Science*, 341(6151):1235367.

Hsiang, S. M., Burke, M., Miguel, E., Meng, K. C., and Cane, M. A. (2015). Analysis of statistical power reconciles drought-conflict results in Africa.

International Union for the Conservation of Nature (2015). Conservation successes overshadowed by more species declines – IUCN Red List update — IUCN.

IPCC (2014). *Climate Change 2014: Mitigation of Climate Change. Contribution of Working Group III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA.

Iyigun, M., Nunn, N., and Qian, N. (2017). Winter is coming: The long-run effects of climate change on conflict, 1400-1900. *NBER Working Paper*, (January):51p.

Jack, B. K. (2011). Constraints on the Adoption of Agricultural Technologies in Developing Countries. (White paper prepared for the Agricultural Technology Adoption Initiative, JPAL (MIT) / CEGA (Berkeley)).

Jacoby, E. H. (1959). *Land Consolidation in Europe*. International Institute for Land Reclamation and Improvement, Wageningen, The Netherlands.

Jaffe, A. B., Newell, R. G., and Stavins, R. N. (2005). A tale of two market failures: Technology and environmental policy. *Ecological Economics*, 54(2-3):164–174.

Jenerette, G. D. and Shen, W. (2012). Experimental landscape ecology. *Landscape Ecology*, 27(9):1237–1248.

Jonsen, I. D. and Fahrig, L. (1997). Response of generalist and specialist insect herbivores to landscape spatial structure. *Landscape Ecology*, 12(3):185–197.

Kareiva, P. (1987). Habitat fragmentation and the stability of predator-prey interactions. *Nature*, 326(6111):388–390.

Kareiva, P. M. (1990). Population dynamics in spatially complex environments: theory and data. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 330(1257):175–190.

Karp, D. S., Chaplin-Kramer, R., Meehan, T. D., Martin, E. A., DeClerck, F., Grab, H., Gratton, C., Hunt, L., Larsen, A. E., Martínez-Salinas, A., O'Rourke, M. E., Rusch, A., Poveda, K., Jonsson, M., Rosenheim, J. A., Schellhorn, N. A., Tscharntke, T., Wratten, S. D., Zhang, W., Iverson, A. L., Adler, L. S., Albrecht, M., Alignier, A., Angelella, G. M., Zubair Anjum, M., Avelino, J., Batáry, P., Baveco, J. M., Bianchi, F. J. J. A., Birkhofer, K., Bohnenblust, E. W., Bommarco, R., Brewer, M. J., Caballero-López, B., Carrière, Y., Carvalheiro, L. G., Cayuela, L., Centrella, M., Ćetković, A., Henri, D. C., Chabert, A., Costamagna, A. C., De la Mora, A., de Kraker, J., Desneux, N., Diehl, E., Diekötter, T., Dormann, C. F., Eckberg, J. O., Entling, M. H., Fiedler, D., Franck, P., Frank van Veen, F. J., Frank, T., Gagic, V., Garratt, M. P. D., Getachew, A., Gonthier, D. J., Goodell, P. B., Graziosi, I., Groves, R. L., Gurr, G. M., Hajian-Forooshani, Z., Heimpel, G. E., Herrmann, J. D., Huseth, A. S., Inclán, D. J., Ingrao, A. J., Iv, P., Jacot, K., Johnson, G. A., Jones, L., Kaiser, M., Kaser, J. M., Keasar, T., Kim, T. N., Kishinevsky, M., Landis, D. A., Lavandero, B., Lavigne, C., Le Ralec, A., Lemessa, D., Letourneau, D. K., Liere, H., Lu, Y., Lubin, Y., Luttermoser, T., Maas, B., Mace, K., Madeira, F., Mader, V., Cortesero, A. M., Marini, L., Martinez, E., Martinson, H. M., Menozzi, P., Mitchell, M. G. E., Miyashita, T., Molina, G. A. R., Molina-Montenegro, M. A., O'Neal, M. E., Opatovsky, I., Ortiz-Martinez, S., Nash, M., Östman, Ö., Ouin, A., Pak, D., Paredes, D., Parsa, S., Parry, H., Perez-Alvarez, R., Perović, D. J., Peterson, J. A., Petit, S., Philpott, S. M., Plantegenest, M., Plećaš, M., Pluess, T., Pons, X., Potts, S. G., Pywell, R. F., Ragsdale, D. W., Rand, T. A., Raymond, L., Ricci, B., Sargent, C., Sarthou, J.-P., Saulais, J., Schäckermann, J., Schmidt, N. P., Schneider, G., Schüepp, C., Sivakoff, F. S., Smith, H. G., Stack Whitney, K., Stutz, S., Szendrei, Z., Takada, M. B., Taki, H., Tamburini, G., Thomson, L. J., Tricault, Y., Tsafack, N., Tschumi, M., Valantin-Morison, M., Van Trinh, M., van der Werf, W., Vierling, K. T., Werling, B. P., Wickens, J. B., Wickens, V. J., Woodcock, B. A., Wyckhuys, K. A. G., Xiao, H., Yasuda, M., Yoshioka, A., and Zou, Y. (2018). Crop pests and predators exhibit inconsistent responses to surrounding landscape composition. *Proceedings of the National Academy of Sciences of the United States of America*, 115(33):E7863–E7870.

