Title
Essays on the Economics of Behavioral and Environmental Changes

Author
Annan-Phan, Sébastien

Publication Date
2021

Peer reviewed|Thesis/dissertation

Essays on the Economics of Behavioral and Environmental Changes

by

Sébastien Annan-Phan

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Agricultural and Resource Economics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Solomon Hsiang, Co-chair
Professor Maximilian Auffhammer, Co-chair
Professor Meredith Fowlie
Professor Joshua Blumenstock

Summer 2021

Essays on the Economics of Behavioral and Environmental Changes

Abstract

Essays on the Economics of Behavioral and Environmental Changes

by

Sébastien Annan-Phan

Doctor of Philosophy in Agricultural and Resource Economics

University of California, Berkeley

Professor Solomon Hsiang, Co-chair

Professor Maximilian Auffhammer, Co-chair

Climate change is set to have a significant economic impact, affecting our behavior and changing the ways we produce and consume. This dissertation illustrates different channels through which environmental change has already been impacting our economy and highlights potential measures and market designs to mitigate future costs. The first two chapters of this dissertation leverage historical weather data to assess the market and behavioral responses to extreme temperature shocks. Chapter 1 utilizes European electricity markets data to assess the non-linear relationship between weather and energy cost. The analysis of spot market electricity prices reveals that estimates of the energy cost due to climate change which rely solely on consumption are significantly biased downward. In addition, the creation of a "unified" European market through regional market integration has contributed to drastically lower the impact of extreme temperatures on electricity price. Chapter 2 demonstrates that warm temperatures are associated with a rise in violent crime in the US. I further investigated the effect of extreme weather on police officer behavior using original and novel crowd-sourced data on civilian deaths involving police. The results indicate that fatal shootings increase proportionately to the growing number of violent crimes during warmer days. Taser use and physical restraints — two controversial uses of force — increase significantly on such days regardless of the threat level, indicating a need to reevaluate their usage. Finally, chapter 3 explores the distribution of human attention to moments in time using Google search data. Attention is a key parameter in the field of environmental economics as public policies enacted today are a direct function of how much people value the past, current, and future states of the economy. The results suggest strong common patterns of thought with respect to time, however recent trends and regional variations indicate that some non-biological factors can alter these patterns slightly.

Contents

Acknowledgments

I am sincerely grateful for the many peers, mentors, friends, and family who made this doctoral thesis possible and gave me the strength to pursue a PhD so far from home.

First and foremost, I am deeply thankful to my mentors Solomon Hsiang and Maximilian Auffhammer. Sol's support originated during my time at the University of Chicago, before I even started on my PhD, and I can only hope it will continue for many years. His advice and insights on research and life are without a doubt the most precious thing I got from this journey. Max convinced me to join the program at Berkeley and kept supporting me from the first semester until the last; he would be the perfect advisor were it not for his dubious love for the wrong soccer team.

I thank the entire environmental and resource economics community at UC Berkeley, the Global Policy Lab, and the 2017 ARE cohort for the critical feedbacks on my research and continuous emotional support during the four PhD years. I thank in particular Meredith Fowlie, who would always provide me with insightful comments during our electricity market discussions, and who would always go the extra mile to help me grow as a researcher. I thank Bethy Sadoulet for lending a shoulder and an ear at any time. Livia and Alejandro have been the pillars of my time as a PhD student, and I thank them for their friendship and support.

I am also indebted to my co-authors and good friends Bocar Ba and Léopold Biardeau. Bocar, you played a major role in my development as a researcher, from our first discussions at the Rubenstein Library back in 2014 to our casual phone calls. I am proud to be your friend and colleague.

Thank you to my mother, Florence, for her unconditional support, love, and kindness. I thank my father, James, and my siblings for being there for me, no matter what. I have the great fortune of being part of another family: thank you David, Etienne, Florimond, Gaëtan, Laurent, Stefan, and Thibaut, for all the "fishing", the chatters, the laughs, and the memories I share with you every day.

Last, I am forever grateful to my wife and partner in life Najwa. Like many — many — things in my life, this dissertation would have simply not come into existence without you. The continuous support and happiness you give me made this thesis possible, but it is our intellectual connection and never-ending discussions that made it a reality.

To my daughter and my sunshine, Selma.

Overview

"*Ce sont les différents besoins dans les différents climats qui ont formé les différentes manières de vivre; et ces différentes manières de vivre ont formé les diverses sortes de lois*"

"Different needs in different climates formed different ways of living, and from these different ways of living originate different kind of laws."

— De l'esprit des lois (Baron de Montesquieu, 1872)

Climate is changing and so should our laws and institutions, but also the way our markets operate. As soon as we introduce an externality[1] — such as climate change — Adam Smith's famous "invisible hand" begins to fail, and it is no longer socially optimal to leave markets on their own. The increase in frequency and magnitude of extreme weather events is known to have severe impact on many aspects of the economy, including health (Carleton et al., 2018), labor productivity (Graff Zivin and Neidell, 2014), and agricultural production (Schlenker and Roberts, 2009). These costs, as well as potential benefits, need to be precisely estimated in order to be internalized by the markets through public policy (e.g. carbon tax). Being able to quantify these impacts is a key to building today's policies and to promoting an efficient path to adaptation, which requires both planning and timely investments. Decades of variation in weather patterns at the regional level provide valuable information for establishing our systems' best responses to environmental change. We cannot perfectly predict the future state of the economy nor the kind of technology available, but we can leverage data from the past to estimate the effect that extreme weather has on the economy and thus identify efficient ways to mitigate future impacts associated with climate change. This dissertation mobilizes and combines datasets from various fields including economics and social sciences (e.g. wholesale price, population density) as well as physical science (e.g. weather forecasts). It also utilizes publicly available data from private entities and crowd-sourced information. Taken together, the essays in this dissertation illustrate different ways of combining these original datasets with insights and methodologies from the economic science to advance our knowledge of the economics of climate change.

[1] A perturbation not taken into account by the market.

Chapter 1 explores the impact of extreme weather events on wholesale electricity markets. Studies on the impact of climate change on the energy sector have been focusing on the quantity effect, namely the fact that extreme weather is associated with higher energy consumption (Deschenes and Greenstone, 2011b; Auffhammer et al., 2017). Chapter 1 goes beyond the *quantity effect* and analyzes the price response to changes in temperature. My findings indicate that failure to take into account the *price effect* has led previous studies to significantly underestimate the impact of extreme weather on energy cost. The results presented in this first chapter indicates that extreme weather not only increases energy consumption but is also associated with significantly higher wholesale prices. Using a "natural experiment" in Central Europe, I also estimated the potential of using market integration as a means of mitigating the effect of extreme weather events. It appears that market integration of the national wholesale markets in Belgium, France, Germany, and the Netherlands drastically mitigated the effect of weather shocks, reducing the marginal effect of a day at 25^oC by 14 to 28 percentage points (relative to a day at 15^oC). Enlarging regional markets through regulation or private initiative is both efficient and considerably less expensive than a physical extension of the grid.

Chapter 2, co-authored with Bocar Ba, studies the non-linear impact of outdoor temperature on the behavior of police officers. Using FBI data combined with crowd sourced data on officers involved in civilian deaths in the US from 2000 to 2016, we demonstrate that both violent crimes and the number of officers assaulted or killed increase on warmer days (days with average daily temperature above 17^oC). We take this result as an indicator of greater personal danger on such days. Consistent with higher threat levels, we found suggestive evidence that fatal shootings also increase during warmer days. However, when accounting for a surge in officer-civilian interaction, we found no additional effect of high temperatures on fatal shootings, indicating a lack of behavioral or physiological response from the officers. When investigating other causes of death, we found that during "extremely warm" days (average daily temperature above 20^oC), the number of casualties associated with Taser use and physical restraint was significantly higher independent of increased interaction between officers and civilians. The results suggest a need to reevaluate the use of Tasers and physical restraint techniques — two controversial uses of force — to prevent unintended deaths.

Chapter 3, co-authored with Léopold Biardeau and Solomon Hsiang, investigates the distribution of human attention toward different moments in time. It is unknown how much individuals naturally think about the past, present, and future. This chapter provides the first evidence of a coherent probability distribution governing attention across moments in time by deconvolving how billions of individuals query the *Google Search* engine. We have discovered consistent and generalizable structure to the distribution of attention across time; regardless of individuals' language or country. We estimate that the present captures roughly 25% of time-related attention, while all the past and future moments combined capture roughly 39% and 36% of attention, respectively. Almost no attention is detected more than 200 days beyond the present moment. Despite consistency in the shape of the distribution

of attention around the world, we were able to identify regional patterns. In addition, it appears that during the period 2008–2018. the share of attention dedicated to "the present" increased gradually at the expense of attention to the past. Together, these results suggest strong common patterns of thought with respect to time, however recent trends and regional variations indicate that some non-biological factors can alter these patterns slightly.

Chapter 1

Adaptation through Market Integration: Mitigating the Impact of Climate Shocks on Electricity Expenditure

1.1 Introduction

The literature documenting the impact of climate change on the economy has grown considerably in the past decades. The effect of extreme temperatures on various sectors has been well established, in particular on health (Deschenes and Greenstone, 2011b), agricultural yield (Schlenker and Roberts, 2009), labor productivity (Graff Zivin and Neidell, 2014), crime level and conflict (Ranson, 2014b; Hsiang et al., 2013), but also energy consumption (Deschenes and Greenstone, 2011b; Wenz et al., 2017). The electricity sector is a particular one. On the one hand, it is the main source of green house gas emissions[1], meaning it generates a strong feed back effect. On the other hand, it is associated with considerable benefits and is tightly connected to other sectors because of the defensive property of air conditioning (AC). For instance, with regard to health effects, Barreca et al. (2016) demonstrated that people can adapt to climate change and significantly mitigate the effect of hot days on mortality by using AC (see Carleton et al. (2018) for a global assessment of adaptation). AC also lowers the effect of weather variation on labor productivity as demonstrated by Graff Zivin and Neidell (2014) in relation to indoor vs outdoor substitution. Thus, as electricity becomes more needed in a warmer world it is crucial to analyze precisely the associated changes in cost, as well as possible mitigating mechanisms. In this chapter, I analyze how weather shocks not only impact electricity consumption but also increase the marginal cost of provision. Using detailed European market data, I estimated the causal effect of market integration on weather shocks mitigation using an "interacted" *difference-in-differences* approach. I found

[1]Based on US numbers for the period 1990-2017 (EPA, 2019)

that relative to a "normal day" at 15^{o}C, market integration mitigated the effect of a 25^{o}C day on the price of electricity by roughly 20 percentage points. The results also suggest that limiting the scope of the analysis to consumption would underestimate the expected damages of climate change on energy expenditure.

Going beyond the sole effect of rising temperatures on aggregate consumption is necessary to document the increase in energy provision cost. The cost of provision matters for calibrating future climate policies, for instance distinguishing the effect of temperature within a day provides insights on the type of investment needed. If temperature only affects consumption during a limited period of the day, one might prioritize the development of energy storage capacity over the construction of new units to take advantage of the existing float of power plants. Limiting the focus on quantity would thus mask the heterogeneous effect of temperature on marginal cost at different times of the day. Auffhammer et al. (2017) found that peak hours are more responsive to extreme temperatures, meaning that one might underestimate the energy cost of climate shock by ignoring the shift in marginal cost. Similarly, one might want to go beyond the quantity effect when estimating the effect of weather shocks at the international level. Different countries have heterogeneous responses to weather shocks. For example, French consumers are mostly using electric heating while Germany relies more on gas, the same demand shock would thus lead to very different increase in marginal cost in the two countries if they were in autarky. Estimating international heterogeneity in response to temperature changes would enlighten the discussion on market integration: the more connected the grids, the more efficient the generation allocation in reaction to temperature shocks.

The effect of extreme temperatures on the energy sector has been widely studied, but existing studies tend to focus on the effect of temperature on electricity consumption. Deschenes and Greenstone (2011b) were the first to document using long time-series how energy consumption changes with temperature, accounting for temporal and geographical unobservables. Using semi parametric regression, they estimated the non-linear effect of temperature on residential energy consumption, the so-called *dose-response* function. They found that for each additional day above 90^{o}F, the annual energy consumption increases by roughly 0.4% compared to a day in the range $50^o - 60^o$. The dose-response function is usually U-shaped and reflects that both cold and hot days are associated with an increase in electricity consumption[2]. The U-shape form is explained by the use of electric heating and AC on cold and hot days, respectively. However, these studies only looked at the demand part of the problem neglecting the supply side of the market. The estimates are all expressed in MW/h (i.e. in *quantity*) whereas one might also be interested in prices and the overall economic cost. There is no reason to assume that (1) the marginal cost is constant (i.e. that a 1% increase in consumption is equivalent to a 1% increase in expenditure), and (2) the cost function itself is not affected by change in temperature. Fewer papers have investigated the effect of extreme temperature on power generation cost. Thermal efficiency

[2]Deschenes and Greenstone found a consistent U-shaped relationship — in particular a day below 10^{o}F would increase annual consumption by 0.32%

is a direct function of ambient temperature, so one would expect the efficiency of power plants to decrease when temperatures rise. In addition, cooling systems depend heavily on nearby streams, which are subject to regulation limiting the maximum temperature of the water. Linnerud et al. (2011) estimated using time series analysis that nuclear generation decreases by 0.5–2% for each additional degree rise in temperature; however, this study and other studies in the engineering literature are limited by a small number of locations due to data access restrictions for reasons of national security. McDermott and Nilsen (2014) investigated water scarcity to estimate a more global impact on the electricity market. Their results suggest that an additional degree in stream temperature above 25°C would increase wholesale prices by roughly 2.5%, but they did not separately identify the effect of temperature on energy demand and supply. This distinction is at the core of this chapter since it is necessary to analyze both sides of the market to estimate changes in cost. In their survey on the measurement of climatic impact on energy demand, Auffhammer and Mansur (2014) focused on the distinction between the intensive and the extensive margin (AC adoption). Interestingly, *energy consumption* and *energy expenditure* are used interchangeably in the article although little is actually said about the cost and the change in prices. Mideksa and Kallbekken (2010) specifically reviewed the impact of climate change on *electricity markets* and, highlighted that existing studies seem to adress only the demand or the supply effect. With regard to market integration, the authors highlighted that price differences will emerge between Northern and Southern Europe, increasing the need for interconnection capacities, which has motivated this study.

This chapter contributes to the literature by focusing on the effect of temperature on electricity prices. Going beyond the average effect of temperature, I propose two distinct methods to estimate the causal effect of market integration on the *price temperature* relationship. First, I used a simple reduced form estimation of the effect of temperature on electricity prices. This method still lacks the tractability to identify whether price changes are due to movements in supply or demand. Unless I assume that demand is purely inelastic, moreover it does not allow for an accurate estimate of change in consumer expenditure. However, the simplicity of the reduced form model and the granularity of the price data add some flexibility to the analysis. In particular, I used the reduced form to estimate the causal effect of a market integration policy on the temperature-price relationship. I applied an augmented *difference-in-differences* estimator to identify the causal effect of such a policy on the magnitude of the marginal impact of temperature. To assess the benefit of releasing the physical congestion constraint (intensive margin), I estimated the reduced form model on counterfactual equilibrium prices accounting for grid expansion.

The remainder of this chapter is structured as follows: section 1.2 provides some background information about electricity markets and presents a theoretical model to evaluate trade in the context of weather shocks. Section 1.3 describes the data and section 1.4 focuses on the identification strategy. Section 1.5 and 1.6 present and discuss the results, respectively.

1.2 Background

Electricity markets in Europe operate according to the zonal pricing rule. Italy, Sweden, and Denmark use zonal prices but have multiple bidding zones[3]. *Zonal pricing* means that price will be the same at every nodes of the grid within the bidding zone, which for most European countries is the entire national territory (eg. France, Belgium, Poland). Some bidding zones are sometimes cross-countries, that is the case for both the German-Austrian-Luxembourg market and the Iberian market[4]. Zonal pricing is to be differentiated from *nodal pricing* where each node can have a different price based on grid congestion. In the zonal pricing, the grid operator internalizes the grid congestion and unsure a uniform price over the bidding zone. Congestion costs are centralized by the grid operator and are included in an annual fee only varying in the mid-long run. The use of zonal pricing in Europe instead of nodal pricing matters in term of climate change impacts. With a zonal pricing system, an exogenous shock on demand (*eg.* heat wave) will be smoothed out at the country level whereas the same shock in a nodal system would significantly increase the price in the sub-region leaving the neighboring regions roughly unaffected. Nodal pricing is beneficial in the long run as it provides the right incentive to invest in new transmission lines. On the opposite, one could argue that such infrastructure might be oversized if those shocks are rare events. If that were the case, it might be more efficient to smooth at the national level those events to keep the burden on the consumers at a minimum level. One other advantage of zonal pricing is that it limits market power. By definition nodal pricing is splitting the market into much smaller sub markets, reducing the cost of transaction associated with collusion and increasing the market power of large utilities (see Wolak (2011) for more details).

Electricity supply needs to equal the demand at all time and at each node for the grid to remain safe, which necessitates distinct electricity markets for different time periods. From long-term to short term, actors are involved in OTC transactions, Future market, Day-ahead and Intraday market (spot markets), and finally the balancing market sometimes called "real-time". Day-ahead (DA) markets are by far the most used and I focus on this one for this study. DA are operated every day at 11:00 am and define hourly price and quantity for the next 24 hours. Actors must submit their book order before the market close (11:00 am) to the unique market operator that is in charge of clearing the market using auction rules. The detail of the auction rule is described further in the book order section. Since weather predictions for the next day are reliable and easily accessible, day-ahead prices are expected to react to daily change in temperature. Trade between zonal markets exist but are limited by (1) the physical capacity of the interconnection lines and (2) the auction system itself. If market operators are not coordinated the market is said to be *not coupled*, meaning that each actor has to bid on a separate market to book a specific amount of the interconnection capacity. The amount of available transfer capacity (ATC) depends on long term contracts and safety measures related to grid overload. When markets coordinate their

[3]Four in Sweden, 2 in Denmark and 5 in Italy

[4]Composed of Spain and Portugal

operations and auction rules, they are "coupled". I will use market coupling as a proxy for market integration in the difference in differences estimation. I now derive a simple model to describe how trade matters when estimating the cost of climate change on electricity provision.

Modeling the cost of climate change on an interconnected grid

Electricity is a textbook example of microeconomics theory applied to trade, it is perfectly homogeneous and has a marginal transport cost close to zero. For the purpose of this study I focus on the electricity sector of two neighboring countries, which are both in perfect competition and are described by representative firms with linear marginal cost such that $MC_1(q) = \alpha MC_2(q)$ and $MC_i(0) = 0$. I assume for now that firms in both countries have the same marginal cost ($\alpha = 1$) and demand is fully inelastic[5] in both countries ($D_1 = D_2 = E$). Markets are clearing at the marginal cost so prices are equals in both market to P^*:

$$P^* = MC_i(E) \tag{1.1}$$

Since equilibrium prices are equals, there is no need for trade and the value of the interconnection line is zero. I can introduce variation in temperature as a demand shifter, whenever there is either a need for electric heating (ambient temperature below 5^oC) or air conditioning (temperature above 20^oC) the demand is increased by a fixed parameter ω:

$$D_i = E + \omega \, \mathbb{1}\{T_i \leqslant 5\} + \omega \, \mathbb{1}\{T_i \geqslant 20\} \tag{1.2}$$

For a sufficiently big shock in temperature δ in country 1, if markets are in Autarky the prices at the equilibrium will be different:

$$\begin{cases} P_1^A = MC(E + \omega) \\ P_2^A = MC(E) \end{cases} \quad \Rightarrow P_1^A > P_2^A \tag{1.3}$$

Figure 1.1 panel A illustrates the additional cost due to δ in the shocked country in the Autarky case. It is a direct function of MC and ω, and is given by:

$$\int_E^{E+\omega} MC_1(q)dq = \underbrace{\omega \, MC_1(E)}_{\text{A}} + \underbrace{\frac{\omega \, MC_1(\omega)}{2}}_{\text{B}} \tag{1.4}$$

The additional cost can be divided in two areas A and B. The former denotes the increase in cost due to a change in the quantity supplied, whereas the later is coming from the necessity to fire up new power plants that are more costly to operate (if marginal cost were constant B would be null). By definition, area A does not depend on the marginal cost of new units,

[5]Electricity demand is often represented as purely inelastic at least in the short run, here it does not change the results of the model and greatly simplified the graphical representation.

therefore it cannot benefit from trade since at the equilibrium we already have $MC_1 = MC_2$. One then needs to focus on area B to describe how trade reduces the cost induced by δ. In this framework, country 1 has to feed the grid with an extra quantity of electricity ω and could either face it alone and bear the whole cost, or import electricity up to the point where the marginal costs are equals. The additional cost can then be rewritten as:

$$\int_{E}^{E+\omega} MC_1(q)dq = A + B_a + B_t \tag{1.5}$$

Where B is now split between the inevitable cost that would still occurs in free trade (B_a) and the dead weight loss due to absence of trade (B_t). The dead weight loss B_t can also be seen as the shadow value of the interconnection line. Figure 1.1 panel B illustrates the described model with trade when the marginal cost in country 2 is lower[6] than in country 1 (*ie.* $\alpha > 1$), however as we can intuitively guess graphically, we would still have gain from trade even if the two marginal costs where equals ($\alpha = 1$). More formally, in the linear case and if we add a management constraint ($\theta \leqslant 1$[7]) on the interconnection capacity we obtain:

$$B_t = \frac{MC_1(\omega) + MC_1(1-\gamma) - MC2(\gamma)}{2}\gamma \tag{1.6}$$

$$with\, \gamma = \theta\omega\frac{\alpha}{\alpha+1} \tag{1.7}$$

Rearranging the right hand side gives

$$B_t = \omega\frac{MC_1(\omega)}{2}[\theta\alpha(1 - 2\frac{\theta\alpha}{(\alpha+1)^2})] \tag{1.8}$$

and partial derivatives

$$\frac{\partial B_t}{\partial \omega} > 0\ , \frac{\partial B_t}{\partial \theta} > 0\ , \frac{\partial B_t}{\partial \alpha} > 0\ , \frac{\partial B_t}{\partial MC_1(\omega)} > 0$$

Which gives us four main intuitive results: the value of the interconnection increases with (1) the quadratic of the demand shift[8] and (2) a better interconnection management. (3) The higher the marginal cost in the shocked country, the more it benefits from trade, and (4) this effect is amplified by the relative difference in marginal cost between countries (parameterized with α).

[6]Precisely $MC_2(q) = \alpha MC_1(q)$ only for $q > E$. Before the temperature shock δ the two markets had similar prices and marginal costs ($MC_1(E) = MC2(E)$)

[7]For example if $\theta = 0.8$ we can only effectively import 80% of the optimal traded quantity at the equilibrium. In the linear case, we know that if $MC_1 = \alpha MC_2$ then country 2 should provide $\frac{\alpha}{\alpha+1}$ of the total electricity ω needed (*eg.* if the marginal cost in country 1 is twice the marginal cost in country 2, then country 1 imports 2/3 of the extra electricity needed)

[8]In the linear case $MC_1(\omega) = a\omega + b$, so B_t is a funtion of ω^2

Figure 1.1: Additional cost in (A) Autarky and (B) free trade

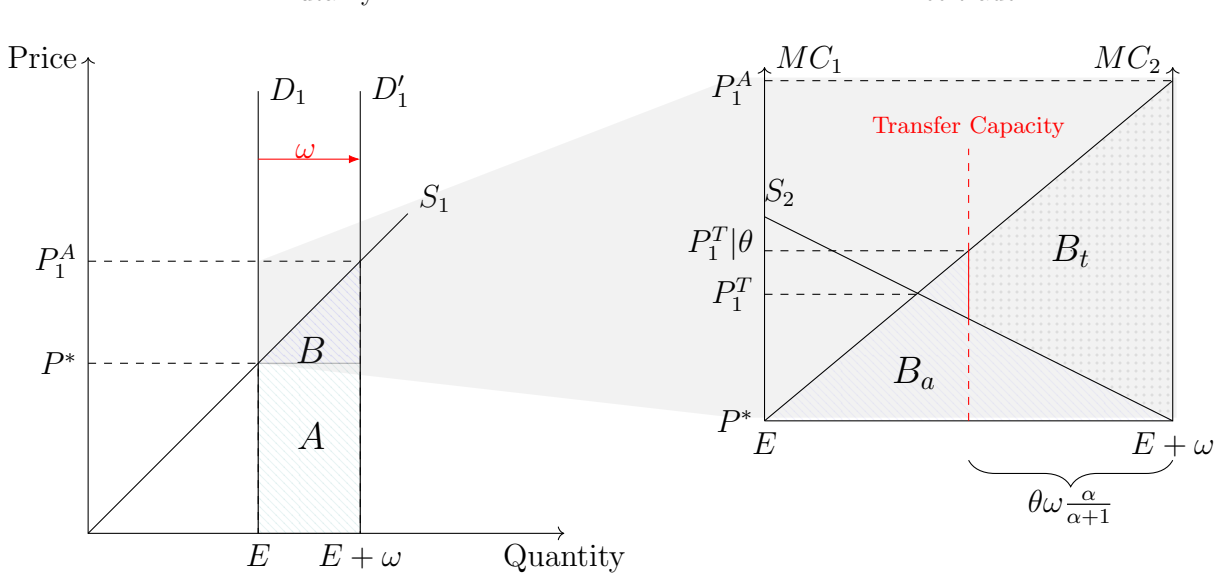

As a result of trade, cost is minimized and price drops from P_1^A to $P_1^T|\theta$. The simple model above highlighted that trade has a separate and identified value when a country faces an exogenous shock in temperature that shifts the demand. This result is true even when marginal cost are equals, only the magnitude of the effect is lowered[9]. One could argue that MC_1, MC_2, and the interconnection capacity itself might be varying with temperature shock δ, however regardless of the aggregated effect, an increase in the interconnection management θ is expected to lower the price of the shocked country. the identification strategy and the proxy variable for the parameter θ will be detailed in section 1.4.

Ultimately, if one cares about climate change impact on welfare, it is necessary to model both the demand and the supply as a function of δ. Marginal cost of country i is now:

$$MC_i = f(q) + \phi \mathbb{1}\{T_i \geqslant \tau\} \tag{1.9}$$

where ϕ is a fixed damage associated with loss in thermal efficiency when temperature T goes above a given threshold τ. The extra cost associated with a large enough temperature shock δ is now affecting every unit of electricity produced. The gain from trade are increasing since MC_1 is now larger, but there is also a much bigger share of the additional cost that cannot be reduced as shown by A' in figure 1.2. However, it does not change the direction of a marginal change in θ on price. For a given effect of temperature shock on price ($\frac{\partial P}{\partial \delta}$), price change is a sufficient statistic to estimate the mitigation impact of trade ($\Delta\theta$). Future research projects will focus on analyzing the nonlinear effect of temperature on supply and demand, which means estimating ω and ϕ. With those parameters, one could perform cost

[9]In fact the result holds even if $MC_1 < MC_2$, as long as α is such that $MC_2 < MC_1(E+\omega)$

and welfare analysis for counterfactual weather on the training sample to better understand the impact of climate change.

Figure 1.2: The total impact of temperature on cost when both sides of the market are impacted

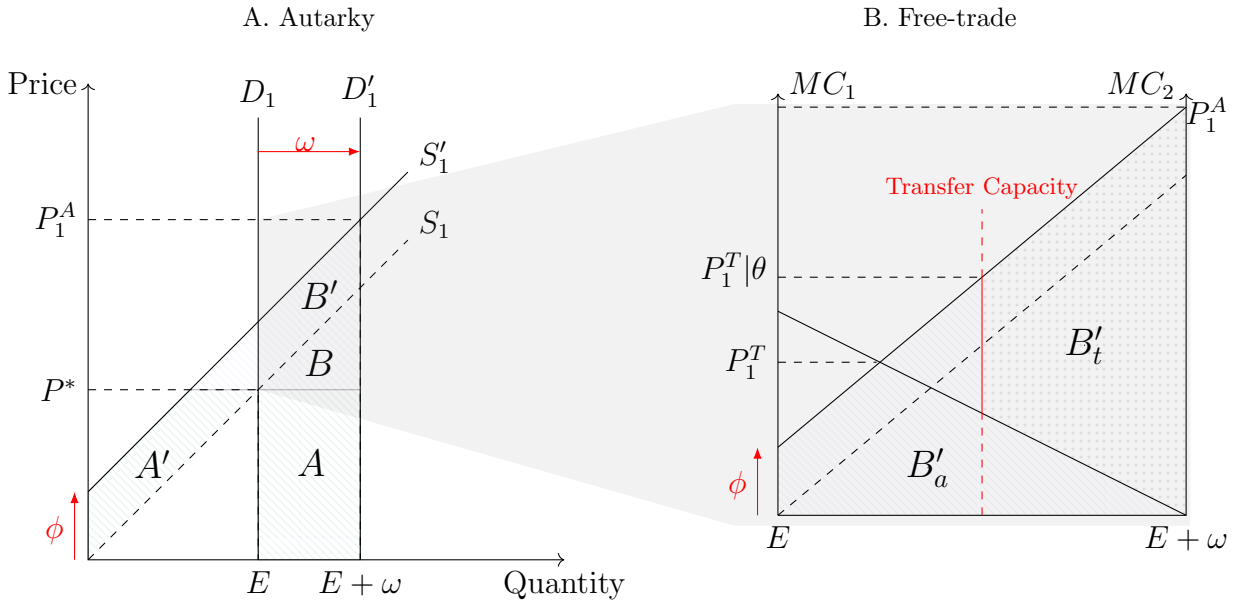

Market coupling

As described in the theoretical model, the level of interconnection management θ is expected to lower the impact of temperature on electricity market prices. I use a change in market operators coordination level to estimate how market integration mitigates the marginal impact of temperature shocks on prices. *Market coupling* is the name given to the "merger" between regional operators. National markets often have slightly different operating rules (timing of the bids, minimum and maximum price etc.) which makes difficult for the agents to both bid on the foreign spot market and book capacity rights at the interconnection. With market coupling, markets operators are simultaneously and cooperatively clearing the market, there is no need for parallel interconnection capacity booking as the operators clear both markets at the same time. Figure 1.3 summarizes the direction of international flows between France and Spain before and after the market coupling on May 13th, 2014. Each point reports interconnection usage in percentage of the available capacity, points in quadrants I and III are "counter intuitive" flows where the exporting (importing) country has the lowest (highest) price. Even within quadrants II and IV, all points should be on the -100/100% lines as it is not efficient to leave some capacity unused if the price spread is different than zero (away from the vertical dash line). It appears clearly from figure 1.3 that post market coupling the interconnection is used in a much more efficient way (equivalent to an increase in θ),

which makes it an ideal candidate for a difference-in-differences estimation. In particular, market coupling had a overnight and sharp impact on the interconnection usage. Figure A.3 presents the time series of counter intuitive flows and overall border usage, it appears clearly that market coupling had an instantaneous and lasting effect. For a more detailed discussion on market coupling, see (Keppler et al., 2016).

Figure 1.3: Directions of commercial flows before and after market coupling

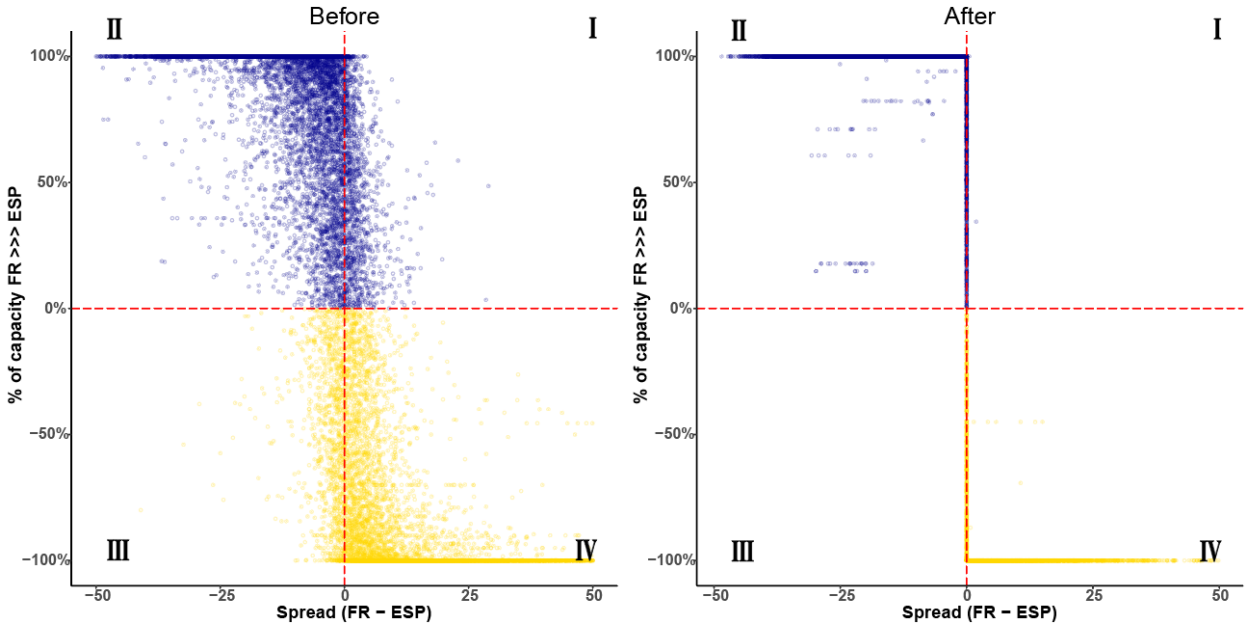

Several market coupling happened in Europe since the early 2000', this chapter focuses on the main one in Central Western Europe (CWE) in November 2010 between Austria, Belgium, France, Germany, Luxembourg, and the Netherlands. The introduction of market coupling is not random, the four countries were already trading with each others although market coupling is only ensuring that existing interconnector are properly used. Several reasons lead us to think that market coupling adoption process is not a threat to the identification strategy. First, it is possible to test for the parallel trend assumption since I observe long time series before and after the adoption of market coupling. Even if treated countries were more inclined to benefit from trade and had a higher sensitivity to temperature variation, it is not a threat to identification as long as the difference in baseline with the control group remained constant before the adoption. Second, weather variation is only predictable on the short run and there is no reasons to believe that treated countries implemented the coupling to prevent the adverse effect of future shocks. Finally, one could argue that market coupling adoption and temperature are both correlated with the significant renewable development in Germany. However the wind generation in Germany at that time was comparable to other neighbors (eg. Spain, and to a lower extent the UK and Denmark) who did not

take part in the market coupling, figure A.4 in appendix reports no significant differences in trends.

1.3 Data

In this study I combine wholesale electricity markets data with actual grid measurement. European national electricity markets were operated by monopolies for over a century until the European commission started massive deregulation and market integration policies in the 1990'[10]. Thanks to the member states' efforts, most of electricity data in Europe is now centralized and well accessible. In addition to electricity related data, I collected weather variables at a refine temporal and spacial level for each bidding zone.

Spot markets

Equilibrium Prices

I collected hourly prices for all the bidding zones providing public access data in Europe. Figure 1.4 shows the 24 countries (31 bidding zones) in the sample and specify whether book order data are available (see *supra*). The analysis is flexible enough such that I can either pooled the average daily price per bidding zone or to treat the 24 different hours as different "individuals" since they are defined simultaneously[11]. Figure 1.5 plots the monthly average price in Euro for each bidding zone with the thick dark line being the load weighted average price. Prices are following similar seasonal patterns but with different magnitudes. They are also converging through time, although the global price convergence is increasingly threaten by intermittent renewable development starting in the 2010'. Table 1.1 summarizes day-ahead prices' series for each bidding zone. We see that average price is similar in central Europe and significantly lower than in Italian bidding zones. As discussed in section 1.2, when markets are not coupled they can have different operation rules, for instance table 1 shows that minimum and maximum legal bidding prices are different. In the Italian markets actors cannot bid negative prices contrary to the German-Austrian-Luxemburg market (the minimum price is -56.87 on December 26, 2012). Negatives prices usually arise when it is cost efficient to stay turned-on due to very high ramping cost, the unit is suffering losses at t but should still be profitable on aggregate. Finally, data availability is different for each bidding zone, figure A.1 in annex reports the time coverage by country and whether we have the detailed book order data. The next section explains what exactly is the aggregated book order data and the auction system.

[10] Directive 96/92/EC

[11] There is no additional information in the price at 1:00 pm with respect to the price at 2:00 pm since they are determined at the same time, each day at 11am.

Table 1.1: Price summary statistics per bidding zone

	Mean	Std. dev.	Minimum	Maximum
Belgium	46.39	18.45	−40.99	314.27
Czech Republic	43.24	19.87	−24.61	203.38
Estonia	36.22	9.62	6.30	124.77
Finland	34.81	8.99	6.30	94.78
France	45.57	21.37	−40.99	612.77
Germany Austria Luxembourg	42.49	17.46	−56.87	301.54
Hungary	44.98	26.41	−6.71	1 147.96
Latvia	44.10	12.36	15.91	126.32
Lithuania	44.37	12.36	15.91	126.32
Netherlands	47.28	16.25	14.83	277.41
Norway	34.85	14.24	2.07	95.76
Poland	37.46	10.61	16.57	94.53
Portugal	46.83	14.72	0.00	93.35
Romania	41.39	13.32	2.99	91.62
Slovakia	39.05	12.03	−42.72	99.77
Slovenia	45.33	13.72	0.00	137.92
Spain	45.62	13.68	0.00	93.11
Switzerland	49.65	18.44	4.16	179.90
UK	43.29	8.12	12.15	169.65
Sweden				
SE1	30.18	9.53	3.51	88.10
SE2	30.18	9.53	3.51	88.10
SE3	38.59	15.65	3.51	231.51
SE4	39.00	15.58	3.51	231.51
Italy				
Center South	63.13	17.12	0.00	146.85
Center North	63.18	16.71	0.00	146.85
North	62.26	16.24	0.00	146.85
Sardaigna	68.72	25.07	11.15	273.85
Sicilia	80.51	27.43	0.00	273.64
South	61.75	17.19	0.00	146.85
Denmark				
West	38.00	13.88	−38.43	178.20
East	40.53	16.01	−38.38	231.51

Figure 1.4: Spatial coverage

Book order

Market operators in central western Europe (EPEX), Great-Britain (APX), Scandinavia & Baltic (Nordpool), Iberian Peninsula (OMIE), and Italy (GME) granted exceptional access to aggregated curve data. Aggregated supply and demand curves are constructed from the detailed book orders. Market operators first sum the quantity asked and offer at a given price, and then clear the market at the point where supply meets demand[12]. The full book order by agent is only available for a limited number of markets (eg. Iberian Peninsula) but most of the operators only communicate the aggregated supply and demand curves. The sole difference is that it is more complicated to analyze market power if one only observes the market's aggregated curves. However, it can be shown that market power is a limited concern in Europe since electricity is traded through zonal pricing. Market size is then at the country level and it is more difficult to use internal congestion to artificially increase market power.

Book order and aggregated curves can be used to analyze market integration. Note

[12]linear interpolation is sometimes necessary

Figure 1.5: Monthly price per bidding zone

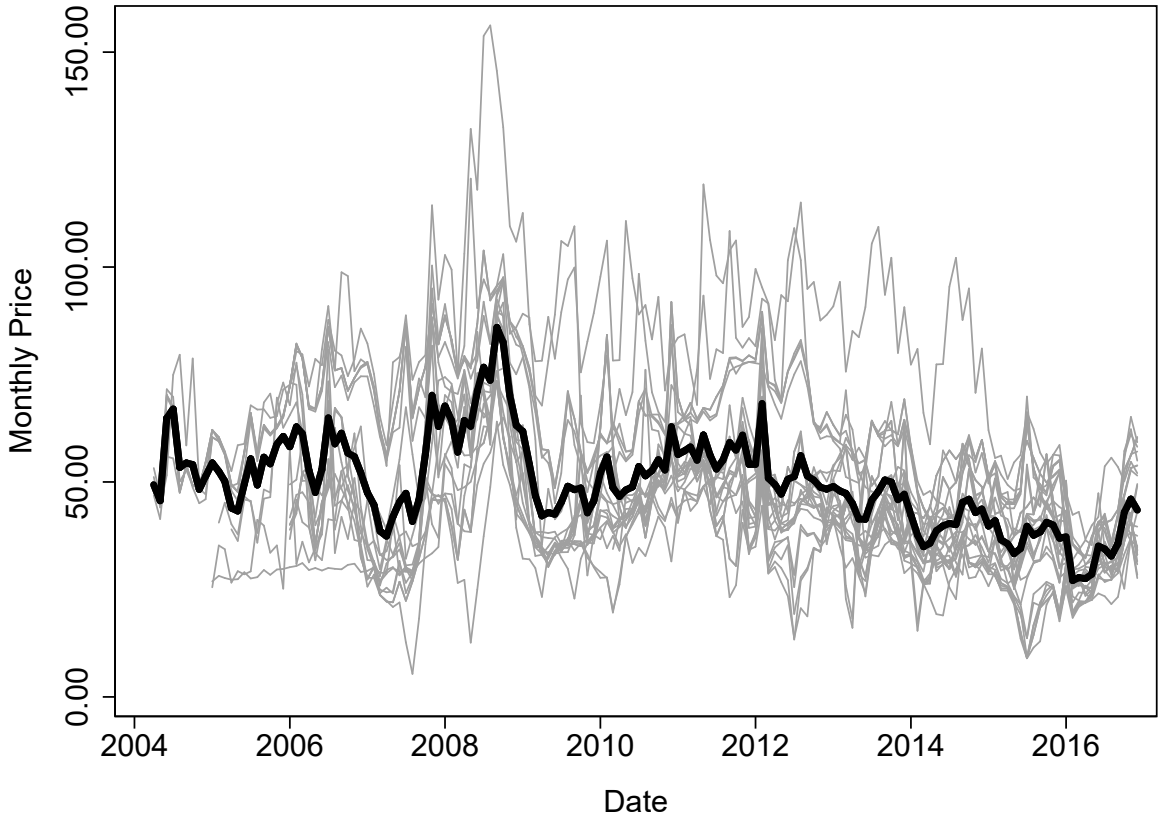

that even if national markets are connected by cross-border lines, the two zonal prices will diverge as soon as the transmission line is congested. One could simulate alternative prices for interconnection expansion scenarios by shifting the aggregated demand curve of the exporting country to the right and the supply curve of the importing country by the same amount. Figure 1.6 illustrates a concrete example for May 5th, 2013 at 10:00pm. At this specific hour, the German price was higher than the French price (which necessarily means congestion at the interconnection). I reconstructed prices for an expansion of 500MW by simultaneously comparing the French theoretical price for a demand 500MW higher with the German price if the supply was 500MW higher. Note that in this example, increasing the interconnection capacity by 500MW would not have been enough to reach a full price convergence but a 1000MW expansion would have been sufficient (and possibly oversized). The effect of increasing the interconnection capacity is different in the two countries: in this example it has a small impact in Germany but French prices are affected to a higher extent (plus 4 euros). This would have generated a surplus gain both for German consumers and French producers, and would have lowered the *transmission system operators*' congestion

rent[13].

The whole demand and supply curves are needed to precisely estimate the effect of temperature changes on electricity cost. In the theoretical model, I assumed that a temperature shock δ would affect homogeneously the demand (supply) by a common factor ω (ϕ). In reality, temperature will affect heterogeneously the different agents in the market. Hydro unit and nuclear power that need cooling water will be much more affected than other type of power plant. As shown in the theoretical model, this could matter as only the units above the marginal cost at the equilibrium (pre-shock) determine the value of the interconnection. If base-load[14] power plant are affected by temperature but marginal producers are not, then one should not expect a change in price. To estimate the effect of temperature shocks on the full curves, it requires to decompose the curves in 1€ uniform slices for every *zone-hour-day* in order to compare the curves' shape over time and space. Concretely, I first linearly interpolate the multiple bid and ask so that I get smooth demand and supply curves and then I create time series of comparable quantity offers and demanded for every 1€ steps using hourly data from 2006 to 2016.

Book order data are also useful to compute theoretical counterfactual prices. For instance, one can find for each hour of the sample the theoretical price that would have cleared the market if there were no congestion on the grid, *ie.* the equilibrium price in a pure open trade framework ($\theta = 1$). The congestion free price is simply given by the intersection of

[13]Congestion rent is collected equally by the two TSOs and is equal to the price difference multiplied by the interconnection capacity. Therefore, reducing the price spread by increasing the interconnection capacity would lower the TSOs' rent but not necessarily their utility since their main objective is to maximize global welfare

[14]Low marginal cost unit such as Nuclear power plants that always produce

Figure 1.6: Interconnection expansion between France and Germany - May 5th, 2013

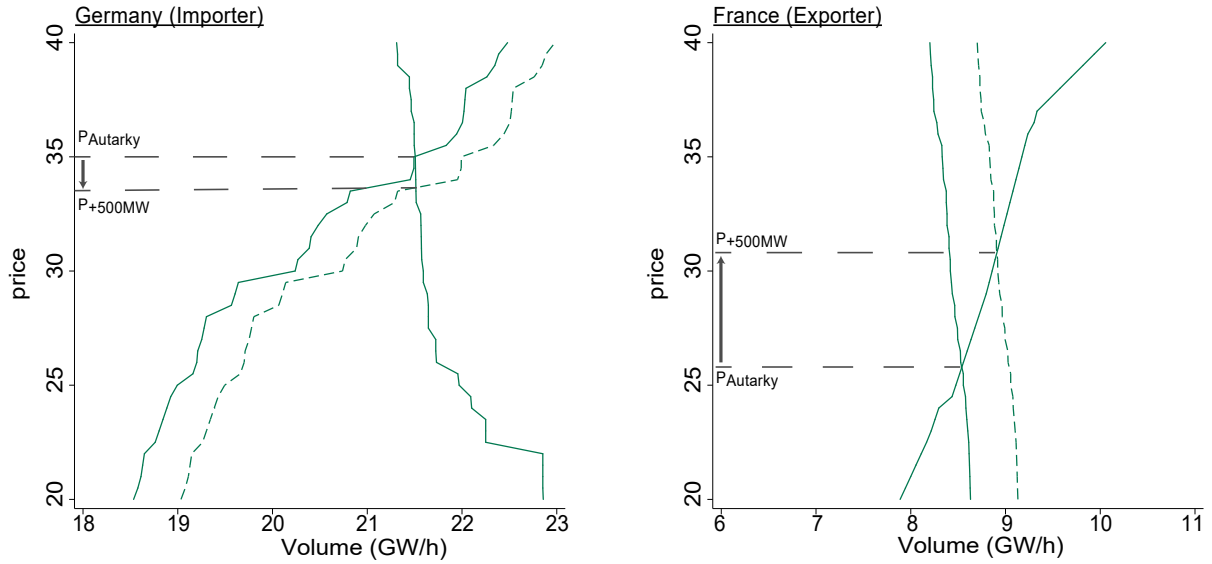

the aggregated demand and supply curve of the two (or more) countries, for the example illustrated by figure 1.6 the congestion-free price is €33.4. As an example, figure A.2 in appendix illustrates the four steps process for the German demand curve at 5:00 pm on January 1^{st}, 2006; in particular I had to use linear interpolation when data are not granular enough.