Kelley, C. P., Mohtadi, S., Cane, M. A., Seager, R., and Kushnir, Y. (2015). Climate change in the Fertile Crescent and implications of the recent Syrian drought. *Proceedings of the National Academy of Sciences of the United States of America*, 112(11):3241–3246.

Kumar, V., Engel, R. P., Currie, R., Jha, P., Stahlman, P. W., and Thompson, C. (2019). Dicamba-resistant kochia (Bassia scoparia) in Kansas: Characterization and management with fall- or spring-applied PRE herbicides. *Weed Technology*, 33(2):342–348.

Lark, T. J., Mueller, R. M., Johnson, D. M., and Gibbs, H. K. (2017). Measuring land-use and land-cover change using the U.S. department of agriculture's cropland data layer: Cautions and recommendations. *International Journal of Applied Earth Observation and Geoinformation*, 62(June):224–235.

Larsen, A. E. (2013). Agricultural landscape simplification does not consistently drive insecticide use. *Proceedings of the National Academy of Sciences of the United States of America*, 2013:1–6.

Larsen, A. E., Gaines, S. D., and Deschênes, O. (2015). Spatiotemporal variation in the relationship between landscape simplification and insecticide use. *Ecological Applications*, 25(7):1976–1983.

Larsen, A. E. and Noack, F. (2017). Identifying the landscape drivers of agricultural insecticide use leveraging evidence from 100,000 fields. *Proceedings of the National Academy of Sciences of the United States of America*, 114(21):5473–5478.

Libecap, G. D. and Lueck, D. (2011). The Demarcation of Land and the Role of Coordinating Property Institutions. *Journal of Political Economy*, 119(3):426–467.

Lobell, D. B., Bänziger, M., Magorokosho, C., and Vivek, B. (2011). Nonlinear heat effects on African maize as evidenced by historical yield trials. *Nature Climate Change*, 1(1):42–45.

Luterbacher, J. and Pfister, C. (2015). The year without a summer. *Nature Geoscience*, 8(4):246–248.

MAE (2012). *Manual Operativo: Unificado Proyecto Socio Bosque.*

Manski, C. F. (1993). Identification of Endogenous Social Effects: The Reflection Problem. *The Review of Economic Studies*, 60(3):531–542.

Manuelli, R. E. and Seshadri, A. (2014). Frictionless Technology Diffusion: The Case of Tractors. *American Economic Review*, 104(4):1368–1391.

Martin, A. (2013). Le Partenariat Européen pour l'Innovation (PEI) : la mise en réseaux comme levier de l'innovation en agriculture.

Mather, A. S. (1992). The Forest Transition. *Area*, 24(4):367–379.

McCown, S., Barber, T., and Norsworthy, J. K. (2018). Response of Non-Dicamba-Resistant Soybean to Dicamba As Influenced by Growth Stage and Herbicide Rate. *Weed Technology*, 32(5):513–519.

McGarigal, K. and Cushman, S. A. (2002). Comparative evaluation of experimental approaches to the study of habitat fragmentation effects. *Ecological Applications*, 12(2):335–345.

McGarigal, K., Cushman, S. A., Neel, M. C., and Ene, E. (2002). FRAGSTATS: Spatial Pattern Analysis Program for Categorical Maps.

McGuirk, E. and Burke, M. (2017). The Economic Origins of Conflict in Africa. *NBER Working Paper*, (w23056).

Meehan, T. D. and Gratton, C. (2015). A consistent positive association between landscape simplification and insecticide use across the Midwestern US from 1997 through 2012. *Environmental Research Letters*, 10(11):114001.

Meehan, T. D. and Gratton, C. (2016). A Landscape View of Agricultural Insecticide Use across the Conterminous US from 1997 through 2012. *PLoS ONE*, 11(11):e0166724.

Meehan, T. D., Werling, B. P., Landis, D. A., and Gratton, C. (2011). Agricultural landscape simplification and insecticide use in the Midwestern United States. *Proceedings of the National Academy of Sciences of the United States of America*, 108(28):11500–11505.