Grid data

Aggregated Consumption

Since demand and supply have to be equal at all time, I can proxy the equilibrium demand using the total load on the grid. Total load is by definition a subset of book order data but it is available for a much larger period of time. I use load data to estimate an equilibrium quantity dose-response to temperature *à la* Auffhammer et al. (2017). ENTSO-E is reporting hourly load data for every European countries but only for the period post 2015. I used data from market operators and national TSO to cover as many years as possible, but estimating the electricity consumption response to temperature shock is not the primary interest here as it has been well estimated in the literature (eg. Deschenes and Greenstone, Auffhammer et al., Barecca et al.). In the robustness analysis load data is used as a weighting scheme. Note that load and consumption data are not perfect substitutes, the load is the current amount of electricity supported by the grid at a specific point in time, it does not disentangle between consumption, generation, losses, and international trade.

Generation

Temperature and wind patterns are closely related and could threaten the identification strategy. Intermittent renewables are automatically dispatch to the grid but are still taken into account in the book order. Solar and wind generation usually bid at zero (or negative) prices if the market operator allows it.

Wind generation in particular has to be accounted for since it is a complex function of wind speed, air pressure (temperature) and humidity. The renewable data comes from each country's Transmission System Operator (TSO).

Weather

European Centre for Medium-Range Weather Forecasts (ECMWF) Climate variables are based on reanalysis data from the European Centre for Medium-Range Weather Forecasts (ERA-Interim), which is based on a climate model combined with observational data. I used their 0.25 x 0.25 gridded data on daily temperature and precipitation to generate aggregated daily temperature at the bidding zone level using population weights[15]. Population weights ensure to report an average of the temperature for places that matter for this study. Table

[15]Cross-section Landsat data, 2011

?? in appendix is reporting summary statistics for temperature and precipitation at the bidding zone level. Climate variation is roughly similar throughout continental Europe, the warmest region is southern Italy and the coldest northern Sweden. Temperature shocks are expected to be correlated over space but using daily level data provides additional variation.

1.4 Identification strategy

Estimating the effect of market integration

The theoretical model predicts that market integration lowers electricity price of the shocked country regardless of the relative difference in marginal cost. To test this prediction I first need a credible and varying measure of market integration. I chose to use the market coupling adoption in a *difference-in-differences like* design as a proxy for market integration since it is strongly correlated with trade efficiency (see figure 1.3 and Keppler et al. (2016)). The data span from 2006 until the end of 2016 and gives us almost five years of pre-treatment data since market coupling occurred in November 2010. I used temperature and prices moments of the 31 bidding zones to find the closest control to the treated group composed of Austria, Belgium, France, Germany, Luxembourg, and the Netherlands. Figure 1.7 shows very similar temperature and price distribution for the control and the treated group. However considering the large size of the sample, mean daily price and temperature are statistically higher in the treated group (€45.24 vs €41.45, and 10.52^oC vs 9.34^oC).

Figure 1.7: control and group hist

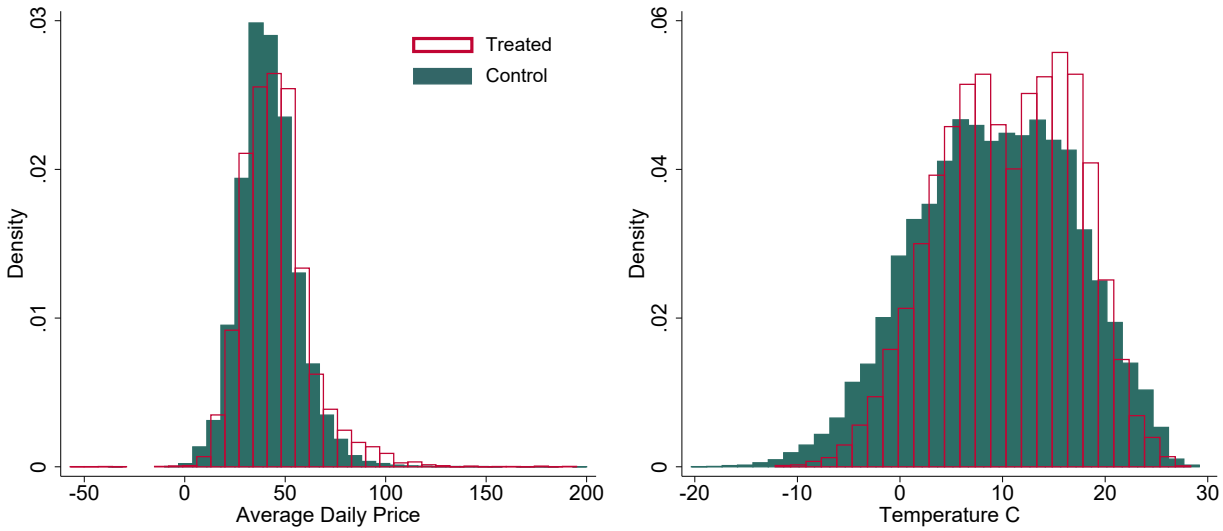

DiD framework does not require the two groups to have similar level of output, what matters for identification is the similarity of the trends prior to the intervention. Let consider

the following price equation:

$$log(P_{z,t}) = \sum_{k=1}^{3}(\beta_k temp^k) + \sum_{k=1}^{2}(\theta_k precip^k) + \psi_{DOW} + \alpha_{cntry,year} + \delta_{cntry,month} + \epsilon_{zt} \quad (1.10)$$

Where price in a given zone z in day t is a nonlinear function of temperature (3^{rd} order polynomial) and precipitation. Once I control for non observable at the space (country by year FE) and time level (day-of-week and month FE), I can estimate the marginal effect of temperature since weather variation is purely exogenous with respect to prices. To estimate the mitigating impact of market integration, I do not want to compare a *difference-in-level-differences* but the difference in "marginal effect" differences before and after market coupling. I am interested in a second order parameter, how the adoption of market coupling mitigated the *marginal effect* of temperature on prices. More formally, I interacted the treat and post dummies with each temperature variable[16]:

$$\begin{aligned} log(P_{z,t}) &= \gamma_1 \mathbb{1}(treat_{z,t}) + \gamma_2 \mathbb{1}(post_{z,t}) + \gamma_3 \mathbb{1}(treat_{z,t})\mathbb{1}(post_{z,t}) \\ &+ \sum_{k=1}^{3}[temp^k(\beta_k + \phi_k^1 \mathbb{1}(treat_{z,t}) + \phi_k^2 \mathbb{1}(post_{z,t}) + \phi_k^3 \mathbb{1}(treat_{z,t})\mathbb{1}(post_{z,t}))] \\ &+ \sum_{k=1}^{2}(\theta_k precip^k) + \psi_{DOW} + \alpha_{country,year} + \delta_{country,month} + \epsilon_{zt} \end{aligned} \quad (1.11)$$

For causal interpretation, it is important to verify that the marginal effect of temperature on price evolved through time (or not) in a similar pattern for both groups. As mentioned in the model section, temperature can affects price through multiple channels: a demand shock, an increase in marginal cost, or even a reduction in transmission line capacity. It is not a threat to the identification strategy as long as these relationships did not evolve differently in the control and treated group. To test the parallel assumption, I limited the sample to the pre-treatment period and regressed electricity prices on temperature variables interacted with a time trend. More formally:

$$\begin{aligned} log(P^g_{z,t}) &= \sum_{k=1}^{3}[temp^k(\beta_k^g + \eta_k^g t)] \\ &+ \sum_{k=1}^{2}(\theta_k precip^k) + \psi^g_{DOW} + \alpha^g_{country,year} + \delta^g_{country,month} + \epsilon^g_{zt} \end{aligned} \quad (1.12)$$

If the difference in the marginal effect of temperature remains stable through time between the two groups g, I expect the trend coefficients (η_k) to be equal between groups.[17]

[16] the un-interracted treat variable is actually dropped to avoid collineratity when interacting with the temperature variables

[17] Precisely it is necessary to verify the parallel stability of the marginal effect all along the temperature distribution (*ie.* is the effect of a day at 25^{o}C stable? 28^{o}C? -10^{o}C? etc.)

Separability of the demand and supply reactions

In most studies about supply and demand estimation, one needs to combine exogenous shifter with equilibrium price and quantity. In this case, I actually observe the whole curves but want to estimate the (temperature) shifters. The empirical strategy is very similar to equation (10), but applied to quantities ask and bid at a given price[18] and hour:

$$\forall p, h \in \{0, 150\} \times \{1, 24\}$$
$$\begin{cases} log(Ask_{z,t}^{p,h}) = \sum_{k=1}^{3}(\beta_k temp^k) + \sum_{k=1}^{2}(\theta_k precip^k) + \gamma^A X_{z,t} + \psi_{DOW} + \alpha_{c,y} + \delta_{c,m} + \epsilon_{zt} \\ log(Bid_{z,t}^{p,h}) = \sum_{k=1}^{3}(\beta_k temp^k) + \sum_{k=1}^{2}(\theta_k precip^k) + \gamma^B Y_{z,t} + \psi_{DOW} + \alpha_{c,y} + \delta_{c,m} + \epsilon_{zt} \end{cases} \tag{1.13}$$

where $X_{z,t}$ and $Y_{z,t}$ are *side-of-market* specific covariates. Considering that I can disentangle the effect of the covariates on supply and demand, I should control in the appropriate equation for the variables correlated with both temperature and one side of the market only. For example, wind generation is expected to shift quantity supplied at every price since wind farms bid at the lowest price possible[19], and wind is correlated with temperature. However, I do not expect quantity demanded to change with wind generation. Other covariates such as solar radiation should be in both $X_{z,t}$ and $Y_{z,t}$, on the one hand it increases supply through solar farm, and on the other hand it decreases the quantity demanded through personal photovoltaic. The set of fixed effect is similar to equation (1.10).

1.5 Results

Net effect of temperature on prices

The theoretical model disentangled the multiple channels through which temperature affects price, in this section I first want to recover the net effect on wholesale price, which is equivalent to estimate equation (1.10). I used the full sample over the period 2006-2016, and the results suggest that temperature has a non-linear effect on prices - just like on consumption. Figure 1.8 plots the price dose-response[20], *ie.* the marginal effect of temperature relative to a "normal" day at 15^oC. In order to easily interpret the results and compare the dose-response with the usual effect of temperature on quantity, I logged the left hand side variable. Results suggest that - everything remaining constant - prices would be 18-26% higher for a day at 25^oC (relatif to a day at 15^oC) and 34-51% for a day at -10^oC. In comparison, the dashed line in figure 1.8 represents the response of quantity due to change in temperature, a day at 25^oC would only increase quantity at the equilibrium by 4-8% and a cold day at -10^oC

[18] I limited the study to the points between €0 and €150

[19] The marginal cost of wind farms is zero.

[20] Estimated coefficients are not easily readable considering the nonlinear relationship. Table of result are available for reference in appendix.

by 19-26%[21]. First, previous studies showing that consumption react to temperature in a non-linear u-shaped way are confirmed by these results. Second, it appears that all along the temperature distribution, prices react in a larger magnitude that quantity. If the demand only were to change when temperature changes, it would means that the slope of the supply curve is at least greater than 1. Since I only observe the net effect, I cannot make any claim about the marginal cost. Still, the result suggest that the existing literature that solely focus on change in quantity might underestimate the energy cost of climate change.

Figure 1.8: Daily price sensitivity to temperature (pooled model)

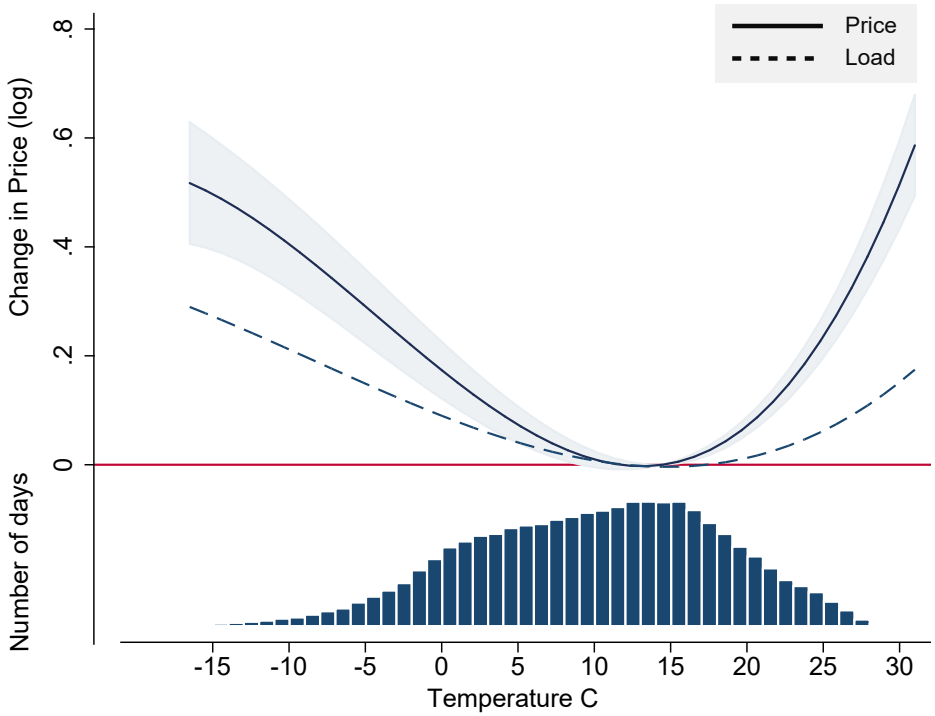

The non linear relationship is robust to the choice of functional form, figure A.5 in appendix shows similar dose-response when I use other semi-parametric method such as bins regressions (scatter) and piece-wise linear (dashed). Polynomial allows curvature in the dose-response and keep the number of parameter to estimate to a much lower level compared to bins regressions. The last point is critical now that I move to the interacted model.

Adaptation through market integration

I now turn to the estimation of equation (1.11), figure 1.9 reports the dose-response before (dashed line) and after market coupling in both the control and the treated group. These

[21]Confidence interval are not reported on the figure for simplicity purpose but the range in the text correspond to the 95% CI.

dose-responses are plotted based on sub-sample estimation of equation (1.10), it confirms the absence of trend in the control group as the dose-responses on the first panel are almost the exact same before and after market coupling. Quite differently, the second panel confirms the theoretical model, the adoption of market coupling appears to mitigate the marginal effect of temperature all along the temperature distribution. To formally estimate the difference-in-differences, I need to estimate equation (1.11). Results are summarized in table 1.2, the marginal effect of temperature is predicted at 0^oC and 25^oC to easily interpret the non-linear interacted coefficients. As suggested by the plot of the sub-sample dose-response, the interacted DID coefficients are significant and decrease the magnitude of the temperature effect. The results suggest that market coupling decreased the effect of a hot (cold) day by 14-28 (5-11) percentage points. Note that the market coupling also had a pure effect on the level of prices, reducing on average the wholesale prices by 9.3 percentage point - regardless of the temperature. It is a result in line with the theory, under the existence of some heterogeneity in the marginal costs, any increase in trade opportunity is cost efficient. Quantitatively, the base effect of market integration is comparable with (Cicala, 2017) which found using market deregulation in the US as a quasi-experiment that trade would decrease the cost of provision by roughly 20%.

Figure 1.9: Daily price sensitivity to temperature (DID)

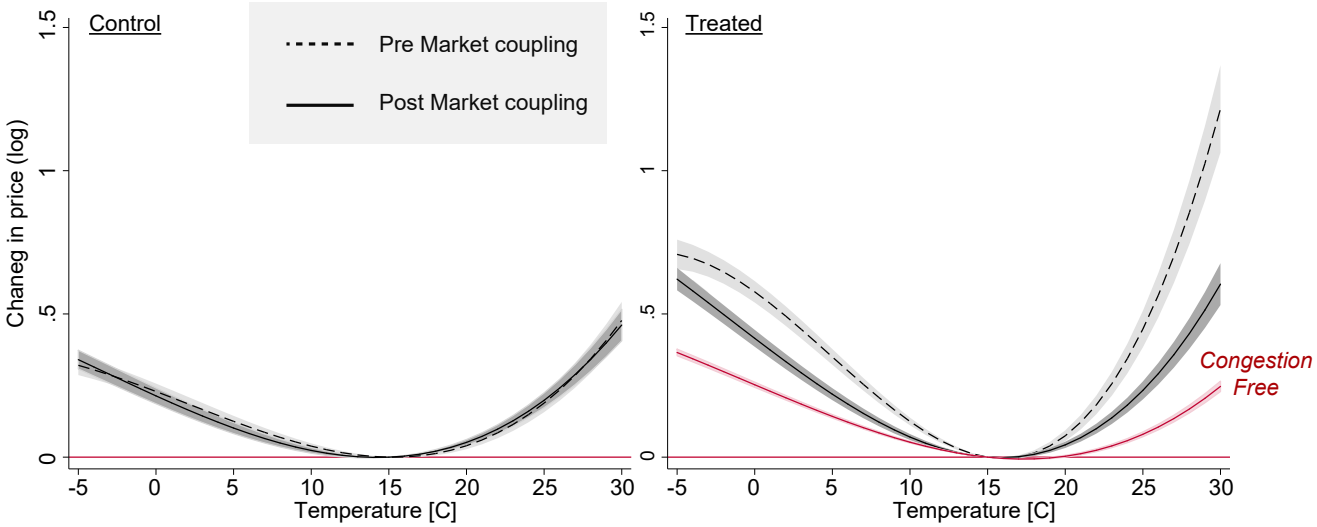

One could qualify the adoption of market coupling as the *extensive margin* of trade, in comparison with an increase of the interconnection capacity. An imperfect measure of the *intensive margin* can be estimated using book order data. I aggregated the nationals' supply and demand curves for each member of the treated group and re-estimated equation (1.10) using the theoretical congestion-free prices[22]. To facilitate the comparison with the

[22]This method is actually closer to the extensive margin as it does not look at small change in intercon-

adoption of market coupling I restrained the sample to the post coupling period only, the second panel of figure 1.9 has side by side the dose-responses for the real prices (blue) and the counterfactual (red). The magnitude of the drop in marginal effects are roughly of the same order. It definitely plays in favor of market coupling considering the prohibitive construction cost of interconnection lines. In addition, the *congestion free* scenario is particularly extreme, it is purely theoretical and would be found way oversized in a cost benefit analysis.

Table 1.2: Difference-in-*marginal-effect*-differences

	Log(Price)		Predicted marginal effect (ref. point 15^oC)	
VARIABLES	coef	se	$T = 0^o$C	$T = 25^o$C
$\mathbb{1}\{Treat\}$	omitted			
$\mathbb{1}\{Post\}$	0.0623***	(0.0155)		
$\mathbb{1}\{Treat\} \times \mathbb{1}\{Post\}$	-0.0927***	(0.0226)		
$Temp$	-0.0192***	(0.000859)		
$Temp^2$	-0.000196***	(6.46e-05)	$[0.20, 0.24]$	$[0.21, 0.27]$
$Temp^3$	3.55e-05***	(2.64e-06)		
$Temp \times \mathbb{1}\{Treat\}$	-0.0196***	(0.00224)		
$Temp^2 \times \mathbb{1}\{Treat\}$	-0.00125***	(0.000264)	$[0.25, 0.32]$	$[0.15, 0.29]$
$Temp^3 \times \mathbb{1}\{Treat\}$	8.12e-05***	(9.60e-06)		
$Temp \times \mathbb{1}\{Post\}$	-0.00758***	(0.000870)		
$Temp^2 \times \mathbb{1}\{Post\}$	0.000369***	(8.14e-05)	$[0.01, 0.05]$	$[-0.05, 0.01]$
$Temp^3 \times \mathbb{1}\{Post\}$	-9.98e-09	(3.16e-06)		
$Temp \times \mathbb{1}\{Treat\} \times \mathbb{1}\{Post\}$	0.00412*	(0.00243)		
$Temp^2 \times \mathbb{1}\{Treat\} \times \mathbb{1}\{Post\}$	0.000961***	(0.000287)	$[-.11, -.05]$	$[-.28, -.14]$
$Temp^3 \times \mathbb{1}\{Treat\} \times \mathbb{1}\{Post\}$	-5.97e-05***	(1.03e-05)		
$Precipitation$	-0.00972***	(0.000845)		
$Precipitation^2$	0.000249***	(4.93e-05)		
Fixed effects	Day-of-Week, Country-by-Month, Country-by-Year			
Observations	45,393			
R-squared	0.567			

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

nection capacity - which is possible - but solely at the extreme case.

Gains from trade

The difference-in-differences approach suggests that market coupling decreased average daily price by roughly 9% in the treated countries. Using detailed data on book order I can backed out the potential gains from trade that would have occurred in a congestion-free world. Using Shephard's Lemma and assuming no income effect, the change in consumer surplus is roughly defined as the price difference between the two states times consumption. On the supply side, assuming perfect competition and applying Hotelling's Lemma I approximated the change in profits as the change in price times the production volumes. Table 1.3 reports change in surplus and net gains for each country. At the exception of Germany, consumers would have experienced net gains whereas producers would suffer losses. This is directly coming from the significant flows of intermittent electricity produced in Germany that make Germany the biggest exporter in the region. One needs to be cautious about the interpretation given to those numbers as wholesale electricity market prices are not directly charged to the end users and similarly a decrease in producer surplus can hide long run efficiency gain of the supply side. Next section discusses more extensively the welfare and policiy implications of these results, in particular with regards to climate change.

Table 1.3: Gains from trade in a congestion-free state (11/2010-01/2016)

	Euros (millions)		
	Δ Consumer Surplus	Δ Producer Surplus	Net Gains
Belgium	704	-664	40
France	1,170	-939	231
Germany	- 3,760	3,930	170
Netherlands	1,330	-1,190	140

1.6 Discussion

Interpreting the gains from trade necessitates to understand the challenges faced by energy actors in each country. Germany has been struggling with a surplus of intermittent electricity depressing wholesale prices toward zero and sometimes even below. More conventional power plants need to run a certain number of hours to be profitable, which means that intermittent excessively low prices diminish the incentive to invest in conventional plants without providing a security of supply during peak hours[23]. The positive (negative) externality associated with conventional (intermittent) units is not currently internalize by the market. However, integrating the different regions by coupling the markets would mitigate the intermittent excessive price fall and benefit in the long run (a) the consumers through

[23]In Central Western Europe, peak hour is often between 6pm and 8pm, when solar production is low or null

security of supply and (b) the producers that need a stable price signal to invest and remain profitable. As shown by Annan-Phan and Roques (2018) market integration not only lowers the global average price but also decreases price volatility in each market.

The change in consumer surplus itself has to be interpreted with caution. Wholesale and end-user electricity prices are only related in the long run and are certainly not equal on average. As explained above, near zero day ahead prices affect long run supply and thus will likely be associated with an increase in end-user prices. Out of the 3.8 billions euros of change in German consumer surplus, only a fraction would actually be passed on the end-users. 22.8% of the German consumers' losses (€716 millions) would come from a price re-evaluation when prices were excessively low ($\leqslant$ 20€/MWh). Considering those factors, the interpretation of surplus variation is limited and one should focus on the efficiency gains coming from the global reduction in marginal cost. Comparing the slope of the supply curves around the equilibrium[24] gives the price reaction associated with a small shock in consumption (*eg.* climate shock). Assuming linear cost around the equilibrium, I can infer the increase in marginal cost induced by a given demand shock. Figure 1.10 reports the effect of a 500 MWh demand shock in each country over time in comparison to the congestion-free reaction. On average, the shock would increase the price in France by €7, the same shock in Netherlands or Germany would have increased price by €13 and €4, respectively. In a congestion-free world, a single shock in any of those country would only increase price - everywhere - by 50 cents.

Figure 1.10: The price reaction of a 500MW demand shock (2011-2016)

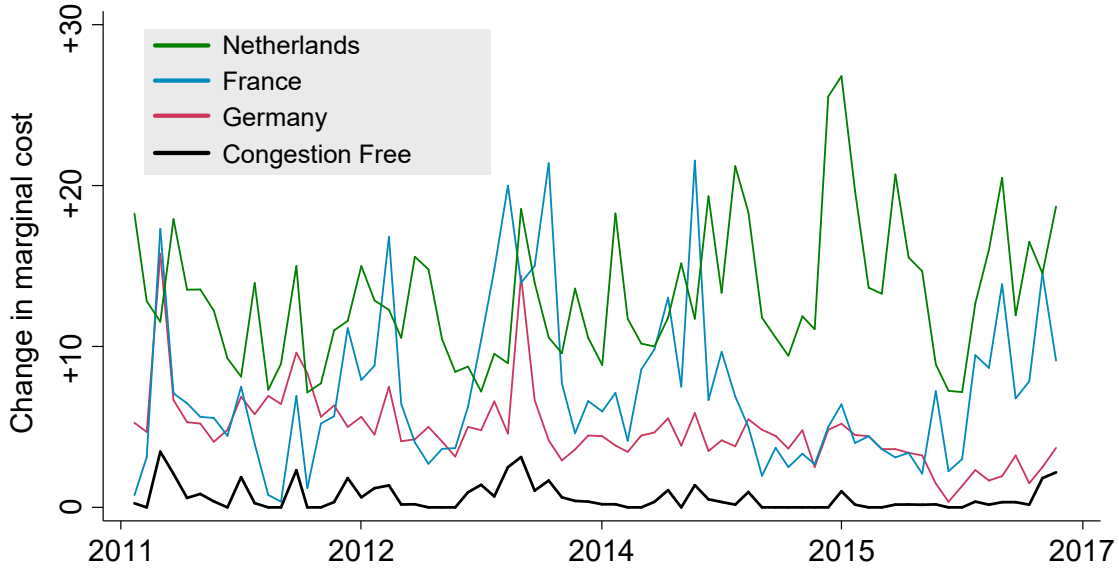

What are the implication for climate change mitigation? Just like in the model the

[24]I estimated the slope of the supply curve each day for prices ranging from -25% to +25% around the equilibrium price

demand is affected not only by warm days but also during cold shocks. Results indicate that trade - either through market coupling or grid expansion - would mitigate the provision cost of such events. It is not clear whether climate change would exacerbate or diminish the need for market integration. On the one hand, more demand shock are expected during the summer but on the other hand, the heating demand (which is roughly as sensitive to temperature as illustrated in figure 1.8) will likely decrease. The shift in temperature distribution is not necessarily homogeneous and a recent studies suggests that weather variance will increase (Trenberth et al., 2015) or that cold shocks might arise in Western Europe and Northern America due to the slowdown of oceanic stream (Rahmstorf et al., 2015). In addition, the sample of countries used in the gains-from-trade calculation are not very sensitive to hot temperatures compared to other countries in the world, even in Europe (*eg.* Greece). The symmetry of the temperature dose-response suggests that market integration is mitigating in the same fashion hot and cold days, meaning that gains from trade with respect to climate change might be much larger for countries only sensitive to hot days.

In this chapter, I presented a new method to study the impact of climate change on electricity expenditure through the analysis of wholesale prices. I showed that relative to consumption, prices were reacting in a larger proportion to temperature shocks. Supported by a theoretical framework, I then estimated the benefit of market integration with regards to climate damages mitigation. Applying a *difference-in-differences like* method I estimated the causal impact of market coupling on the temperature dose-response. Market coupling adoption reduced the impact of a day at 25^{o}C on prices by 14-28 percentage point. The effect is statistically and economically significant, in particular if compared to the dose-response of the congestion-free counterfactual where interconnection are assumed unlimited. Market coupling effect has roughly half the magnitude of the unlimited grid expansion, although achieving a congestion free world would be prohibitively costly. More research need to be conducted to disentangle precisely the different channels trough which temperature affects supply and demand in order to derive welfare conclusions. However, the reduced form estimation documented the potential efficiency gains of trade, in particular with regards to climate change.

Chapter 2

Hot Temperatures, Aggression, and Death at the Hands of the Police: Evidence from the U.S.

2.1 Introduction[1]

For every ten murders in the U.S., one civilian dies as the result of an encounter with the police.[2] Police killings have increasingly captured the attention of the public, with much of the focus on the racial disparities in police shootings and the role of race in influencing police action (MacDonald and Fagan (2019); Goncalves and Mello (2020)) and in particular the use firearm (Ross (2015); Fryer (2018a,b); Edwards et al. (2018); Knox et al. (2019)). The different types of threats faced by police officers, however, have been left unexplored, as have the reactions of the police to these threats. One threat in particular, hot weather, is often associated with an increase in violent behavior (Bell and Baron (1976); Jacob et al. (2007); Burke et al. (2015)). This chapter explores the risk associated from higher temperatures to the use of deadly force by the police.

We exploit county and monthly variation in temperatures to estimate the effect of weather on civilian deaths by firearm, conducted electrical weapons (CEWs, also known as Tasers), less-than-lethal weapons, and physical restraint. We document that an increase in temperature is associated with a higher level of risk for officers and bystanders. Consistent with the increase in danger and opportunity, we find suggestive evidence that fatal shootings also increase with temperature. However, when controlling for the increase in civilian-officer interactions, we find no additional impact of temperature on the number of fatal shootings. Extremely warm days have an additional impact on other types of force as well: we find that fatal interactions involving a CEW or physical restraint are significantly higher for any

[1] The material in this chapter was co-authored with Bocar Ba.

[2] Authors' calculation, See fig A.7 in the Appendix.

additional days above 32^oC.[3]

Our results indicate that officers' judgment over their use of firearms is not directly impacted by hot temperatures, contrasting with results obtained during in-training experiments Vrij et al. (1994). In addition, our findings suggest that one should be careful when deploying CEWs or techniques of physical restraints that could result in asphyxia when temperatures increase. Further research is needed in order to understand the physiological relationships between non-lethal tactics and high temperatures that result in an increased likelihood of death. We do not find evidence that temperature impacts the number of civilian deaths by other kinds of less-than-lethal weapons.[4]

Demonstrating these findings requires overcoming two empirical difficulties. First, there is a lack of readily available public data that records the interactions between civilians and police that result in the death of civilians (Banks et al. (2016); Ouss and Rappaport (2019)). Secondly, it is a challenge to disentangle whether these deaths are related to behavior, either from civilians or police officers, or result from an unintended physiological response to temperature, which would indicate that certain weapons are more dangerous during warm days.

This study addressed these challenges in three steps. Firstly, we used novel data from Fatal Encounters, which has crowd-sourced records on civilian-police interactions resulting in a civilian casualty in the United States between 2000 and 2016. Secondly, we used temporal variation in daily temperatures over the different geographical regions in the U.S. to account for county-level unobservables and seasonal patterns in crime data. Finally, we disentangled the nonlinear impact of temperature for each type of force in order to evaluate whether to interpret our results as a behavioral or physiological effect.

We first establish the link between temperature and threat level for police officers and civilians. We defined incidents as cases in which a suspect poses a serious threat of injury or death to officers or bystanders. We argue that these incidents, potentially dangerous, are captured by (1) the number of violent crimes and (2) the number of officers who are assaulted and/or killed. Variation in situational environment such as darkness (Doleac and Sanders (2015); Chalfin et al. (2019)) and weather ((Ranson, 2014a)) have been linked to changes in crime level. We show that temperature has a positive and statistically significant impact on the number of assaults on both civilians and police officers. This is consistent with previous studies that suggest that higher temperatures are associated with more aggressive behavior (Jacob et al. (2007); Burke et al. (2015)). We hypothesized that, in response to an increased threat level in warmer weather, there should be a corresponding increase in the number of civilian deaths.

In the second part of our analysis, we explore the effect of temperature on police-related deaths, controlling for the threat level and taking into consideration the officers' exposure to risk. We found suggestive evidence that fatal police shootings is impacted by hot tem-

[3]For the rest of the chapter, we define 'warm'to refer to temperatures between 17 and 24^oC (i.e. 63 and 75^oF); we use 'hot' to refer to temperatures between 24 and 32^oC (i.e. 75 and 90^oF); and 'very hot' or 'extremely warm' indicate to temperatures above 32^oC

[4]See Section 2.2 for a discussion of our categorization of force options.

peratures. The magnitude of the estimated effect is similar to the increase in threat levels. Consistent with the fact that higher temperatures are associated with a higher level of threat, officers are more likely to kill a civilian by firearm when temperature increases. Once we account for such increase in threat, temperature does not have an additional effect on the use of firearm. In contrast, our results indicate that the number of civilian deaths caused by the use of CEWs increases by 5.3% on very hot days, independently of the increase in opportunity. The effect of CEWs on casualties is null for days that are not extremely warm, indicating a physiological interaction with temperature only for very hot days. Similarly, we also document that the number of civilian deaths by physical restraint increases by an additional 15.1% during extremely hot days. Although physical restraints and CEWs are not categorized as lethal uses of force, these results suggest that such tactics can have unintended deadly consequences when used in extremely warm weather. The results for incidents involving other kinds of less-than-lethal forces (see Section 2.2) were too imprecise to draw any conclusions.[5]

This chapter contributes to different strands of the literature. First, CEWs are considered a less-than-lethal weapon. Some police departments have recently expanded their Taser arsenals, arguing that the increased availability of CEWs will lead to a reduction in police shootings (Bustamante (2017); Hinkel and Smith Richards (2017)). However, Ba and Grogger (2018) find no evidence that Tasers reduce firearm usage. Using a randomized controlled trial, Ariel et al. (2019) find that the introduction of Tasers leads to more use-of-force incidents.[6] This chapter provides additional evidence for the potential dangers of Taser use by demonstrating the possible consequences of using CEWs in hot weather. Second, despite widespread findings that higher temperatures are related to higher crime rates and more aggressive behavior, we show that higher temperatures have no statistical effect on the number of casualties from police shootings when controlling for the increase in opportunity.

2.2 Background

Use-of-force options

Most law enforcement agencies in the U.S. have policies that guide their use-of-force. The International Association of Chiefs of Police have defined use-of-force as the "amount of effort required by police to compel compliance by an unwilling subject." Depending on the

[5]We also found that temperature modestly impacts vehicular accidents, but our results were not statistically significant (See Section A.2 in the Appendix).

[6]Ba and Grogger (2018) find similar results in Chicago; however, their findings might be sensitive to their interpretation of the level of severity of Tasers, in the use-of-force model. The authors claim that Tasers can be used on active resisters; however the Chicago Police Department's use-of-force model indicates that Tasers can be used on both active resisters and assailants (CPD (2017)). As a result, the authors lowered the level of severity of Tasers and might be mischaracterizing the level of danger of Tasers. We later document that Tasers are the third leading cause of civilian deaths when interacting with the police behind firearms and vehicular accidents.

level of danger, the continuum of force includes verbal commands, physical restraint, less-than-lethal force and lethal force. Less-than-lethal force involves technologies to gain control of a situation. Deaths by vehicle are not part of the continuum.

Our data is crowdsourced information from public news articles that do not rely on administrative information from police departments. This explains why causes of death do not perfectly fit the use-of-force model in our case. Below, we provide definitions of the causes of death that are part of the use-of-force model in order to provide some context necessary to interpret our results.

Physical restraint We define physical restraints to be soft and hard techniques such as grabs, holds, joint locks, punches, or kicks to restrain or subdue the subject. The casualties from physical restraints correspond to "asphyxia/restrained" in the Fatal Encounters' classification scheme.

Less-than-lethal We define less-than-lethal force as involving intermediate weapons such as an impact weapon (e.g., baton) or chemical weapons. Although CEWs are considered a less-than-lethal weapon, we analyze the use of this weapon separately. Our definition of less-than-lethal corresponds to "Beaten or Bludgeoned with instrument" and "Chemical agent or Pepper spray" in the Fatal Encounters' classification. Officers may use less-than-lethal weapons if the subject is physically aggressive or exhibits assaultive behavior with an immediate likelihood of injury to self or others (Use of force project (2017)). Non-lethal force options are used to limit the escalation of conflict where employment of lethal force is undesirable or prohibited.

CEW CEWs are considered a less-than-lethal force. CEWs discharge a high-voltage and low-amperage jolt of electricity through a dart. The electrical charge overrides the subject's nervous system and should temporarily incapacitate him. The use of CEWs requires training and periodic recertification.

Firearm Firearms are considered lethal use-of-force. Officers are allowed to use lethal weapons, i.e., firearms, if they reasonably believe that a suspect poses a serious threat to the officer or another individual. Many police departments collect data on officer-involved shootings and systematically investigate them. The use of firearms requires training and periodic recertification.

CEWs

The most common CEWs used by law enforcement agencies are manufactured by Taser International (now Axon).[7] In addition to selling the device, Axon provides training and

[7]According to Forbes Taser International holds 95% of market share.

certification to sworn law enforcement officers, military personnel, and licensed professional security employees.

According to Axon's training material, CEWs have limited effectiveness on loose or thick clothing, low nerve or muscle mass, and obese subjects (Axon (2018)). This suggests that officers might be less likely to use CEWs when the temperature is low, because people tend to wear thicker clothes in such weather. Axon also recommends against using CEWs on elderly, pregnant, or low Body Mass Index (BMI) individuals, warning that use on such subjects could increase the risk of death or serious injury to the subject.

Due to the electrical discharge, CEWs frequently cause subjects to fall, thereby increasing the risk of bodily trauma, particularly when the fall occurs on a hard surface, such as a sidewalk. Other potential effects of CEW use that could increase the risk of sudden death include changes in blood chemistry, blood pressure, respiration, heart rate, adrenaline, and stress hormones. Last but not least, CEWs are also liable cause to cardiac effects. Additional factors that increase cardiac risks associated with the use of CEWs are the duration of the delivered electrical charge and the distance to the heart from where the dart impacts the body (Axon (2018)).

Axon strongly recommends that officers keep a detailed account of any incident in which they deploy a CEW. The company suggests the officer note their own actions as well as the subject's and document the subject's medical status. For law enforcement agencies, it is important to keep this information in case of a lawsuit or for internal investigations.

2.3 Data

Source

Our primary analysis used Fatal Encounters data spanning from January 2000 to December 2016. We supplemented this dataset with combined daily climate data from the European Centre for Medium-Range Weather Forecasts. Finally, we obtained the FBI's Uniform Crime Reporting Data System (UCR) from Kaplan (2018), who constructed a monthly dataset on a police-department-level basis for index crimes from 1968 to 2016. We merged climate data and Fatal Encounters data with FBI UCR data by county and month.

Fatal Encounters is a crowdsourced database that relies on publicly available news sources to track deaths that have occurred as a result of police interaction in the United States since 2000. Every entry is manually checked and properly sourced with an accompanying police report, when available, or other official statements. The dataset includes key variables such as the date of the incident, county and cause of death. The data also indicates whether the subject displayed symptoms of mental illness or substance abuse. However, this information is reported by the person who submitted the incident to Fatal Encounters and may therefore be unreliable or subject to reporting bias. Our final sample includes only fatal incidents that can be tied to police use-of-force or those caused by vehicles.[8]

[8]Cases of death by vehicle are included because they represent a large share of the RAW data (19.22%).

There is no federally managed national dataset on police killings in the United States. The Bureau of Justice Statistics (BJS) acknowledges their current data-collection methodology for arrest-related deaths is defective and could be improved by crowdsourcing (Arrest-related deaths program Assessment, March 2015). In particular, the BJS states that Fatal Encounters "most closely matched the Arrest-related deaths program scope" (Banks et al. (2016)). To our knowledge, this is the best source of information on this subject and has yet to be stringently analyzed.

Summary statistics

Table 2.1 presents the distribution of type of death from 2000 to 2016. Firearms and vehicular accidents represent the most common cause of death with 73.4% and 21.1% of the final sample, respectively. Death by CEWs is the third most common fatal use-of-force death and represents about 3.7% of arrest-related deaths. Meanwhile, less-than-lethal force and physical restraints correspond to about 1% of the sample. We also disaggregate the distribution according to age, gender, race, and symptom of mental illness or substance abuse of the civilian who died during the interaction. About 5-6% of the civilians who died by firearm, less-than-lethal weapon, and physical restraints are female, while only 2.5% of those who died by CEW are female. CEWs have the lowest share of juvenile and elderly in the sample, which is consistent with the safety guidelines provided to police officers by Axon.[9] About 31% of the civilians who die in police shootings are white, which represents the largest proportion for identified race for this type of death. In contrast, the largest share of civilians who die as a result of less-than-lethal force or physical restraints are black. CEWs and physical restraints are also the most common causes of death for subjects who exhibited symptoms of mental illness or substance abuse.

2.4 Empirical strategy

Figure 2.1 provides the geographical distribution of temperatures across the U.S. We divided the U.S. into ten climate zones and assigned each county to a climate zone based on its mean average daily temperature throughout the year. This map shows a high degree of heterogeneity in temperature across the U.S. The northern parts of the country are more likely to have lower temperatures (i.e., less than 13^oC), whereas the southern areas tend to have warmer weather (15^oC and above). Figure 2.1 shows the average yearly number of civilian deaths when interacting with police per 100,000 capita for each county in the U.S. We did not find any obvious relationship between the temperature zones and civilian death rates.

We ignore the following categories of cause of death: drowned, drug overdose, fell from height, stabbed, undetermined, burned/smoke inhalation, and others.

[9]See Taser use guidelines: cga.ct.gov/2007/rpt/2007-R-0068.htm

To identify the causal effect of temperature on civilian death as a result of police interaction, we used the exogenous daily variation in ambient temperatures at the county level. Daily variation matters if we believe that temperature has a non-linear effect on fatal encounters. Binned regression allows for heterogeneous response at different parts of the temperature distribution. This method has been successfully used to estimate the non-linear effect of temperature on mortality (Deschenes and Greenstone (2011a)), the crime rate (Ranson (2014a)), civil conflict (Hsiang et al. (2011)), and agricultural yield (Schlenker and Roberts (2009)).

Fatal encounters are relatively rare events at the county level and need to be estimated using a count model. The Poisson regression approach accounts for the skewed distribution of fatal encounters. Our data is aggregated by type of death τ for every year y at the month m and county c level. The number of deaths $Y_{\tau cym}$ follows a Poisson distribution with a probability density function and mean $\lambda_{\tau cym}$ such that:

$$f(Y_{\tau cym}|\lambda_{\tau cym}) = \frac{\exp(-\lambda_{\tau cym})}{Y_{\tau cym}!}\lambda_{\tau cym}^{Y_{\tau cym}} \tag{2.1}$$

To ensure that the mean $\lambda_{\tau cym}$ is strictly positive, we estimate the following parametric specification:

$$\begin{aligned}\lambda_{\tau cym} = \exp(\sum_{\tau}^{5}\sum_{b}^{8} Type_{\tau} \cdot Temp_{cym}^{b}\beta_{\tau b} + h(Precip_{cym})\beta_{p} \\ + X'_{cym}\beta_{x} + \gamma_{\tau cy} + \delta_{state} + \varepsilon_{\tau cym})\end{aligned} \tag{2.2}$$

where $Type_{\tau}$ are dummy variables for the distinct types of force (cause of death) which are Firearm, CEW-related, Vehicle, Physical restraint, and Less-than-lethal weapon. Following Deschenes and Greenstone (2011a), we capture the full distribution of monthly fluctuations in weather. The variables $Temp_{cym}^{b}$ denote the number of days where the daily mean temperature is in one of the eight bin variables[10] b in county c within a month m and year y. This method allows us to combine daily variation in temperature with more aggregated variables such as monthly death count. We omit the bin $[12; 17)$ from the regression and used it as the reference category for the temperature effect.[11] The interaction between type of force and temperature, $Type_{\tau} \cdot Temp_{cym}^{b}$, enabled us to recover separate effects of temperature on the number of civilian deaths for each cause of death. Hence, the parameters of interest are given by the whole vector $\beta_{\tau b}$.

The term X_{cym} represents a vector of time-varying explanatory variables that may influence the probability of afatal interaction between police and civilians. A list of these variables is provided in the notes to Tables A.1-A.5 in Appendix A.2. These controls include

[10]The bins are in $5^{o}C$ increments from 2 to $32^{o}C$, including one for temperatures lower and greater than this range.

[11]The choice of the reference bin only affects the interpretation.

the level of precipitation in county c during month m and year y. To ensure that our results are not biased by using a restrictive form of precipitation, we model it using a polynomial function (quadratic), $h(Precip_{cym})$. The term δ_{state} is a state-by-season fixed effect and $\gamma_{\tau cy}$ is a vector for each type of use-of-force county-by-year fixed effect, which control for space and time varying non-observables. As controlling for crime variation was a key component of our analysis, we performed estimation with multiple types of crime being either directly included in the vector X_{cym} or defined as the exposure variable. Setting crimes and arrests within a month as the exposure variable adjusted our estimate for the amount of "opportunity" officers had to use any kind of force. In other words, we account for crimes and arrests as risk sets. Because the risk set itself is impacted by temperature, we have the total number of arrests set as an exposure variable in our preferred specification:

$$\begin{aligned}\lambda_{\tau cym} = \exp(\sum_{\tau}^{5}\sum_{b}^{8} Type_{\tau} \cdot Temp_{cym}^{b}\beta_{\tau b} + h(Precip_{cym})\beta_{p} \\ + X'_{cym}\beta_{x} + \gamma_{\tau cy} + \delta_{state} + \varepsilon_{\tau cym}) \cdot Arrests_{cym}\end{aligned} \tag{2.3}$$

The identifying assumption for our analysis was that, after controlling for county-year-type of death and state-by-season fixed effects, differences in weather and crime between months within a county represent the true effect -$\beta_{\tau b}$- of weather on the number of civilian deaths by type of force used. Moreover, this specification allows for a joint estimation of each type of force which accounts for correlated shocks among type of death. Unless specified, we clustered standard errors at the year-by-type of death.

2.5 Results

Effect of temperature on the level of threat

We began our analysis by examining the effect of temperature on the level of threat for police officers and bystanders. Figure 2.2 presents the coefficients from estimating a simplified version of equation 2.2 on the number of officers assaulted or killed and on the number of violent crimes[12].

These figures indicate that there is a positive and statistically significant impact of temperature on the number of violent crimes and the number of officers assaulted or killed when the temperature is 17^oC or higher. An extra day with a temperature above 17^oC leads to 0.2%-0.5% more assaulted or killed officers, and to 0.28%-0.62% more violent crimes[13]. As pointed out by Ranson (2014a), these coefficients seem small; however, they represent the

[12]We do not observe any kind of type of force used for the violent crimes nor assaults on officers, therefore the equation is simplified and only has binned temperature, spatial and temporal fixed effects.

[13]As a reference, Chalfin and McCrary (2018) find that a one percent increase in the size police causes violent crimes to go down by 0.29 to 0.34%.

effect of a single day of weather per month, and yield important effects in the aggregate. For instance, in a spring month with ten unusually warm days ($[22^oC, 27^oC)$), violent crimes and assaulted or killed officers would be 4.3 to 4.4% higher than the number of violent crimes and assaulted or killed officers relative to the reference bin.

The nonlinear impacts of temperature on violent crimes[14] and officers assaulted or killed indicate that the level of threat for both officers and civilians is higher when the temperature is 17^oC or higher. The level of danger for an officer is statistically similar between extremely hot days and hot days.