Mesinger, F., DiMego, G., Kalnay, E., Mitchell, K., Shafran, P. C., Ebisuzaki, W., Jovic, D., Woollen, J., Rogers, E., Berbery, E. H., Ek, M. B., Fan, Y., Grumbine, R., Higgins, W., Li, H., Lin, Y., Manikin, G., Parrish, D., and Shi, W. (2006). North American Regional Reanalysis. *Bulletin of the American Meteorological Society*, 87:343–360.

Meyfroidt, P. and Lambin, E. F. (2009). Forest transition in Vietnam and displacement of deforestation abroad. *Proceedings of the National Academy of Sciences of the United States of America*, 106(38):16139–16144.

Missirian, A. and Schlenker, W. (2017). Asylum Applications and Migration Flows. *American Economic Review: Papers & Proceedings*, 107(5):436–440.

Munro, A. (2008). The spatial impact of genetically modified crops. *Ecological Economics*, 67(4):658–666.

Obama, B. (2015). Remarks by the president at the United States Coast Guard Academy Commencement. (United States Coast Guard Academy. New London, Connecticut).

Olimpi, E. M., Garcia, K., Gonthier, D. J., De Master, K. T., Echeverri, A., Kremen, C., Sciligo, A. R., Snyder, W. E., Wilson-Rankin, E. E., and Karp, D. S. (2020). Shifts in species interactions and farming contexts mediate net effects of birds in agroecosystems. *Ecological Applications*.

Oppenheimer, M., Campos, M., Warren, R., Birkmann, J., Luber, G., O'Neill, B. C., and Takahashi, K. (2014). Emergent Risks and Key Vulnerabilities. In Field, C. B., Barros, V., Dokken, D., Mach, K., Mastrandrea, M. D., Bilir, T., Chatterjee, M., Ebi, K., Estrada, Y., Genova, R., Girma, B., Kissel, E., Levy, A., MacCracken, S., Mastrandrea, P., and White, L., editors, *Climate Change 2014: Impacts, Adaptation, and Vulnerability. Part A: Global and Sectoral Aspects. Contribution of Working Group II to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, chapter 19, pages 1039–1099. Cambridge University Press, Cambridge, UK.

Ostrom, E. (2000). Collective action and the evolution of social norms. *Journal of Economic Perspectives*, 14(3):137–158.

Perz, S. G. and Skole, D. L. (2003). Secondary forest expansion in the Brazilian Amazon and the refinement of forest transition theory. *Society and Natural Resources*, 16(4):277–294.

Piguet, E. and Laczko, F. (2014). *People on the Move in a Changing Climate. The Regional Impact of Environmental Change on Migration*. Springer, Dordrecht; New York.

Pimm, S. L. and Raven, P. H. (2000). Extinction by numbers. *Nature*, 403:843–845.

Powles, S. B. (2008). Evolved glyphosate-resistant weeds around the world: lessons to be learnt. *Pest Management Science*, 64(4):360–365.

PRISM Climate Group at Oregon State University (2019). Gridded Climate Data.

Reifman, A. S., Larrick, R. P., and Fein, S. (1991). Temper and Temperature on the Diamond: The Heat-Aggression Relationship in Major League Baseball. *Personality and Social Psychology Bulletin*, 17(5):580–585.

Roberts, M. J. and Schlenker, W. (2013). Identifying Supply and Demand Elasticities of Agricultural Commodities: Implications for the US Ethanol Mandate. *American Economic Review*, 103(6):2265–2295.

Robinson, R. A. and Sutherland, W. J. (2002). Post-war changes in arable farming and biodiversity in Great Britain. *Journal of Applied Ecology*, 39(1):157–176.

Rosenzweig, M. L., Brown, J. S., and Vincent, T. L. (1987). Red Queens and ESS: the coevolution of evolutionary rates. *Evolutionary Ecology*, 1(1):59–94.

Rusch, A., Chaplin-Kramer, R., Gardiner, M. M., Hawro, V., Holland, J., Landis, D. A., Thies, C., Tscharntke, T., Weisser, W. W., Winqvist, C., Woltz, M., and Bommarco, R. (2016). Agricultural landscape simplification reduces natural pest control: A quantitative synthesis. *Agriculture, Ecosystems and Environment*, 221:198–204.

Sacks, W. J., Deryng, D., Foley, J. A., and Ramankutty, N. (2010). Crop planting dates: An analysis of global patterns. *Global Ecology and Biogeography*, 19(5):607–620.

Schlenker, W. and Lobell, D. B. (2010). Robust negative impacts of climate change on African agriculture. *Environmental Research Letters*, 5:014010.

Schlenker, W. and Roberts, M. J. (2009). Nonlinear temperature effects indicate severe damages to U.S. crop yields under climate change. *Proceedings of the National Academy of Sciences of the United States of America*, 106(37):15594–15598.

Schlenker, W. and Walker, W. R. (2016). Airports, air pollution, and contemporaneous health. *Review of Economic Studies*, 83(2):768–809.