This confirms that higher temperatures are associated with more aggressive behavior from civilians (Burke et al. (2015)). Thus, we deduce from Figure 2.2 that it is more dangerous for officers to perform their job when the ambient temperature is higher. Because the level of threat increases above a certain temperature, we expect officers to be more likely to discharge their weapon during warm days. As a result, we hypothesized that, in response to an increase in the level of threat during warm days, there should be an increase in the number of civilian deaths by police shootings (Use of force project (2017)). We test that hypothesis estimating equation 2.2. Figure AA.8 compares the coefficients for fatal shooting with the level of threat. Although results are underpowered, we find that point estimates are consistently increasing in a similar pattern for higher temperature and in particular above 27^oC. These results suggest that, if anything, fatal police shootings increase proportionately to the threat level.

However, the large confidence intervals in figure AA.8 indicate that police officers might not respond to higher levels of threat by using their firearms. They could, for example, use de-escalation techniques, such as restraining methods or less-than-lethal weapons. Although these approaches are not intended to be fatal, there may be some unintended consequences of using these types of force, e.g., the death of the suspect or a bystander. Restraining techniques, for example, can result in death by asphyxia.[15] For these 'less-than-lethal' typse of force, it is important to disentangle the direct impact of temperature from the mechanical effect ascribable to the increase in interaction.

Effect of temperature on civilian deaths

Figure 2.3 presents the coefficients from estimating equation 2.3 which accounts for the different causes of death and the threat level. Allowing for heterogenous effects per causes of death in addition to controlling for the civilian-officers interactions demonstrates that there is a precise null impact of temperature on the number of civilian deaths by firearm. In other words, we found that all coefficients of temperature are close to zero, with small standard errors. The fact that temperature did not impact the number of civilian deaths by shooting might indicate that officers do not suffer directly from warmer temperatures.

[14] Our results are consistent with those found in Ranson (2014a). Our analysis uses the same data source and extends it to include recent years.

[15] For e.g., Eric Garner's 2014 death in New York City, which resulted from a chokehold.

This result is surprising as it contradicts the experimental results from Vrij et al. (1994) who found that police officers are more likely to fire their weapons at assailants during a training simulation conducted in hot temperatures. Our results raise some concerns about the external validity of their findings. On the one hand, it may be very difficult to design a credible experiment that has such high and potentially real-life stakes for both the officers and the civilians involved. On the other hand, in our data we only report fatal interactions; i.e., we do not have data on use-of-force incidents that did not result in the death of a civilian. Thus, we cannot rule out with certainty some behavioral hypothesis.[16]

We now discuss our results for other forms of less-than-lethal force in turn. In terms of the use of less-than-lethal force, we find that there is a significant impact of 'very hot' days on the number of deaths by CEWs and physical restraint, regardless of opportunity (tables A.1 and A.4). The results for less-than-lethal force are too imprecise to draw any conclusions (table A.2).

For CEWs, we found that the monthly number of civilian deaths increased by about 5.3% for every additional very hot day compared to a day in the $12 - 17^oC$ range. The results for very hot days ($> 32^oC$) are highly significant (p<0.001) and unlikely to be a false discovery[17]. The incidence ratio is about 1.0 when the temperature is between 2^oC and 32^oC with small standard errors; in other words, there is a null impact of temperature on the deaths by CEWs. When the temperature is lower than 2^oC, the number of deaths by CEWs is lower compared to the reference temperature; however, the results are not statistically significant. Results in Figure 2.3 also suggest that law enforcement officers are unlikely to substitute CEWs for their firearms during hot days. Indeed, more than the absence of a positive effect for gunshot fatalities, weapon substitution during hot days would imply a reduction in the number of deaths caused by firearms (i.e., a negative coefficient or incidence ratio below 1).

In addition to CEW-related deaths, we found that deaths caused by physical restraints tended to increase significantly during 'extremely warm' days. The incident rate ratio is about 1.15 when the temperatures are at least 32^oC, meaning that the number of civilian deaths by physical restraints increases by 15.1% on 'very hot' days. Due to the low number of cases, the effect of an additional day at 32^oC or above is much less precise for incidents in which physical restraints was used, barely above the 5% threshold (t-statistic = 1.92). The effect at $27 - 32^oC$ is much lower at 4.8% but statistically significant (p-val ¡ 0.05). The physical restraints of an individual during arrest can sometimes leads to fatal hyperthermia and asphyxiation if the victim is subject to excited delirium or controlled in the wrong position. Though the number of asphyxiation and restraint-related deaths only accounts for 1% of all fatal encounters, police officers could nonetheless take simple precautions on hot days to avoid civilian deaths (e.g., limit the use of handcuffs and other restraints when possible).

[16]For instance, it is possible - although unlikely - that police officers do shoot more often and are more aggressive during warmer days but that it somehow does not impact the number of civilian deaths.

[17]With a t-statistic of 4.18, the estimate is also statistically significant at the 1% level when adjusting for multi hypothesis using the very conservative Bonferroni correction ((Rupert Jr et al., 2012)).

For types of less-than-lethal force, excepting CEWs, we found that the effect of temperature on civilian deaths is non-monotonic, meaning that temperature does not increase (or decrease) the number of civilian deaths for these types of force. Moreover, we did not find a statistically significant impact of temperature on civilian deaths for these types of force. The results have large standard errors.

Robustness Testing

We present alternative specifications in Appendix A.2. High temperatures are associated with numerous confounding factors, either societal (*eg.* more crime) or physiological such as discomfort or lack of accuracy ((Qiu and Zhao, 2019)). Our main results are intact when accounting for the various measure of crime level (such as number violent crimes, number of property crimes, and number of assaulted or killed officers).

Appendix A.2 investigates some behavioral mechanism, in particular whether our estimates are mechanically driven by the fact that CEWs usage is strongly correlated with the time of year. Weather might alter officers' choice of force: CEWs are less effective when used on individuals wearing thicker or more layers of clothing. Therefore, it is important to account for the fact that officers are less likely to use their CEWs when temperatures are low. In order to address this, we repeat our analysis on counties where the annual average temperature is at least 19^oC. This approach provides an upper bound on civilians' choice of clothing and officers' force options throughout the year with respect to temperature. We confirmed that temperature has no effect on civilian deaths due to police shootings for warmer regions in the U.S. We show that the number of civilian deaths by CEWs increases by about 6.7% during 'very hot' days only. Our preferred interpretation of this result is that the use of CEWs on hot days triggers unknown physiological factor(s) that increase the risk of death.

Finally in an effort to control for variation in the behavior of civilians, appendix A.2 reports the effect of temperature when officers are facing erratic behavior. We explored the fact that an officer might perceive civilians who exhibit symptoms of mental illness or substance abuse as less predictable or cooperative. We find similar results for both groups for the deaths by firearm and CEWs. Temperature did not impact the number of deaths by firearm, which might indicate that officers exhibit self-control in their decision to use lethal force when facing less predictable civilians or, alternatively, that officers do not perceive such civilians as sufficiently threatening to warrant force. For CEWs, we confirmed that high temperatures increase the odds of a lethal interaction regardless of the mental status of the civilian.

2.6 Discussion

In contrast to the documented racial differences in police shootings (Ross (2015); Fryer (2018a,b); Edwards et al. (2018); Knox et al. (2019)), our research focuses on the following question: how does the level of threat faced by police officers or civilians impact the use

of deadly force by the police? To answer this question, we first show that police officers and civilians face a higher level of threat due to higher temperatures. We exploit this result to test whether the number of police-involved civilian fatalities increases with temperature and, by extension, test the relationship between the use of deadly force and the level of threat faced by officers. We focus on two types of force: (1) lethal force (firearms) and (2) non-lethal force, namely Tasers and forms of physical restraints.

We find suggestive evidence that during hot days the number of civilian deaths from police discharging their firearms increase due to a higher threat level. Our interpretation of this result is that, if anything, officers use their force proportionately to the threat level. However, once accounting for the total number of arrests ("risk set"), we find no additional effect of temperature, contradicting the existing literature on the link between temperature and aggressive behavior. When considering that an officer might perceive as less predictable or cooperative civilians who exhibit symptoms of mental illness or substance abuse, we found a null impact of temperature on death by firearm for both groups. This indicates that officers exercise restraint even in situations in which they (1) suffer high temperature and (2) might reasonably see themselves facing a seemingly increased level of threat from unstable civilians.

These results bear consideration in light of recent high-profile trials for police officers charged with killing civilians in the U.S. (McWhorter (2016)). The officers on trial cited self-defense in response to a perceived threat. Given our identification strategy based on temperature, our results, insofar as they undermine a direct link between the level of threat and the officer's decision to use deadly force, indicate that more research is needed to explain the relationship between threat and police use-of-force. Such research may have important consequences for future legal and policy decisions regarding officer liability with regard to officer-involved civilian shootings.

Secondly, the significant number of unintended deaths from CEWs and physical restraints in hot temperatures may incentivize police departments and policymakers to reevaluate their guidelines around the use of less-than-lethal types of force. Less-than-lethal types of force are not, by definition, intended to be deadly, although these deaths may be related to unknown physiological/biological factors. Guidelines restricting the use of CEWs on vulnerable members of the population—namely, pregnant women, children, and the elderly—already recognize this. However, to the best of our knowledge, there is no study that links CEWs and temperature. Unfortunately, because our study relies on observational data, it is difficult to identify the potential mechanisms behind our results. We believe that more research is needed in this matter from medical and biological science experts in particular, in order to identify the physiological processes that explain the increased probability of dying from CEWs under these conditions.[18]

[18]Potential channels that have been suggested to us from discussions with physicians are that tasers and high temperatures might have an impact on sweat, humidity, hypertension, or electrolytes. However, we do not have the expertise or medical knowledge to evaluate these theories.

Figure 2.1: Yearly average of mean temperature and death rate in U.S. counties

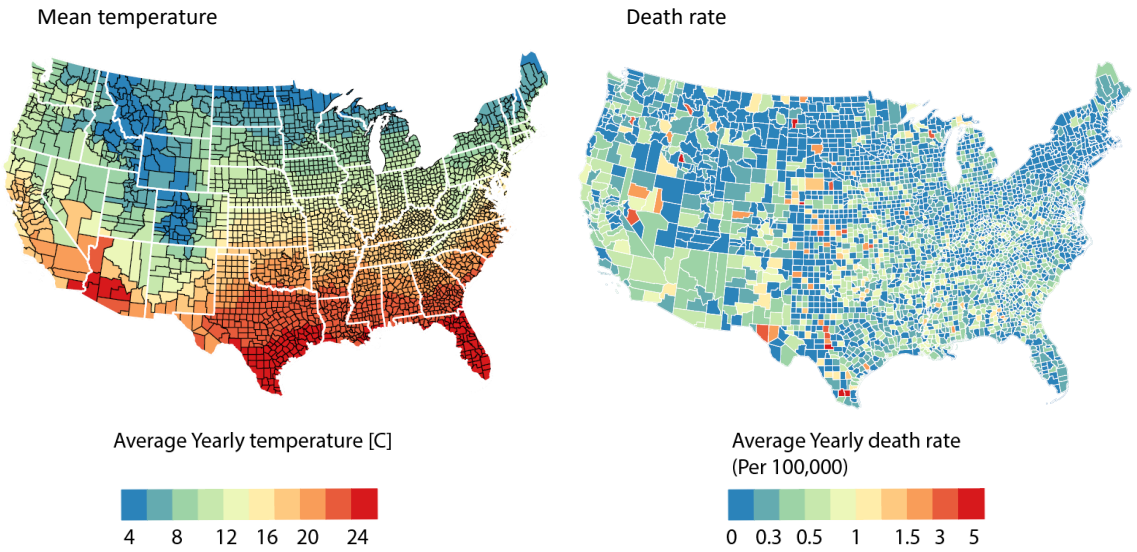

Notes: These maps present the mean temperature and death rate (number of death per 100,000 residents) in U.S. counties from 2000 to 2016.

Conclusion

We provide evidence that the level of risk faced by police officers increases with temperature and that police officers are thus more likely to shoot and kill civilians under hotter weather. However, once we account for the opportunity increase we do not find any additional effect of temperature, indicating that police officers are capable of exercising restraint in difficult conditions. These results also hold when civilians exhibited symptoms of mental illness or substance abuse. In addition, although Tasers are considered a less-than-lethal weapon we demonstrate that the number of deaths by CEWs significantly increases on extremely warm days. Our results suggest that deaths by CEWs on very hot days are driven by compounding physiological factors, and not by an increase in Taser usage from officers. In light of the magnitude of our results, the use of CEWs and physical restraint techniques in extreme heat should be reconsidered.

Figure 2.2: Effect of temperature on the level of threat

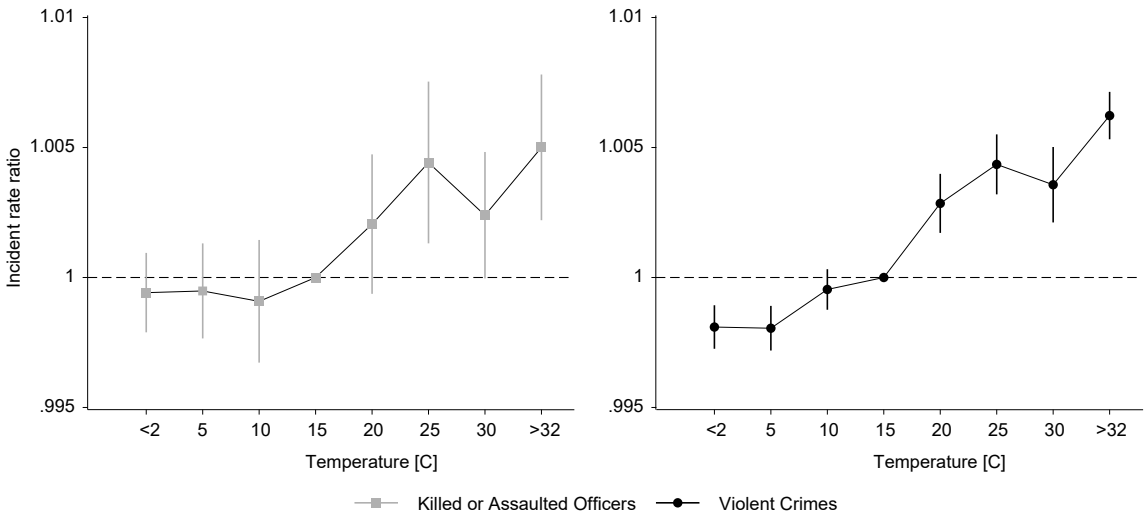

Notes: These figures present the estimated effect of temperature on the number of killed/assaulted officers and violent crimes. The specification controls for precipitation, state-by-season fixed effect, and county-by-year fixed effect. For interpretation, we report the incident rate ratio on the y-axis. Standard errors are clustered at the state level. We report the 95% CI. We report the midpoints for bins with temperature between 2 and 32^oC, with 5 degree increments. Bins -5 and 32+ report temperatures lower and greater than the 2, and 32^oC range.

Figure 2.3: Effect of temperature on the number of civilian deaths by type of force

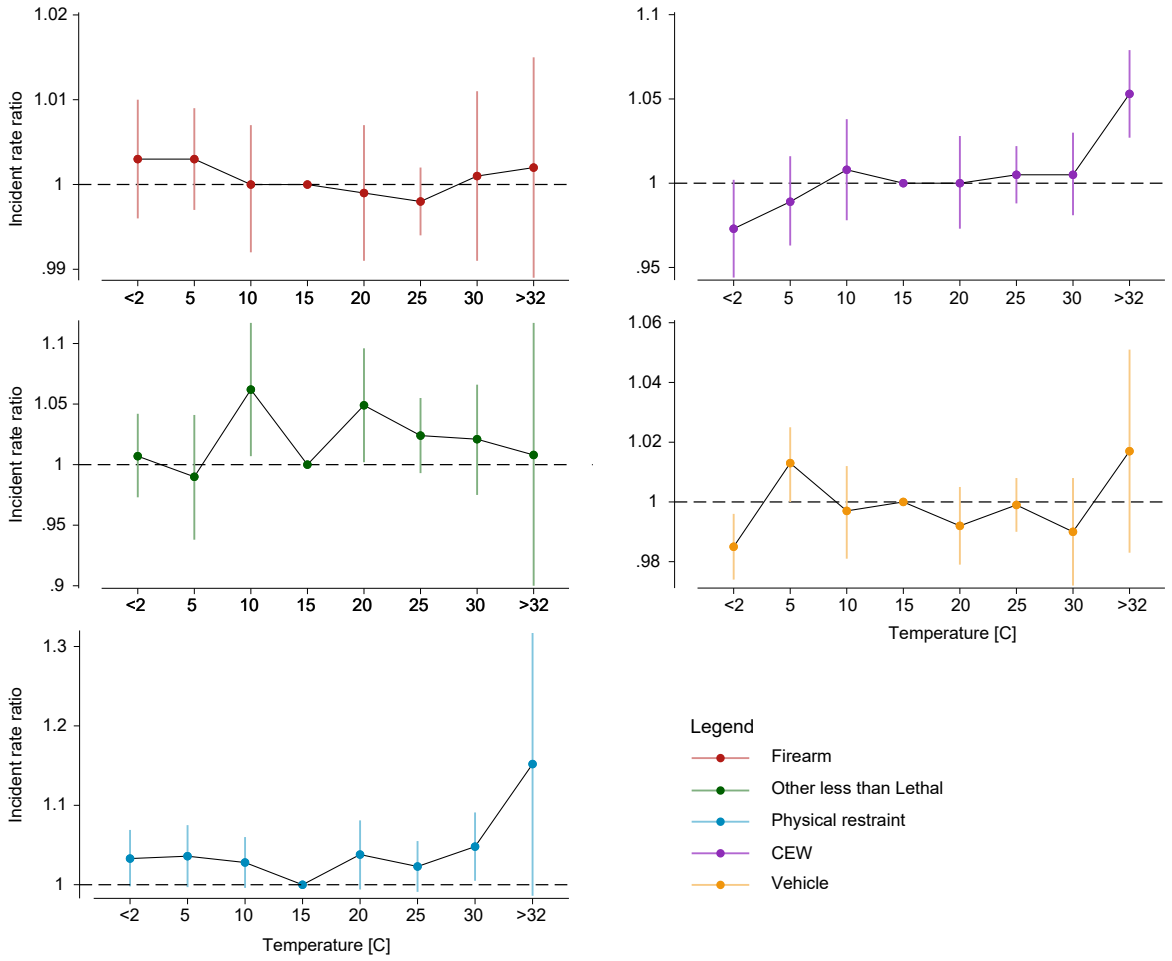

Notes: These figures present the estimated effect of temperature on the number of civilian fatalities by type of deaths. Section 2.2 provides details about the different causes of death. The specification controls for precipitation, state-by-season fixed effect, and county-by-year fixed effect. For interpretation, we report the incident rate ratio on the y-axis. Standard errors are clustered at the year-by-type of death. We report the 95% CI. We report the midpoints for bins with temperature between 2 and 32^oC, with 5 degree increments. Bins -5 and 32+ report temperatures lower and greater than the 2, and 32^oC range.

Figure 2.4: Effect of temperature on the number of civilian deaths by type of force and symptom status

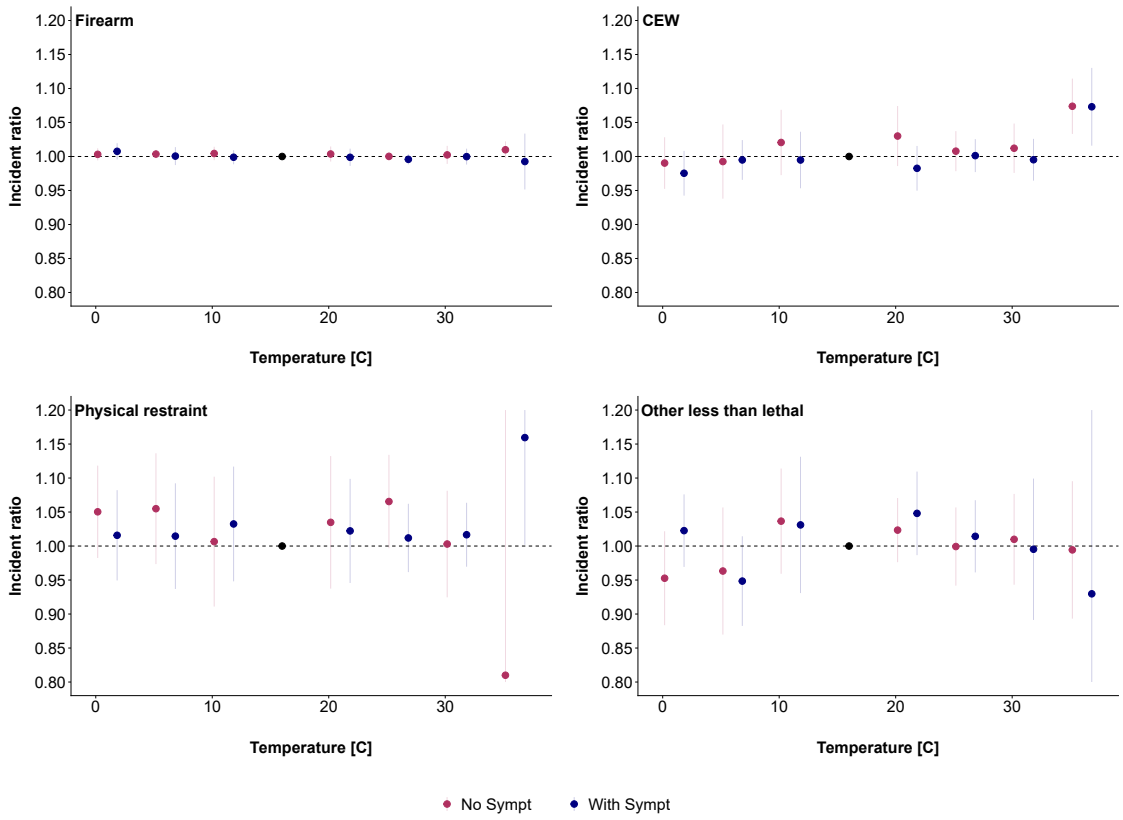

Notes: These figures present the estimated effect of temperature on the number of civilian fatalities by type of deaths and symptom status. Section 2.2 provides details about the different causes of death. The specification controls for precipitation, state-by-season fixed effect, and county-by-year fixed effect. For interpretation, we report the incident rate ratio on the y-axis. Standard errors are clustered at the year-by-type of death. We report the 95% CI. We report the midpoints for bins with temperature between 2 and 32^oC, with 5 degree increments. Bins -5 and 32+ report temperatures lower and greater than the 2, and 32^oC range.

Table 2.1: Summary statistics from 2000 to 2016

	Firearm	CEW	Other less-than-lethal	Physical restraint	Vehicle
Age					
<16	1.8%	0.3%	1.9%	1.4%	9.5%
[16;30)	40.0%	26.7%	24.7%	25.4%	44.7%
[30;65)	53.5%	70.8%	70.3%	68.1%	37.4%
65 and more	4.8%	2.2%	3.2%	5.2%	8.4%
Male	94.7%	97.5%	94.3%	94.4%	74.2%
Race					
Black	20.3%	29.6%	31.7%	28.2%	15.2%
Hispanic	12.8%	10.8%	14.6%	12.2%	10.1%
White	30.8%	24.9%	21.5%	24.4%	24.5%
Unknown	33.7%	33.3%	27.9%	33.8%	48.4%
Symptom of mental illness or substance abuse					
Without symptom	59.1%	31.5%	39.2%	29.1%	85.4%
With symptoms	18.8%	45.0%	32.3%	42.7%	5.6%
Total	15175	771	158	213	4355
Percent	73.4%	3.7%	0.8%	1.0%	21.1%

Notes: This Table presents the summary statistics of the civilians our sample by cause of death from 2000 to 2016. Section 2.2 provides details about the different causes of death.

Chapter 3

A Distribution of Human Attention to Moments in Time

3.1 Introduction[1]

In the course of our individual lives, how much do we think about the past, present, and future? How much conscious attention is focused on events tomorrow compared to the day after, or a date three years ago? How we as individuals distribute our attention across moments in time affects numerous aspects of society and is intrinsic to the human experience, yet we know little about it quantitatively. For example, we do not know answers to simple questions, such as whether individuals naturally think about the past more often than they think about the future, or sophisticated questions, such as the rate at which attention to the future decays. Extensive laboratory-based research has made progress understanding certain aspects of how humans perceive and navigate the passage of time (Eagleman et al., 2005; Bueti et al., 2008; Nobre and Coull, 2010), how neurological systems encode time and orient focus in time(Lewis and Miall, 2003; Olivers and Meeter, 2008; Paton and Buonomano, 2018), or how individuals (McClure et al., 2004; Angeletos et al., 2001) and groups (Gollier, 2013; Millner and Heal, 2018) make choices that tradeoff outcomes that occur at different moments in time. Furthermore, an expansive computational literature has documented how the fraction of collective attention dedicated to specific events (García-Gavilanes et al., 2016; Fanta et al., 2019), individuals (Roediger and DeSoto, 2014) or cultural products (Lorenz-Spreen et al., 2019; Candia et al., 2019) (e.g. songs) decays over time as they are forgotten. Yet, to our knowledge, no study has attempted to measure, either experimentally or computationally, the natural probability distribution of all attention across all moments in time, including the future—perhaps because asking a subject whether they pay attention to something (e.g. "tomorrow") focuses their attention on that thing, thereby biasing their response. Thus, while it is understood that there is a limited "budget" of conscious attention (Kahneman, 1973; Eagleman et al., 2005; Nobre and Coull, 2010; Lorenz-Spreen et al., 2019)

[1]The material in this chapter was co-authored with Léopold Biardeau and Solomon Hsiang.

such that an individual cannot pay attention to all moments in time with equal focus, how this budget is allocated remains an open question.

Here we propose a quantitative model for how a population of individuals distribute their attention across moments in time and we provide the first measurement of this distribution under natural conditions. As individuals direct the focus of their attention towards different objects, ideas, tasks, and events—which we refer to collectively as *targets*— some will be unrelated to any particular time, such as a poem or abstract life goal, while others are associated with a specific moment in time, such as a birthday or project deadline. When individuals think about these time-associated targets, some fraction of these thoughts will lead them to execute an internet search that is related to the target, and among these internet searches, a subset will utilize a query whose text contains information that identifies a time period associated with the target (e.g. '2012' or 'Memorial Day'). By studying patterns in search volume for such time-associated terms, we are able to recover the fraction of attention a population directs towards targets that occur at fixed moments in time.

In this study, we use publicly available data compiled from *Google Search* (Google) to study how populations direct their attention at different targets. This approach allows us to study a large portion of all humankind, since roughly 1.8B individuals execute a *Google* search each day (totaling 63,000 searches per second (ArdorSeo)), and it allows us to observe natural patterns of attention passively without interfering in normal behavior. *Google Search* data clearly do not capture most thoughts or ideas experienced by individuals, as only a small fraction trigger an internet query. However, because a large number of individuals utilize *Google Search*, it has been shown to be a highly sensitive measure of changes in real world attention, despite the low probability that any individual executes a query at any given moment. For example, it has been demonstrated that *Google Search* volume corresponds with aggregated income levels (Preis et al., 2012; Noguchi et al., 2014), unemployment (Ettredge et al., 2005) and individual wellbeing (Algan et al., 2019), travel planning and consumer confidence (Choi and Varian, 2012), attention to movies, video games, and music (Goel et al., 2010), housing sales and prices (Wu and Brynjolfsson, 2015), the attention of financial investors (Choi and Varian, 2012; Drake et al., 2012) and stock price movements (Preis et al., 2013), as well as thoughts and views that individuals might not voluntarily disclose, such as racial animus (Stephens-Davidowitz, 2014), suicide (Gunn III and Lester, 2013) and sexual preferences (Stephens-Davidowitz and Pabon, 2017). However, to our knowledge, ours is the first study to use these data to measure how attention is allocated across *all* targets based on their position in time.

Importantly, the application of *Google Search* data in this analysis does not equate search queries with attention, nor does it require that queries represent a random sample of all thoughts or forms of attention. Rather, we study query volumes because they are driven by underlying patterns of attention through some unobservable process that might not involve uniform sampling of attention. Our empirical analysis only requires that within each category of target, such as thoughts about a movie or questions about gardening, changes in attention over time are reflected in changes in queries associated with that specific category over time (see *methods*). This assumption is consistent with prior analyses of *Google Search* data

(Ettredge et al., 2005; Algan et al., 2019; Choi and Varian, 2012; Goel et al., 2010; Wu and Brynjolfsson, 2015; Drake et al., 2012; Preis et al., 2013; Stephens-Davidowitz, 2014; Gunn III and Lester, 2013; Stephens-Davidowitz and Pabon, 2017) and can be easily verified for categories of targets not previously analyzed. For example, total queries for "sunscreen" cycle annually, peaking in the Northern Hemisphere summer, queries for "Donald Trump" increased abruptly in 2015 when global attention focused on the then presidential candidate, and queries for "cat" have remained essentially constant for over a decade (ED Figure A.10).

Model: The Kernel of Attention to Time

Our model is described using two parallel timelines: targets are fixed at locations in *absolute time* (θ) while a population passes along in *experienced time*, where they are observed at each moment (t) only transiently. As the population proceeds through experienced time, their attention to targets in absolute time evolves as their focus is constantly re-centered on the moment in time that they inhabit (Figure 3.1**a**). Attention is then allocated across all possible targets depending on their position in absolute time relative to the population's current position in experienced time, a distance we call *relative time* ($\theta - t$) (Figure 3.1**b**). The situation is analogous to how passengers on a moving train might allocate their attention to stationary landmarks that they pass. Similar to how the gaze of passengers move across landmarks depending on how near or far they are in physical distance from the train at a given moment, the attention of subject populations moves across targets depending on their distance in time.

Under natural conditions, attention in a subject population is likely to flicker rapidly across numerous targets located at different times in a seemingly random pattern. For example, it is theoretically possible that the process governing human attention is uniformly random, in the sense that this probability is the same for all targets at each moment — although we do not think this example is realistic. Our hypothesis is that the distribution of this flickering attention is governed by a stable and measurable function defined over moments in relative time, which we call the "kernel of attention to time" (KAT) and denote $\kappa(.)$ (Figure 3.1**c**). Specifically, we hypothesize that instances of human attention flicker across time-associated targets as if they are being drawn randomly, such that their position in relative time follows the marginal probability density function $\kappa(\theta - t)$ (see formulation in *methods* and Eq. 3.1). As a population moves forward through experienced time, the distribution of its attention will shift simultaneously such that the KAT is always anchored to the moment inhabited by the population — where relative time equals zero (i.e. $\theta = t$). In the moving train analogy, our model would suggest that while passenger attention across stationary landmarks may appear random, they actually follow a fixed probability distribution that is always anchored around the current position of the train.

Testing and measuring the KAT using *Google Search*

If the above KAT hypothesis is correct, it generates predictions that can be tested using *Google Search* data and which allow us to estimate the shape of the KAT.

As a population travels through experienced time and "passes" a target that is fixed at a moment in absolute time (Figure 3.1**d**), attention allocated to the target will trace out the KAT as the population approaches and then passes the moment of the target (Figure 3.1**e**). If these changes in attention are reflected in a changing number of search queries, then a time-series of search query volume will be a reversed mirror-image of the KAT, when plotted in experienced time (Figure 3.1**f**). Note that experienced time is the timeline marking when each search query is executed. In cases where a target occupies multiple moments in time (e.g. an event that lasts for a year, Figure 3.1**g**) then aggregate search volume for that target will be a summation of KAT functions that are all flipped and shifted, each corresponding to a different moment that the target occupies (Figure 3.1**h**). Thus, if the KAT hypothesis is correct, total search volume for any target (in experienced time) should be the convolution of the KAT with a time-series of Dirac or indicator-functions that span the absolute time occupied by the target (see Figure 3.1**i**, also ED Figure A.11 and *methods*), scaled by a constant.

Using this result, we should be able to recover the shape of the KAT by deconvolving the envelope of *Google Search* volume time-series (Wiener, 1950; Silvia, 1987; Delaigle et al., 2008), so long as we know when in absolute time the targets occur (see *methods*). Knowing the date of the target is important because we must compute the distance in relative time between when the search query is executed (t) and when the target occurs (θ). We cannot assign the target of most search queries to a date, but we can do this for the subset of queries that contain date-specific strings. Specifically, we estimate the KAT for queries that contain a year in numerical form (e.g. "2016"), a month (e.g. "August"), or a holiday (e.g. "Cinco de Mayo") during the period 2009-2018 (see *methods*). Prior work has studied basic patterns in aggregate search queries for years (Preis et al., 2012; Noguchi et al., 2014), but none has explained nor demonstrated how individual attention and query behavior might generate these data.

Deconvolution of total search volume into KAT and time series components is essential to observing the common underlying structure of attention across different target classes and understanding the origin of these aggregate measures of attention. The envelope of search volume for year, months, and holiday targets show predictable periodicity within each of these classes (Figure 3.2, blue lines), but the connection between the raw envelopes for these classes is neither visually nor quantitatively obvious. For example, the rise in 2014 total search volume for "2015" is nearly exponential, but the run-up in 2014 search volume for "diwali" preceding to the holiday is not. The existence of a generalizable KAT would connect these two time-series, which our model predicts diverge due to differing underlying distributions of targets.

Note also that prior studies (García-Gavilanes et al., 2016; Fanta et al., 2019; Roediger and DeSoto, 2014; Lorenz-Spreen et al., 2019; Candia et al., 2019) of "collective memory"

for specific historical targets (e.g. a specific publication (Candia et al., 2019)) do not recover a probability distribution of attention nor the KAT. Rather, they can be understood as measuring the decay of attention to a single target, analogous to a specific circle or square in Figure 3.1**b**. While this decay might be governed by the structure of the KAT, the KAT cannot be mathematically inferred by this decay since the KAT describes the superposition of attention to all types of targets (see Eq 3.4 in *methods*) and has a future-oriented component that cannot be observed by studying collective memory.

3.2 Results

The temporal structure of *Google Search* volume for holidays, months, and years is well-predicted by the KAT model for human attention (Figure 3.2). After adjusting for secular trends in search volume, which generally reflect expansion of the internet, the KAT model exhibits high predictive accuracy for holidays ($R^2 = 0.81$), months ($R^2 = 0.84$), and years ($R^2 = 0.93$) in our sample.

Figure 3.3**a** displays the estimated average distribution of attention to time (the KAT, using year targets) for all individuals using *Google Search*, based on trillions of queries. We estimate that roughly 25% ($\text{CI}_{95\%} = [24.5\%, 26.1\%]$) of attention is directed at the present period — defined as the shortest contemporaneous interval of time (e.g. day) that we are able to observe — dramatically dominating any individual moment in the future or past. For example, we estimate that on average, attention is 5.7 ($\text{CI}_{95\%} = [5.6, 5.8]$), 18.7 ($\text{CI}_{95\%} = [18.3, 19.1]$), and 28.5 ($\text{CI}_{95\%} = [27.5, 29.4]$) times more likely to be directed at the current day compared to 10 days, 50 days, and 100 days in the future, respectively; and 6.3 ($\text{CI}_{95\%} = [6.1, 6.4]$), 16.0 ($\text{CI}_{95\%} = [15.3, 16.8]$) and 22.8 ($\text{CI}_{95\%} = [21.6, 23.9]$) times more likely than days of corresponding distance in the past (ED Figure A.12). The structure of the KAT we estimate is both statistically precise and highly robust to a variety of different modeling approaches, including statistical models that flexibly account for constant differences between populations, such as culture, and gradual changes in populations over time, such as demographic trends and changing access to technology (ED Figure A.13; see *methods*).

As targets move further in relative time, the rate ($\%\Delta/day$) at which attention declines is not constant. This is true for both the future and the past. Attention to dates in the immediate future or past 7 days decline roughly 25% and 30% per day, respectively (red segments, Figure 3.3**a**). In the more distant future and past, rates of decline fall to 0.6% per day for both the future and past 56 to 150 days (yellow segments). In between is a transition region (1-8 weeks) where rates of change in attention evolve quickly (green and blue segments). For reference, if the KAT were an exponential function, analogous to having a constant discount factor for attention, this rate of decline would be constant.

Although the present period is the single moment receiving the most attention, the total attention allocated to all future or past moments is larger than attention allocated to the present. We compute the total attention allocated to the "future" by summing attention

across all periods in the future, and compute similar sums for all periods in the past. In the global average KAT (Figure 3.3**a**), we estimate that roughly 39% ($CI_{95\%} = [38\%, 39.5\%]$) of attention is allocated to the past, while 36% ($CI_{95\%} = [35.2\%, 36.6\%]$) is allocated to the future (black marker in Figure 3.3**b**, see also ED Figure A.12**a**). It is important to note that these values reflect attention to specific targets associated with identifiable moments in time, but they do not capture attention allocated to abstract notions of "the distant Future" or "the distant Past" that are not associated with specific dates. These abstract concepts are not associated with specific or measurable intervals of time and so are not represented in the KAT model nor in our measurements based on search queries.

The structure of the KAT — a sharp peak in the present, slowing rates of decline with growing temporal distance, and roughly equal mass in the past and future — is broadly reproducible using search data on annual targets from any region of the world, regardless of language, region, or culture. Analyzing data from 181 countries separately, we find that country-specific KATs are tightly clustered around the global average (Figure 3.3**b**). Figure 3.3**c** illustrates separate KAT estimates for a diverse selection of twenty different countries. While differences exist across populations, which we explore below, the overall similarity in how populations everywhere allocate attention in relative time is notable and might suggest a common biological origin.

We compare the average KAT we estimate using language-independent numerical annual targets (e.g. "2012") to those estimated using month targets (defined in the local language) or national holidays (see ED Table A.6). We find that the structure of the KAT is similar regardless of the class of target used (Figure 3.3**d-f**, see ED Figure A.14 for country-specific estimates). However, holiday targets produce a KAT that is relatively more future-oriented, perhaps a result of anticipation specific to holidays, and we have difficulty detecting attention to targets more than 100 days before or after the present using either holiday or monthly targets, due in part to different statistical properties of these data. Nonetheless, all three approaches to estimating the KAT indicate a roughly symmetric structure with consistent mathematical form (described below) that peaks sharply on the present moment.

In seeking an analytical mathematical form that describes the structure of the KAT, we find that it is generally well approximated by two functions, one for the future and one for the past, where each is from the family of rational functions (i.e. the ratio of two polynomials, green lines in Figure 3.3**d-f**; also see *methods* for analytical forms) which intersect at $\theta - t = 0$, i.e. the contemporaneous period. Exponential functions, which have been shown to fit the first year of collective memory for specific targets generally (Candia et al., 2019), cannot fit the structure of the KAT (red lines); while high-order polynomials can approximate the data but have the undesirable feature of being non-monotonic (yellow lines). Following research in behavioral economics and psychology (Laibson, 1997; Findley, 2015),we also attempt to fit "hyperbolic discounting"-like functions ($\frac{a}{1+b(\theta-t)}$) that are within the family of rational functions but restrict the numerator to be constant (note that these are not "hyperbolic functions"). This restriction on the rational function produces visually comparable, albeit poorer, fit to the KAT (blue lines) but results in substantially biased

predictions for total search volume of monthly and yearly targets (ED Figure A.15). Only the unrestricted rational function approximation reliably generalizes across annual, monthly and daily holiday targets.

Although there is consistency in the overall structure of the KAT around the world, we also explore how the estimated KAT differs across countries. First, we estimate a KAT only using search data originating from a single country. We then compute the fraction of attention allocated to the past, present, and future based on that country-specific KAT. To summarize these results, we map the ratio of attention to the future relative to the past (Fig 3.4**a**) and the fraction of attention allocated to the present (Fig 3.4**b**). We observe systematic regional patterns[2] (ED Figure A.16). For example, much of tropical Latin America, coastal Africa, and Asia allocates greater attention to the past than the future, while most of North America, Western Europe, North-Central Africa, and Southern South America allocate more attention to the future than past, as do Australia, Malaysia, Taiwan, and Japan (regional outliers). Geographic patterns of attention to the present are less clear. Future work should seek to understand the origin or consequences of these patterns in attention.

Last, we examine whether the distribution of attention to moments in time has been recently changing in a coherent way around the world. To do this, we estimate a KAT for each year separately using each class of target (years, months and holidays) and plot these trends in Figure 3.4**c**. Regardless of which class of target is used, we find that between 2009 and 2018, the fraction of attention to the present is rising roughly 0.012 per year globally (± 0.006, $p < 0.01$), at the expense of the fraction of attention to the past, which is declining roughly by the same amount (0.012, $p < 0.05$). The fraction of attention allocated to the future appears to have been somewhat stable over this decade. It is possible this trends results from accelerating production of new content that crowds-out attention to historical content, although the magnitude of trends in prior analyses of that phenomenon (Candia et al., 2019; Lorenz-Spreen et al., 2019) do not match our finding. The causes for and consequences of this trend in the distribution of attention is a potentially important question for future investigation.

3.3 Discussion

Here we have proposed that the distribution of human attention to moments in time maintains a coherent probabilistic structure and we recover estimates for this structure, the KAT, for a population of billions of *Google Search* users. We discover that the shape of the KAT is highly consistent around the world, regardless of the measurement procedure, suggesting that it describes a general pattern of thought fundamental to human minds everywhere. However, we also document some modest differences around the world and over time in attentiveness to the past, present, and future — suggesting that social, cultural, environmental, economic,

[2]We included all the countries for which the data were made available on Google's API. However, one should remain aware of the Internet restrictions put in place in several countries (e.g. China, Iran or Syria) which may impact the accuracy of their estimated KAT.

or other factors can influence how attention is distributed across moments. We believe that disentangling the influence of these factors is a potentially important area for future research.

These findings may inform our understanding of questions in social and cognitive sciences. For example, economics research has studied how individuals make tradeoffs over time (e.g. foregoing present consumption to increase future consumption), captured in models using "discounting" and "time preference" parameters (Gollier, 2013; Millner and Heal, 2018). Attention to time does not capture economic tradeoffs directly, but the structure of the future-oriented portion of the KAT could enable better understanding of the origin of time-based preferences. In another example, political science research has examined how constituents retrospectively evaluate leaders for their past performance (Nordhaus, 1975; Findley, 2015), which may depend on the past-oriented portion of the KAT. In a third example, psychology research has considered many ways that thoughts about the past or future influence individuals, such as altering affect (Zimbardo and Boyd, 1999; Boyd and Zimbardo, 2006) or contributing to anxiety (Wohlford, 1966). The methods we present here, perhaps combined with additional data, might provide insights into these relationships or inform the development of clinical tools that support mental health.

Our use of *Google Search* data allows us to measure the KAT *in situ* for the first time, however two remaining challenges might be addressed in future research using other data sources. First, our empirical approach cannot capture attention to abstract notions of "the distant Future" or "the distant Past" if attention directed towards those topics is not associated with specific moments in time that can be queried. For example, the diffuse concept of future generations living centuries from the present will not be well-captured by our approach. Nevertheless, the structure of the KAT we recover might be consistent with how humans think about these distant periods, since finite attention allocated to multi-century periods would likely indicate that only a very small fraction of attention can be spread across individual moments spanning those long intervals.

Second, the sample of aggregated and anonymized *Google Search* queries we use is not fully representative of the global population. For example, access to the internet, and by extension to Google, varies from country to country (Internet World StatsInternet World Stats, 2019) with a higher penetration rate for higher income countries. Google is also used less frequently by older populations (Weber and Castillo, 2010; Weber and Jaimes, 2011), and it is not widely used in countries where access is restricted by governments. Thus, it is possible that some of the differences we observe in aggregated data across locations or over time is due to differences in the composition and behavior of Google users. Nonetheless, the overall consistency of the KAT we recover around the world might suggest that unobserved populations could behave relatively similarly to those we do observe.

It is also possible that *Google* searches may be executed by individuals based on some expectation about the availability of online content (Jara-Figueroa et al., 2019), with individuals searching less for future targets if they believe associated content does not yet exist. While this behavior may influence our results, it cannot account for our main findings. The increase in attention to the present, relative to attention to the future, is too rapid to be explained by changing content: a similarly abrupt increase in online content between

sequential moments would require that total searchable online content grows continuously at a rate of 25% per week, orders of magnitude faster than actual content growth. Additionally, a content-driven bias would suggest past-oriented queries should strongly dominate future-oriented queries, which we do not observe.

We believe that understanding how human attention is distributed across time may inform the design of policies, technologies, interventions, or institutions. For example, the timing of public health reminders (e.g. "use condoms to avoid STD transmission") or natural hazard warnings (e.g. "store clean water") could be designed to maximize their benefits by accounting for populations' overall attentiveness to moments in time, especially when inattention to such messages is costly. In other examples, policies designed to encourage saving money depend on attentiveness to the future while the design of census questionnaires that require accurate recall might depend on attentiveness to the past.

We close noting that in discussions surrounding long-term planning, such as managing the global environment, it is sometimes argued that challenges arise because populations do not think enough about the future (King, 2019). Our results suggest that while some factors may cause populations to think slightly more or less about the future, the overall general pattern of thought in humans everywhere is to allocate essentially no attention to moments in time further than two-hundred days from the present. Thus, these results might suggest that efforts to directly focus large populations on the distant future are unlikely to succeed. As an alternative, societies may need to design institutions that systematically evaluate and address long-term challenges, rather than relying on the attentiveness of individual minds.

3.4 Methods

A Probabilistic Model for Attention to Time

Timelines We assume that targets of attention are associated with fixed positions in time and human populations move through time. Although both the targets and human populations exist in the same single timeline in the real world, we find it helpful to conceptualize two parallel timelines that are indexed separately, but which have a direct correspondence between one another. Targets of attention exist in a static timeline that we refer to as "absolute time" and which is indexed by θ. Human populations travel along a timeline that we call "experienced time" and which is indexed by t. Each position in absolute time θ_0 has one and only one corresponding position in experienced time t_0. Because the two timelines are parallel, we define the distance between a moment t and θ_0 to be the same as the distance between t and t_0 (or θ and θ_0).

Embedding targets of attention in absolute time We discretize time into "moments" that are finite and equal in length. In principle, time can be sub-divided into arbitrarily small moments; although, in our data analysis, we consider discrete dates since that is the resolution of data available to us.

Let ϕ_x be a potential target of attention, and let $\Phi = \{\phi_1, \phi_2, \cdots\}$ be the countably infinite set of all possible such targets. Let the subset $\Phi^{time} \subseteq \Phi$ be those targets associated with specific moments in absolute time, such that for each target in $\phi_x \in \Phi^{time}$ there exists a fixed interval of time Θ_x associated with the target ϕ_x. Specifically, ϕ_x is associated with Θ_x if ϕ_x is associated with some (or all) moments contained in the interval Θ_x. Let the interval Θ_x be composed of a set of M consecutive moments in time $\Theta_x = \{\theta_{x_1}...\theta_{x_M}\}$ and we consider ϕ_x associated with the interval Θ_x if it is associated with some or all moments θ_{x_m} in the interval. Let the notation $\phi_x \mapsto \theta_{x_m}$ mean that the target ϕ_x are associated with moment θ_{x_m} contained in the interval of time Θ_x (define $\phi_x \mapsto \emptyset$ if $\phi_x \notin \Phi^{time}$).