Sciumbato, A. S., Chandler, J. M., Senseman, S. A., Bovey, R. W., and Smith, K. L. (2004). Determining Exposure to Auxin-Like Herbicides. II. Practical Application to Quantify Volatility. *Weed Technology*, 18(4):1135–1142.

Scott, J. C. (1998). *Seeing Like a State. How Certain Schemes to Improve the Human Condition Have Failed*. Yale University Press.

Scott, P. T. (2013). Dynamic Discrete Choice Estimation of Agricultural Land Use (JMP). pages 1–55.

Simpson, E. H. (1949). Measurement of Diversity. *Nature*, 163(1946):688.

Sinha, P. and Cropper, M. L. (2013). The Value of Climate Amenities: Evidence from US Migration Decisions. *NBER Working Paper*, (18756).

Stern, N. (2007). Stern Review: The Economics of Climate Change. Technical report.

Sunding, D. and Zilberman, D. (2001). The Agricultural innovation process: Research and technology adoption in a changing agricultural sector. In Gardner, B. L. and Rausser, G. C., editors, *Handbook of Agricultural Economics, Volume 1A*, chapter 4, pages 207–261. Elsevier B.V.

Tack, J., Barkley, A., and Nalley, L. L. (2015). Effect of warming temperatures on US wheat yields. *Proceedings of the National Academy of Sciences of the United States of America*, 112(22):6931–6936.

Thelin, G. P. and Stone, W. W. (2013). Estimation of annual agricultural pesticide use for counties of the conterminous United States, 1992–2009. *U.S. Geological Survey Scientific Investigations Report 2013-5009*, pages 1–54.

Thévenot, L. (1996). Justification et compromis. In Canto-Sperber, M., editor, *Dictionnaire d'éthique et de philosophie morale*, pages 789–794. PUF, Paris.

Tscharntke, T., Klein, A.-M., Kruess, A., Steffan-Dewenter, I., and Thies, C. (2005). Landscape perspectives on agricultural intensification and biodiversity - ecosystem service management. *Ecology Letters*, 8(8):857–874.

Turner, M. G. (2005). Landscape Ecology: What Is the State of the Science? *Annual Review of Ecology, Evolution, and Systematics*, 36:319–344.

USDA National Agricultural Statistics Service (2018). Cropland Data Layer 2008-2018.

Van Valen, L. (1973). A New Evolutionary Law. *Evolutionary Theory*, 1:1–30.

Wechsler, S. J., McFadden, J. R., and Smith, D. J. (2018). What do farmers' weed control decisions imply about glyphosate resistance? Evidence from surveys of US corn fields. *Pest Management Science*, 74:1143–1154.

Weiss, D. J., Nelson, A., Gibson, H. S., Temperley, W., Peedell, S., Lieber, A., Hancher, M., Poyart, E., Belchior, S., Fullman, N., Mappin, B., Dalrymple, U., Rozier, J., Lucas, T. C., Howes, R. E., Tusting, L. S., Kang, S. Y., Cameron, E., Bisanzio, D., Battle, K. E., Bhatt, S., and Gething, P. W. (2018). A global map of travel time to cities to assess inequalities in accessibility in 2015. *Nature*, 553(7688):333–336.

White, C. A. (1983). *A History of the Rectangular Survey System*. Washington, D.C.

White, M. J. (2004). The "Arms Race" on American Roads: The Effect of Sport Utility Vehicles and Pickup Trucks on Traffic Safety. *The Journal of Law and Economics*, 47(2):333–355.

Wiens, J. A. (1976). Population Responses To Patchy Environments. *Annual review of Ecology and Systematics*, 7:81–120.

Wiens, J. A. (2000). Ecological heterogeneity: an ontogeny of concepts and approaches. In Hutchin, M. J., John, E. A., and Stewart, A. J. A., editors, *The Ecological Consequences of Heterogeneity*, pages 9–31. Blackwell Science, Oxford, UK.

Wunder, S. (2005). Payments for environmental services: some nuts and bolts.

Yackulic, C. B., Fagan, M., Jain, M., Jina, A., Lim, Y., Marlier, M. E., Muscarella, R., Adame, P., DeFries, R. S., and Uriarte, M. (2011). Biophysical and socioeconomic factors associated with forest transitions at multiple spatial and temporal scales. *Ecology and Society*, 16(3):23.

Yang, D. (2006). Why do migrants return to poor countries? Evidence from Philippine migrants' responses to exchange rate shocks. *The Review of Economics and Statistics*, 88(4):715–735.

Yang, M., Blyth, W., Bradley, R., Bunn, D., Clarke, C., and Wilson, T. (2008). Evaluating the Power Investment Options with Uncertainty in Climate Policy. *Energy Economics*, 30(4):1933–1950.