For an example, a birthday party on July 15, 2017 during the time between 10:00 and 12:00 could be a target of attention (ϕ_x). In this case, we might consider the two hours of the party the moment in time directly associated with the target (θ_{x_m}). This party could then also be associated with the date "July 15, 2017", which describes an interval of time containing that moment (Θ_x). In our analysis of internet query data, we sometimes study identifiable periods described by just the month, in which case Θ_x="July", or by just a year, in which case Θ_x="2017". This framework and notation is important because in search data, we often cannot identify the precise moment to which individuals are attentive when they execute a query — e.g. they do not search for "10:00-12:00 on July 15, 2017" — but we can identify information about a larger interval of time containing that moment (e.g. "July"). Nonetheless, our aim is to understand attention to narrower moments using the search query data identifying only longer time intervals, which forces us to consider how attention is allocated to moments which are themselves nested within intervals. Below, we omit the x (and x_m) subscripts of ϕ_x, Θ_x and θ_{x_m} for notational parsimony, although there continues to be mappings from a target ϕ to moments θ contained within time intervals Θ.

The Kernel of Attention to Time At each instant in time t experienced by humans, a randomly selected individual j in a subject population focuses their attention A_j on a target. Let the notation $A_j(t) = \phi$ denote the event that at moment t, individual j directs their attention at target ϕ. Let $\Pr[A_j(t) = \phi]$ denote the probability of this event occurring. We propose that the probability distribution of attention across targets is such that the total attention allocated to all targets associated with the moment θ depends on the temporal distance between that moment and the time t inhabited by the subject population. We call the temporal distance $\theta - t$ "relative time" (see Figure 3.1**a-c**). A positive value for relative time means an event is in the future of the subject population, a negative value for relative time means the event is in the population's past, and a zero value means that the event is in the population's present.

Formally, we hypothesize that attention is allocated across targets such that

$$C \cdot \left(\sum_{\phi \in \Phi^{time}} \Pr[A_j(t) = \phi \mid \phi \mapsto \theta] \right) = \kappa(\theta - t) \tag{3.1}$$

holds (approximately) at each t. Here, $\kappa(.)$ is a probability density function that integrates to one, which we call the *kernel of attention to time* (KAT). Eq. 3.1 states that at time t, the sum of probabilities that individual j is attentive to any target associated with the specific moment θ is described by the kernel function $\kappa(.)$, where the argument is the difference between the time of the targets (θ) and the time at which j might be attentive to the targets (t), i.e. relative time ($\theta - t$). C is a scaling factor:

$$C = \frac{1}{\sum_{t=-\infty}^{\infty} \sum_{\phi \in \Phi^{time}} \Pr[A_j(t) = \phi \mid \phi \mapsto \theta]} \tag{3.2}$$

that ensures $\kappa(.)$ integrates to one. The denominator of C is the total expected fraction of attention that is allocated to all targets associated with θ across all periods, a value that is much smaller than one since most potential targets will not be associated with the specific moment θ, possibly because they are not associated with any moment at all (e.g. a poem). For notational parsimony, here we assume this scaling is constant, but it could be allowed to change as a function of θ if, for example, there are more targets associated with holidays than for other dates.

Note that the expression inside the brackets on the left-hand-side of Eq. 3.1 is the a sum of probabilities that attention is directed at targets associated with the time θ. The expression sums over all possible time-associated targets (it could alternatively be written as a sum over all possible targets Φ). The elements of this sum may change as the time of the subject population (t) moves forward and the expression does not require that the probability of attention to any single target evolves following the shape of $\kappa(.)$. Instead, Eq. 3.1 allows for many possibilities in how the likelihood of attention to individual targets associated with θ evolve, however it constrains their sum to be proportional to the value of the KAT at the corresponding relative time $\theta - t$.

Predictions for Internet Search Volume

Eq. 3.1 generates predictions about the time-series structure of time-associated internet search volume. In each moment, if an individual's attention is allocated to a target, then there is some probability that it will lead to them executing an internet search query related to the target. Of those searches, a fraction will be associated with a specific moment in time. From this fraction, an even smaller fraction will contain a string that identifies a time interval containing the associated moment in time. Thus, we can compute the probability that an individual j allocates attention at time t toward some time-associated target ϕ and we observe an internet search query that identifies the time interval Θ associated with that

target:

$$\Pr[\text{observe } j \text{ search for } \phi \text{ at time } t \text{ and identify } \Theta] =$$
$$\sum_{\theta \in \Theta} \Bigg(\Pr[A_j(t) = \phi | \phi \mapsto \theta, \phi \in \Phi^{time}] \cdot \underbrace{\Pr[\phi \mapsto \theta \mid \phi \in \Phi^{time}]}_{\gamma(\phi,\theta)} \cdot \Pr[\phi \in \Phi^{time}]$$
$$\underbrace{\cdot \Pr[j \text{ searches for } \phi \text{ at time } t \mid A_j(t) = \phi, ...] \cdot \Pr[j'\text{s search query identifies } \Theta \mid j \text{ searches for } \phi, ...])}_{\omega(\phi)} \Bigg). \tag{3.3}$$

Eq. 3.3 decomposes this probability into five (conditional) probabilities for each moment in time θ, which are multiplied and summed over all moments in the identifiable interval of time Θ. The first factor is the probability that j allocates attention to the specific target ϕ at time t, if the target is associated with the specified moment θ. The second factor is the probability that the target of attention is associated with the moment θ, which we denote $\gamma(\phi, \theta)$ for clarity below. The third factor is the probability that the target ϕ is associated with any moments in time – for a known target this value is either one if the target ϕ is in the set Φ^{time} and zero otherwise. The fourth factor is the probability that j executes an internet search for the target ϕ at time t, given that they were attentive to it at that moment. The fifth factor is the probability that the executed search query contains information that identifies the time interval Θ that contains the moment θ associated with the target. The product of the last two factors is assumed to be a constant for each potential target (or category of target) during the period of observation, written as $\omega(\phi)$ for notational simplicity. In principle, gradual changes in internet access or slowly changing patterns of behavior might alter $\omega(\phi)$ because it changes the likelihood that attention to a topic generates an internet query, but we find in practice that such trends do not affect our results (see ED Figure A.13). Note that different targets ϕ could have different values for $\omega(\phi)$, indicating that different types of targets could have different likelihoods of generating a search query (e.g. attention to gardening and attention to a medical appointment may have a different likelihoods of producing a time-associated search query).

We compute the expected total volume of searches executed at time t that explicitly identify time period Θ, denoted $S_\Theta(t)$, by summing across N individuals in population P and summing over all possible targets of attention in Φ:

$$S_\Theta(t) = \sum_{j \in P} \sum_{\phi \in \Phi} \Pr[\text{observe } j \text{ search for } \phi \text{ at time } t \text{ and identify } \Theta]$$
$$= \underbrace{\frac{N}{C}}_{\alpha} \sum_{\theta \in \Theta} \underbrace{\left(\sum_{\phi \in \Phi^{time}} \gamma(\phi, \theta) \omega(\phi) \right)}_{\bar{\delta}(\theta)} \kappa(\theta - t) \tag{3.4}$$

where the second equality follows by substitution of Equations 3.1, 3.2 and 3.3 and rearranging terms. Eq. 3.4 illustrates one way that the function $\kappa(.)$ behaves like a kernel function, hence KAT, since it weights relative search volume for all possible targets based on the distance in relative time. The scaling factor $(\frac{N}{C})$ captures the size of the population (N) and the likelihood that individuals in that population allocate attention to any time-associated targets $(\frac{1}{C})$. For parsimony, we rewrite this scaling factor as α. Eq. 3.4 can be further simplified by observing that the summation across all targets (ϕ) of the joint probability distribution $\gamma(\phi, \theta)$ over targets and moments in time is the marginal distribution over moments in time when targets occur. Multiplication of $\gamma(\phi, \theta)$ by $\omega(\phi)$ in the summation of Eq. 3.4 results in a re-weighting of this marginal distribution based on $\omega(\phi)$ (i.e. the likelihood that attention to each type of target results in an internet search query identifying a period of time). This re-weighted marginal is written $\bar{\delta}(\theta)$ and describes the fraction of all possible time-associated targets that are associated with each specific moment in absolute time and generate a search query identifying that moment in time.

We note that the KAT and the notion of "collective memory" studied elsewhere (García-Gavilanes et al., 2016; Fanta et al., 2019; Roediger and DeSoto, 2014; Lorenz-Spreen et al., 2019; Candia et al., 2019), are fundamentally different concepts. This can be seen in Eq. 3.4, which shows that the KAT describes the distribution of attention across all $\theta \in (-\infty, \infty)$ at each time t, where mass in the distribution results from summing attention across all potential targets $\phi \in \Phi$. Studies of collective memory document how references to a single topic ϕ_x (e.g. a research article or song) decays for moments in time following the time of the target, i.e. $t \geq \theta_x$ (García-Gavilanes et al., 2016; Fanta et al., 2019; Roediger and DeSoto, 2014; Lorenz-Spreen et al., 2019; Candia et al., 2019). Thus, $\kappa(\theta - t)$ cannot be inferred nor reconstructed from studies of collective memory.

The final expression for predicted search volume at time t identifying time period Θ thus has the simplified form

$$S_{\Theta}(t) = \alpha \sum_{\theta \in \Theta} \bar{\delta}(\theta) \kappa(\theta - t). \tag{3.5}$$

Eq. 3.5 describes how total query volume that identifies a period of time (Θ) across all targets should evolve as a population moves through time (t), approaching and then passing the identified moments (θ) that occupy Θ. This query volume is the product of a scaling factor (α) and a time series $(\sum_{\theta \in \Theta} \bar{\delta}(\theta) \kappa(\theta - t))$. In estimation, this scaling factor is not directly observed because the data are normalized by the data provider, but regardless, this normalization is immaterial because α is a nuisance parameter that can be accounted for by allowing for flexible multiplicative trends in the data. The structure of the time series component is the subject of our focus. Eq. 3.5 says that the time series of search volume is a super-position of reflected (with respect to t), shifted (by θ), and rescaled (by α) KAT functions, one for each moment θ in the identified period of time Θ. One KAT function is reflected and shifted so that it is centered at each moment θ (Figure 3.1**f**). These reflected and shifted KAT functions are each multiplied by $\bar{\delta}(\theta)$, the weighted fraction of time-associated targets corresponding to each moment θ, and summed at each value of t, producing a single

time-series observable through *Google* (Figure 3.1**g-i**).

The time series of search volume $S_\Theta(t)$ therefore appears as a super-position of impulse-response functions. The impulses are located at positions in time θ over the interval Θ, with a magnitude $\bar{\delta}(\theta)$. Each impulse-response has the shape of the reflected KAT. Thus, Eq. 3.5 is mathematically equivalent to a re-scaled (discrete form) convolution of the KAT with a function $f_\Theta(t)$:

$$S_\Theta(t) = \alpha\big[f_\Theta(t) * \kappa(t)\big] \tag{3.6}$$

where $f_\Theta(t)$ traces out the values of $\bar{\delta}(\theta = t)$ for all moments during the interval Θ and is zero everywhere else (see ED Figure A.11). Thus $f_\Theta(t)$ describes the number of time-associated targets associated with the period Θ at each time t. If $\bar{\delta}$ is a constant for all moments in Θ, then $f_\Theta(t)$ is a square wave with nonzero values during $t = \theta \in \Theta$. This formulation allows us to recover the form of the KAT impulse-response function using standard signal-processing techniques for deconvolution (Wiener, 1950; Silvia, 1987; Delaigle et al., 2008) applied to internet search data (Eq. 3.7 below).

Google Search Data

We obtain data on queries made to the *Google Search* engine from the *Google Trends API*. These data describe changes over time in the number of queries originating from a geographic region during a discrete time period which contain a specific string. For example, we examine queries that contain the string "2016", which aggregates search volume for queries such as "taxes 2016", "halloween 2016", and "2016 presidential election". Time-series provided by Google have been rescaled by an unknown factor, but this does not affect our analysis because search volume is already predicted to be rescaled by an unknown constant (Eq. 3.5) and these are all removed in the deconvolution. We focus on country-level data, which are available for most countries in the world and use a sample that begins January 1, 2008 to ensure data quality, since earlier data contains numerous irregularities. The temporal resolution of the available search volume data available depends on the length of the sample, with daily searches only available when the sample period does not exceed 270 days. This length of time is sufficient of analysis of monthly and daily targets (e.g. Θ = "July"), but is insufficient for the study of annual targets. Longer time series are provided using weekly search volumes, which we use in our analysis of annual targets, adjusting for the alternative time-step.

Targets corresponding to years We obtain search volume for queries containing the numerical representation of each year (e.g. "2016") between 2009 and 2018, inclusive(e.g. Fig 3.2**a**). Because these numerical values are language-independent, we obtain data from 181 countries representing 96.7% of the global population. This results in a longitudinal sample where an observation corresponds to a unique {country, week-of-query, target-year} triplet. Individual country-weeks appear more than once, because we examine searches for different target years originating on the same date. The final sample in our main analysis contains

246,131 observations. In our main analysis (Fig 3.3), we use data on queries occurring within the year that is the target period as well as 270 days before and after that year. We note that search volume for years exhibit irregular outliers where search volume can surge up to 2,700% for just a week, such as during the Olympics Games or the Football World Cup. To limit the influence of these outliers, pre-processing of raw query data rescales time-series to match medians in each country-by-target-year sample.

Targets corresponding to months We obtain search for queries containing the name of each month (e.g. "january") for each month. We construct samples that include 120 days before and after the endpoints of each target month (e.g. 270 days in total for a 30 day month). We focus this analysis on a more limited sample of 20 countries from five continents that are broadly representative of the global population—selected for breadth of representation across languages, cultures and religions—and where rates of *Google Search* usage are high (Argentina, Australia, Brazil, Chile, Colombia, France, Germany, Hong-Kong, India, Indonesia, Mexico, Netherlands, New Zealand, Norway, Peru, Singapore, Spain, Sweden, the United Kingdom, and the United States) and examine queries where the month is described using the primary language for each country (English, French, German, Spanish, Portuguese, Bahasa Indonesia, Swedish, Norwegian, or Dutch). The resulting sample contains 712,620 observations, each corresponding to a unique {country, day-of-query, target-month} triplet.

Targets corresponding to specific days We identify searches associated with specific dates by collecting search volume for searches containing the name of single-day holidays. We use the same sample of 20 countries used for month-related targets and select three holidays for each each country (e.g. "Cinco de Mayo" for Mexico, "Waitangi Day" for New-Zealand, and "National Day" for Singapore, see ED Table A.6 for a complete list). For each holiday, we examine search queries during a window of 270 days, centered on the holiday, resulting in a sample of 178,762 observations, each corresponding to a unique {country, day-of-query, target-holiday} triplet.

Empirical estimation of the KAT

For each class of target (year, month, or holiday), we estimate the KAT using *Google* search queries in a two step procedure. First, we non-parametrically deconvolve query volume to recover the KAT. This flexible approach recovers the KAT without requiring that we impose any assumptions about its form, however, because we have limited data, there is noise in these estimates which do not result in a parsimonious analytical expression. To address these issues, in the second step we fit smoothing functions to the estimates of the KAT obtained in the first step.

Deconvolution We obtain an estimate for the KAT by deconvolving Eq 3.6, a standard operation that is the inverse of the convolution (Wiener, 1950; Almon, 1965; Silvia, 1987;

Delaigle et al., 2008). The approach isolates an impulse-response function from time-series data when the timing and magnitude of impulses are known. We separately pool all search data for targets associated with years, months or days to recover an average estimate of the KAT corresponding to each class of target. For targets associated with years, we deconvolve these time series by estimating the distributed-lag regression with a weekly time-step:

$$S_{yit} = \sum_{L=-35}^{35} \beta_L f_{yi,t-L} + \epsilon_{yit} \tag{3.7}$$

where observations are indexed by the search target year y (e.g. $y = 2016$ for queries containing the string "2016"), country i where queries originate from, and week that queries are executed t. In the notation of the theoretical framework above, the interval of interest Θ is equal to the target year y. The variable f_{yit} is equal to one if the date of the search t is contained in the year y, and zero otherwise. This dummy variable corresponds to the function $f(t)$ in Eq. 3.6 and appears as a square wave if it is plotted for a single i and y over time (as in ED Figure A.11). Coefficients β_L are estimated for 35 weekly leads and 35 weekly lags of the variable f_{yit}, in addition to its contemporaneous values where $L = 0$, and they flexibly recover the impulse-response function without imposing any assumptions on its form. Because there are a small number of extreme outlier observations, such as when the Olympics or World Cup is occurring (described above), we estimate coefficients in Eq. 3.7 via the robust regression procedure in ref. [(Li, 1985)] that limits the influence of these outlier observations. This procedure (Berk, 1990) first estimates coefficients in Eq. 3.7 using ordinary least-squares, then iteratively re-fits the model by down-weighting observations with large estimated values for $\hat{\epsilon}_{yit}$. The resulting estimated coefficient $\hat{\beta}_L = \kappa(-L)$ corresponds to our estimate for the the KAT at relative time $-L$. The $\hat{\beta}_L$ are displayed as black dots in Figure 3.3**c-e**. Unexplained variations in search volume correspond to the residuals ϵ_{yit}. We also estimate a separate version of Eq 3.7 for monthly targets, where year index y is replaced by month index m and f_{mit} is equal to one if t is in month m and zero otherwise, as well as a version for holiday targets, where y is replaced by holiday h and f_{hit} is one if $t = h$ and zero otherwise. In models with month and holiday targets, t indexes daily observations. Results from these three estimates explain a large amount of variation in search queries for these three classes of target (years: $R^2 = 0.93$; months: $R^2 = 0.83$; holidays: $R^2 = 0.81$) when convolved with their respective versions of $f(t)$, displayed in Figure A.15. To examine the robustness of these results, we also estimate Eq 3.7 including year and/or country-specific constants (fixed effects) that flexibly adjust for differences in search volume between countries, for example accounting for differences in internet access or culture, and for nonlinear trends in search volume over time, for example due to changing usages patterns within countries. We find that adjusting for these factors in Eq. 3.7 has essentially no effect on the estimated structure of the KAT (see ED Figure A.13).

Fitting an analytical form We seek a smooth analytical function that approximates the KAT estimated in Eq. 3.7 by fitting different families of functions, evaluating both their

goodness-of-fit to the coefficients $\hat{\beta}_L$ and evaluating their ability to accurately predict search volume when convolved with $f(t)$. For each family of functions, we independently fit parameters to the future-oriented and past-oriented portions of the KAT, keeping the present in both samples and forcing the two functions to intersect at the present. We estimate exponential functions of the form $e^{a(\theta-t)+b}$, motivated by intuition from economic discounting (Ramsey, 1928; Gollier, 2013) and market interest rates (Fisher, 1930; Von Mises, 2016) where inter-temporal tradeoffs between sequential future decisions should be self-consistent. We also estimate 8^{th} order polynomial functions of the form $\sum_{p=0}^{8}(a_p(\theta-t)^p)$, motivated by their potential parsimony and simplicity. We also estimate functions that resemble models of "hyperbolic discounting" behavior widely documented in psychology and behavioral economics (McClure et al., 2004; Angeletos et al., 2001; Laibson, 1997; Rubinstein, 2003), which take the form $\frac{a}{1+b(\theta-t)}$. Note that while exponential and hyperbolic-discounting-like forms are motivated by prior research studying tradeoffs over time (i.e. "discounting"), there is not a clear theoretical linkage between the form of the KAT and the tradeoffs we expect populations to make over time. Additionally, note that neither of these existing research literature provide any guidance on the form of the past-oriented portion of the KAT. Lastly, we estimate rational functions of the form $\frac{a+b(\theta-t)+c(\theta-t)^2}{1+d(\theta-t)+e(\theta-t)^2}$ where the numerator is allowed to be up to the second order and the denominator is restricted to be second order. The family of rational functions, broadly defined, technically includes the polynomial forms and hyperbolic-discounting-like forms that we examine where each has specific restrictions on the order of the numerator and/or denominator. Comparisons between the fitted functions within each family are shown as colored lines in Figure 3.3**c**, and results from their convolution with $f(t)$ are shown in ED Figure A.15. Fitted values from the rational family of functions minimize RMSE and also predict search volume with lowest errors. Using our sample of annual targets and 181 countries, the average KAT is estimated to have the form:

$$\hat{\kappa}(\theta-t)=\begin{cases}\frac{-0.0037(\theta-t)^2-0.75(\theta-t)+25.35}{-0.207(\theta-t)^2-74.41(\theta-t)+100} & \text{if} \quad (\theta-t)<0 \quad \text{(i.e. the past)}\\ \frac{-0.001(\theta-t)^2+0.13(\theta-t)+25.30}{-0.2(\theta-t)^2+51.1(\theta-t)+100} & \text{if} \quad (\theta-t)>0 \quad \text{(i.e. the future)}\end{cases} \tag{3.8}$$

which is displayed in Figure 3.3**a**.

Uncertainty We compute uncertainty in these results by block-bootstrap resampling of the original data (blocks corresponding with country-and-target-specific time series) and repeating estimation of Eq. 3.7 and re-fitting the rational function analytical form for the KAT each time. Because the range of estimates is extremely narrow, in Figure 3.3**a** we display the full range of values (maximum to minimum) recovered at each moment in relative time (95% confidence intervals, which are standard, are too narrow to see). We also compute the fraction of attention allocated to the past, present, and future in each resampled estimate, as well as the relative fraction of attention, both displayed in Figure A.12.

Computing total attention to the future and past To compute attention to the past or future, we sum total attention to moments in time before or after the present. To do

this, we sum $\hat{\beta}_L = \kappa(-L)$ from $L = -35$ weeks (245 days) to $L = -1$ for the past and from $L = 1$ to $+35$ for the future. We compare these values to the probability mass at $L = 0$, which we interpret as attention to the present. We display these three values in ED Figure A.16**a** (Figure 3.3**a** breaks down attention to the future and past into "distant" and "near" components by introducing an additional cutoff at 100 days in both directions). We use block-bootstrap resampling to compute uncertainty in these measures of attention, estimating Eq. 3.7, fitting an analytical function, and integrating the resulting function following each resampling of the data. ED Figure A.12**a** shows the distribution of total attention allocated to past, present, and future moments using the years targets, and ED Figure A.12**b** details the relative fraction of attention from these re-sampled values.

Country specific estimates To compare the allocation of attention over time across countries, we estimate the KAT only using data from one country at a time. We estimate Eq 3.7 for each country separately and use the country-specific values for $\hat{\beta}_L$ to fit country-specific rational functions. Figures 3.3**b** and ED Figure A.14 show country specific KAT estimates for our sample of twenty representative countries, the later displaying estimates using annual, monthly, and daily targets. We integrate these curves using the approach above to compute total attention to the past, present, and future for each country, using annual targets. The distribution of values for all 181 countries in our sample are displayed in Figure 3.3**b** in three dimensions (also in ED Figure A.16, broken down by region) and in the maps of Figure 3.4**a-b**.

Estimating trends in the KAT To compute trends in attention to the past present and future, we pool query data for targets across our sample of twenty countries shown in Figure 3.3**b**. We then estimate the KAT for each class of target only using data from a single year. We integrate the resulting estimated fraction of attention to past/present/future. These values are plotted in Figure 3.4**c** using models for each class of target for each of the ten years in our sample.

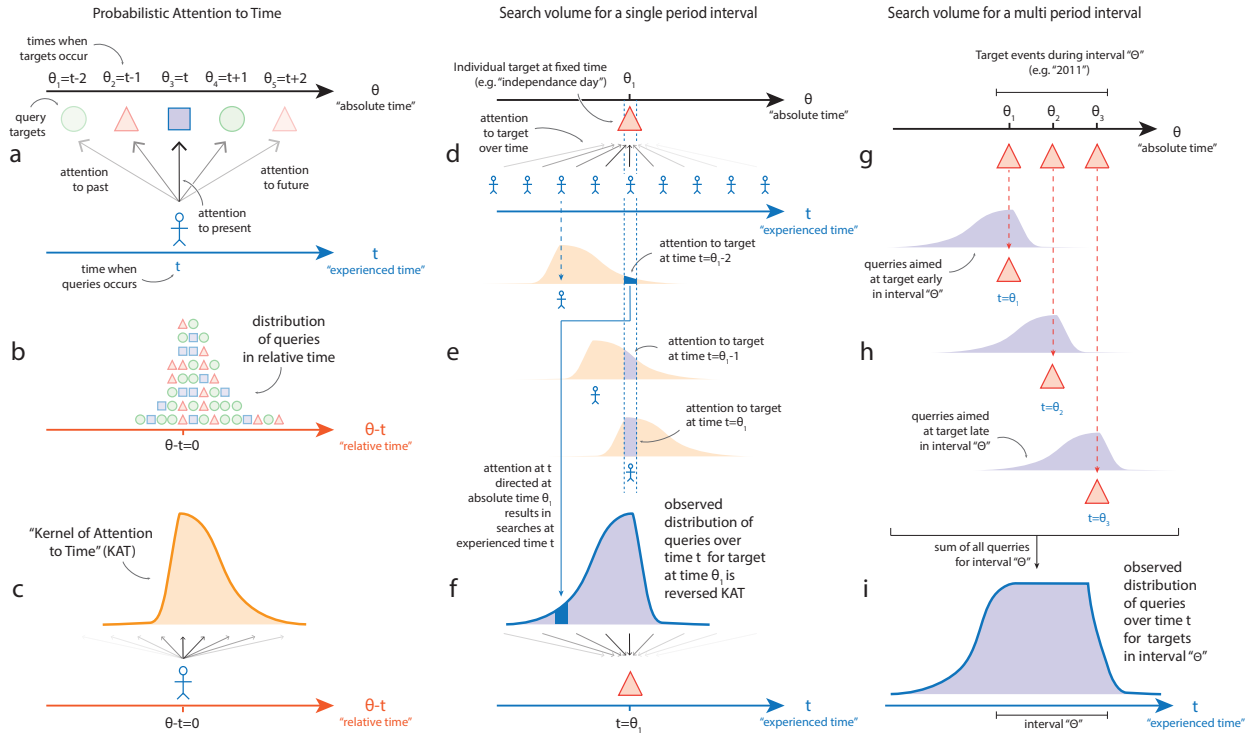

Figure 3.1: **A probabilistic model of attention to time and predictions for *Google Search* queries.** Search volume for time-specific queries captures the distribution of attention to time for a population of users (also see *methods*). **(a)** Users traveling through "experienced time" t conduct search queries for specific targets (e.g. holidays) that occur at moments in "absolute time" $\theta = \theta_1$, θ_2, etc. Queries where $\theta < t$ indicate attention to past events, $\theta = t$ indicates attention to the present, and $\theta > t$ indicates attention to the future. **(b)** A large number of time-related queries can be organized by their position in "relative time" $(\theta - t)$, i.e. the distance in time between when the target occurs (θ) and the time when the query is executed (t). **(c)** We hypothesize that the probability distribution of time-related queries in relative time is approximately stable for a fixed population. We call this probability distribution the *kernel of attention to time* (KAT). **(d)** Data on search queries is organized around the query targets (e.g. searches for "Diwali"), not around the individuals who execute the query. As individuals move through time, passing a target that is fixed in absolute time, attention to the target will evolve as the target is initially in the future, then present, then past. **(e)** The quantity of attention to a target at a moment is scaled by the KAT. As a population approaches and passes a target, attention to the target traces out a KAT distribution that moves through time, centered on the moment in time inhabited by the population. **(f)** Search volume data is recorded based on the time when the search is executed. Thus if search volume scales with the KAT, then a time series of search volume for a single target at a fixed moment will have the same shape as the KAT, but reversed (reflected) in time (e.g. before the event occurs, search volume for the target will reflect attention to the future). **(g)** When the targets of a common search query spans an interval of time Θ, it may not possible to identify which specific moment is the target of a given search query (e.g. queries containing "september" or "2014" could come from targets on any date during those periods). **(h)** Each target identified by a query for Θ will have a distribution of search volume associated with it. **(i)** The resulting search volume associated with interval Θ is the sum of search volume distributions for all targets in interval Θ. This implies that the distribution of search volume for interval Θ is the convolution of the KAT, shown in (c), with the time series of targets in interval Θ, shown in (g).

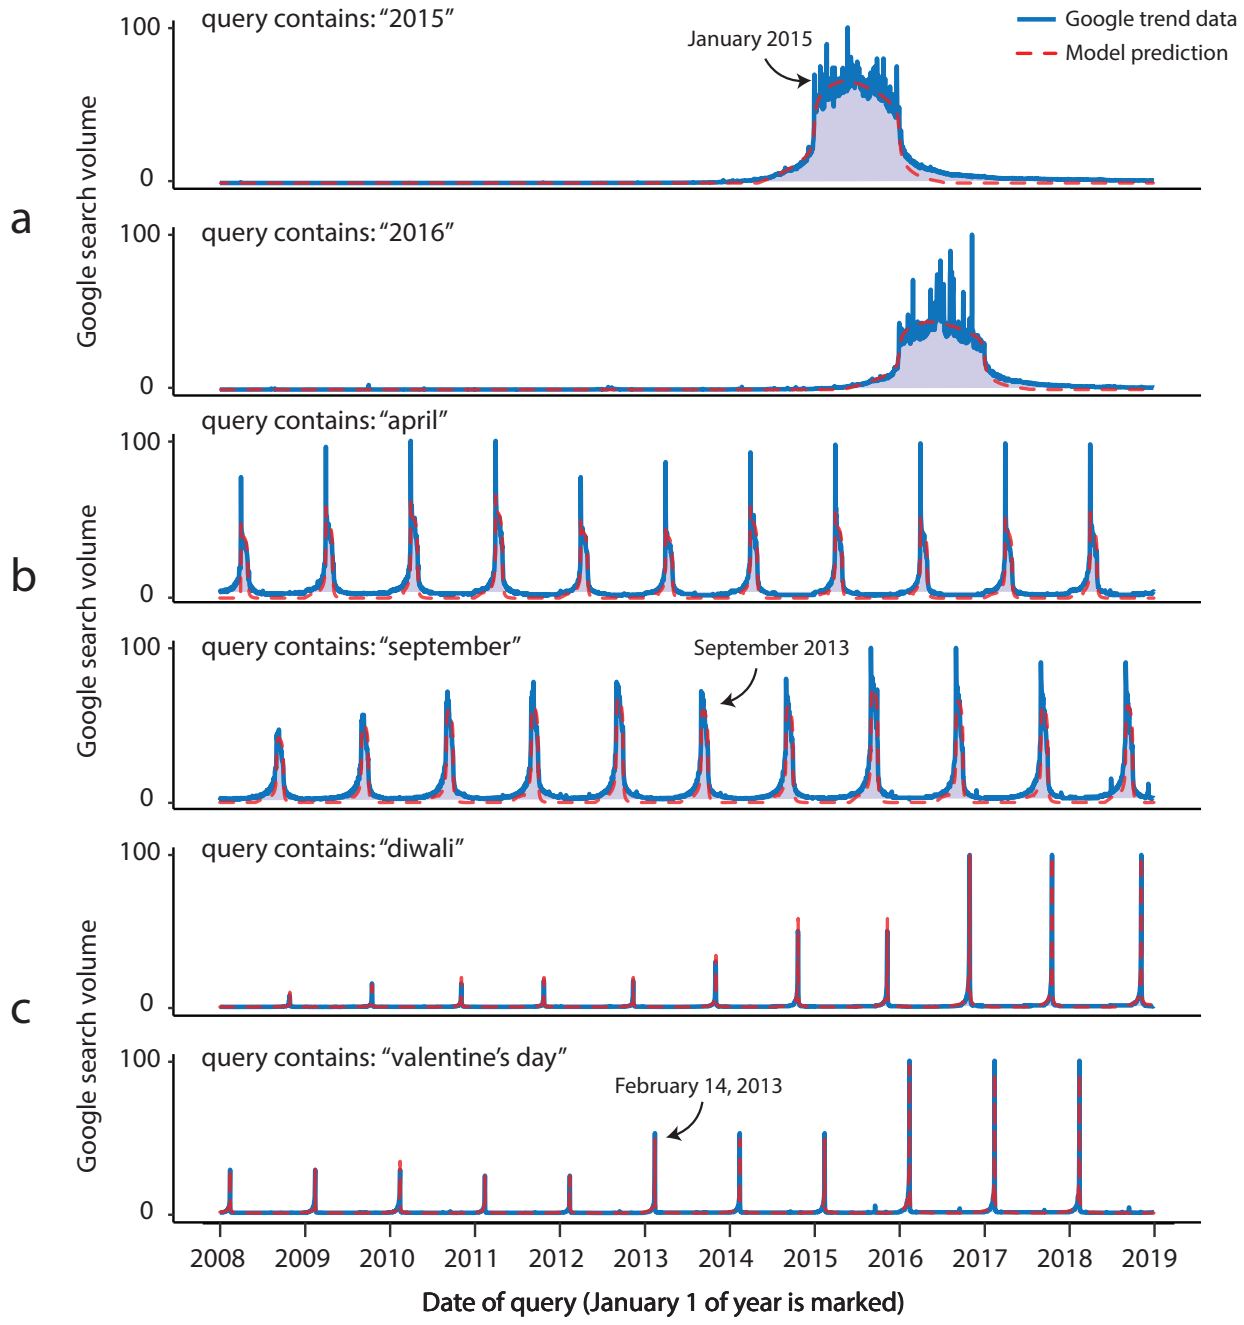

Figure 3.2: **Example *Google Search* volume for time-related queries and KAT-based predictions** **(a)** Global *Google Search* volume for queries containing year targets "2015" and "2016" (blue) and out-of-sample predictions using a single KAT-based model fitted to a 20 countries restricted sample (red line, analogous to Figure 1**i**). **(b)** Same as (a), but for queries containing month targets "april" and "september". **(c)** Same as (a), but for single-day targets "diwali" and "valentine's day".

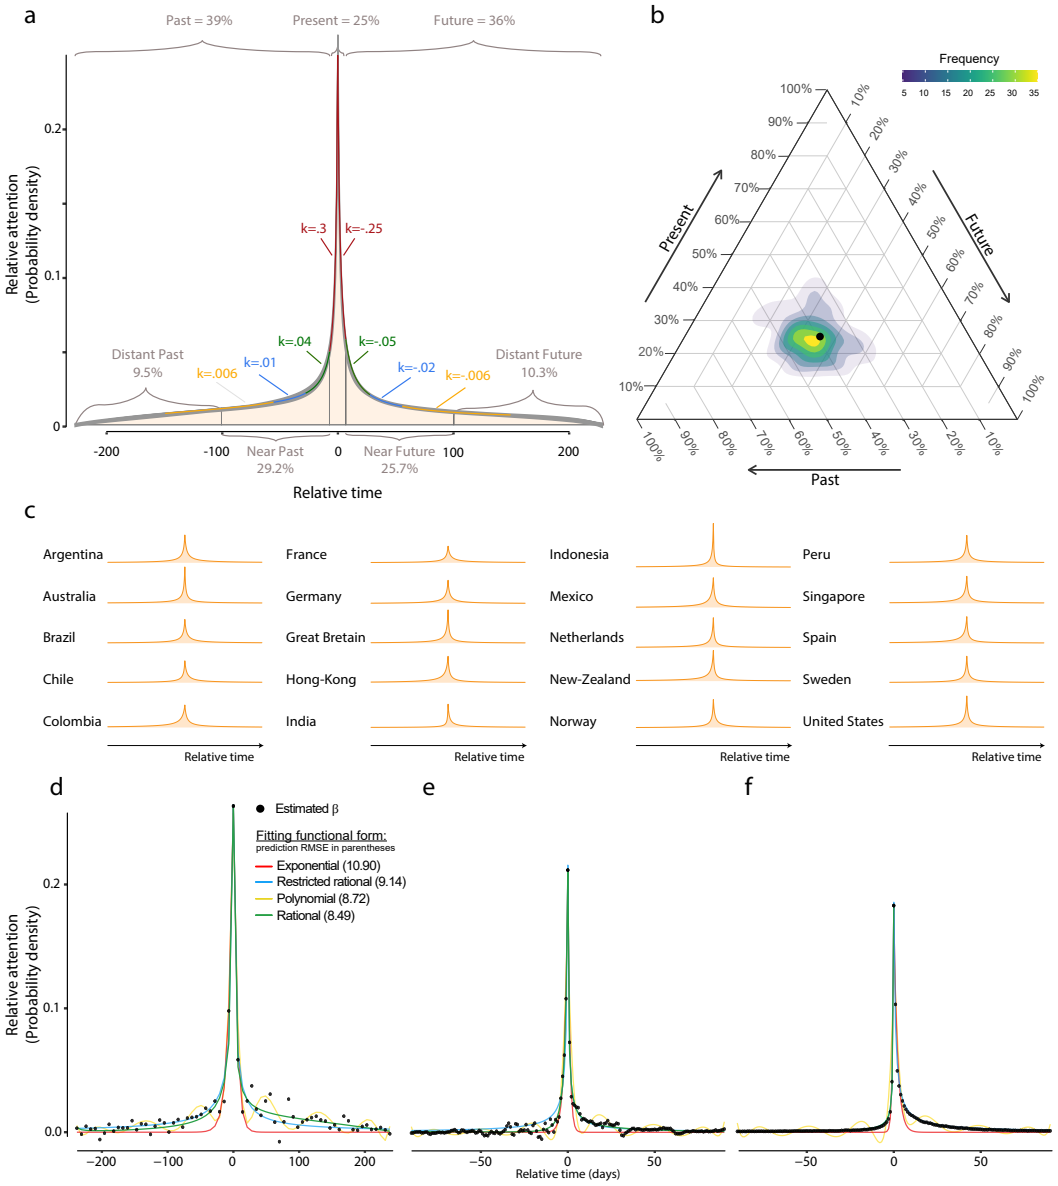

Figure 3.3: **The Kernel of Attention to Time (KAT). (a)** The estimated average probability distribution for the allocation of attention to different moments in relative time (i.e. the KAT) for countries representing 97% of the global population. Estimate is constructed by deconvolving search queries for year targets that are language independent. Colored segments indicate approximate day-over-day rates of decay (k) for each segment as a function of temporal distance from the present. Gray area denotes the range of statistical uncertainty, bounded by the maximum and minimum values estimated using 10,000 block-resampled versions of the data. **(b)** Ternary plot showing the total fraction of attention allocated to the past, present, and future. The black dot corresponds to the global estimate in (a), the heatmap shows the distribution of estimates from 181 different countries. **(c)** KATs independently estimated in each of 20 selected countries using the same year targets. **(d)-(f)** Estimated KAT computed by deconvolving searches from targets of different lengths: **(d)** annual targets (eg. "2016"), **(e)** monthly targets (eg. "april"), and **(f)** daily targets (eg. "diwali", listed in Extended Data Table A.6). (d)-(f) pool data from 20 countries shown in (c) and show coefficient estimates recovered via deconvolution (black dots). Four functional forms are fitted to these coefficients: exponential (red), a "restricted rational" function ($\frac{a}{1+b(\theta-t)}$) similar to hyperbolic-discounting models (Laibson, 1997; Findley, 2015)(blue), eighth-order polynomial (yellow), and general rational functions that are the ratio of two polynomials (green). See Extended Data Figure A.15 for the performance of each functional form in reconstructing search query volume, and Extended Data Figure A.14 for individual KAT estimates for each country and class of target.

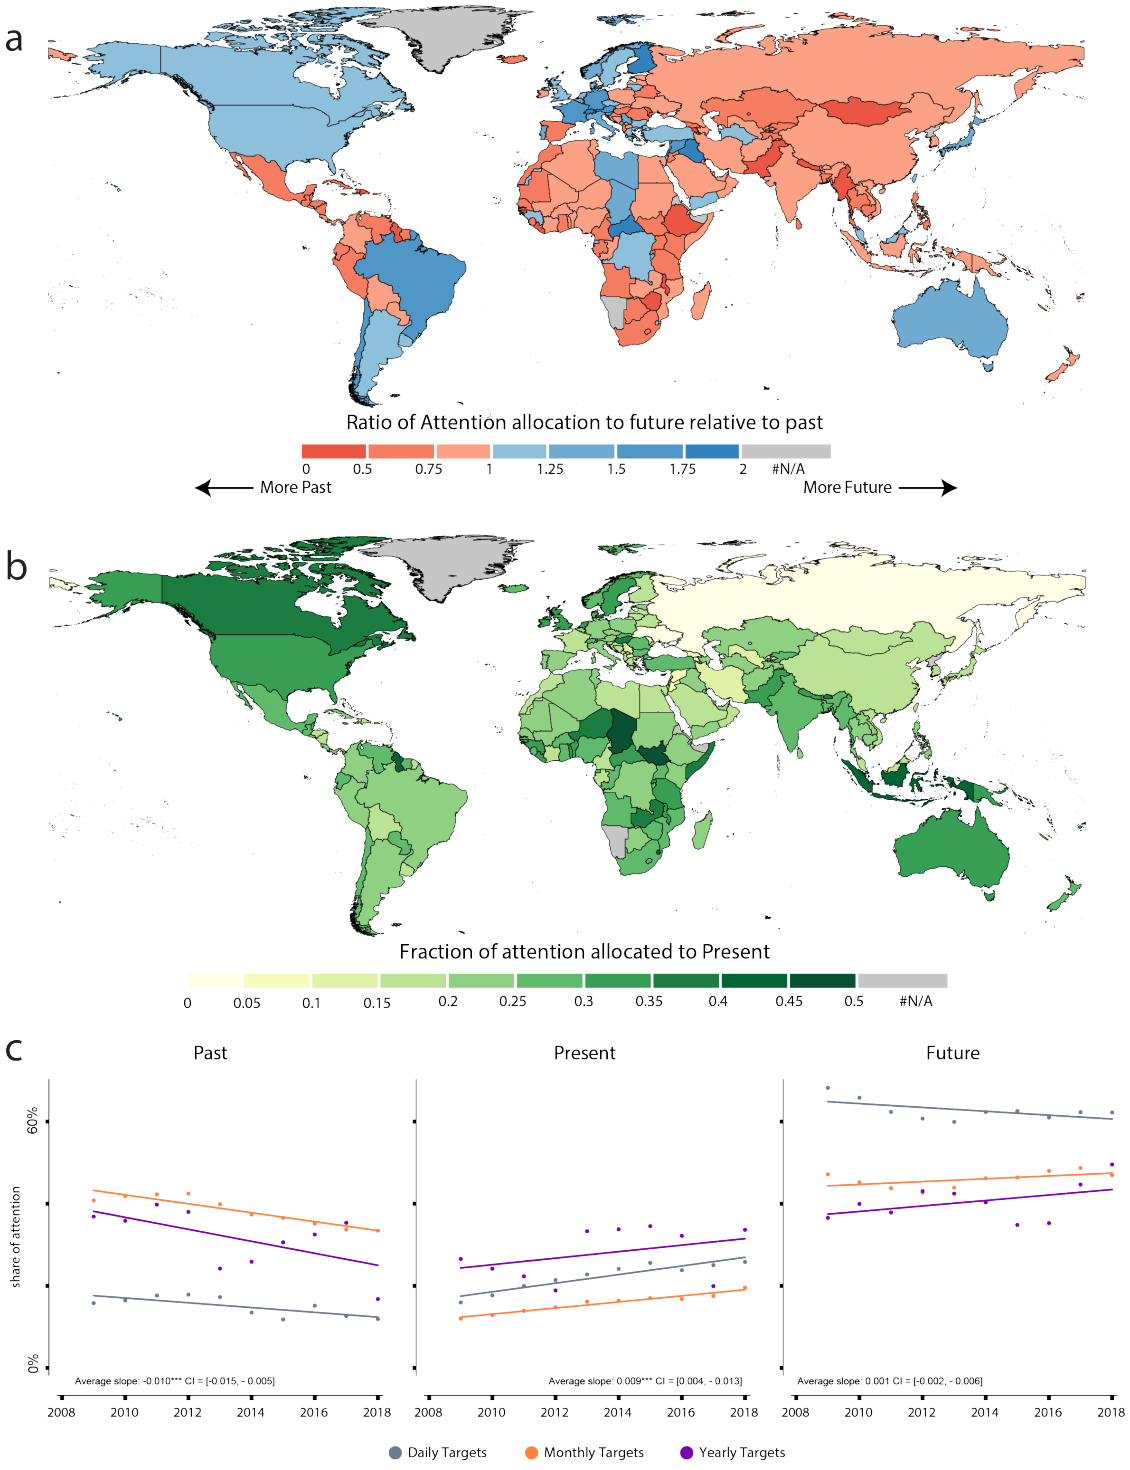

Figure 3.4: **Relative attention to past, present, and future around the world and over time. (a)** The estimated ratio of average attention allocated to the future relative to the past in each country. A ratio greater than one indicates that more attention is allocated to the future, a ratio of one indicates equal attention, a ratio below one indicates more attention is allocated to the past. **(b)** The fraction of attention allocated to the present in each country. All three shares (past, present, future) are shown in ED Figure A.16 for each country and region. **(c)** The evolution of shares of attention allocated towards the past, present, and future over the decade 2009-2018. Trends are seperately estimated for each class of search target (grey=day, orange=month, purple=year). Sample is the set of twenty countries used in Figure 3.3(d)-(f).

Bibliography

Yann Algan, Fabrice Murtin, Elizabeth Beasley, Kazuhito Higa, and Claudia Senik. Well-being through the lens of the internet. *PloS one*, 14(1), 2019.

Shirley Almon. The distributed lag between capital appropriations and expenditures. *Econometrica: Journal of the Econometric Society*, pages 178–196, 1965.

George-Marios Angeletos, David Laibson, Andrea Repetto, Jeremy Tobacman, and Stephen Weinberg. The hyperbolic consumption model: Calibration, simulation, and empirical evaluation. *Journal of Economic perspectives*, 15(3):47–68, 2001.

Sébastien Annan-Phan and Fabien Roques. Market integration and wind generation: An empirical analysis of the impact of wind generation on cross-border power prices. *The Energy Journal*, 39(3):1–25, 2018.

ArdorSeo. How many google searches per day on average in 2020? data retrieved from the Ardor Seo website (accessed March 12, 2020), `https://ardorseo.com/blog/how-many-google-searches-per-day`.

Barak Ariel, David Lawes, Cristobal Weinborn, Ron Henry, Kevin Chen, and Hagit Brants Sabo. The "less-than-lethal weapons effect"-introducing tasers to routine police operations in england and wales: A randomized controlled trial. *Criminal Justice and Behavior*, 46 (2):280–300, 2019.

Maximilian Auffhammer and Erin T Mansur. Measuring climatic impacts on energy consumption: A review of the empirical literature. *Energy Economics*, 46:522–530, 2014.

Maximilian Auffhammer, Patrick Baylis, and Catherine H Hausman. Climate change is projected to have severe impacts on the frequency and intensity of peak electricity demand across the united states. *Proceedings of the National Academy of Sciences*, 114(8):1886–1891, 2017.

Axon. Taser protect life: Annual conducted electrical weapon (cew). user update. 20.2, 2018.

Bocar Ba and Jeffrey Grogger. The introduction of tasers and police use of force: Evidence from the chicago police department. Technical report, National Bureau of Economic Research, 2018.

Duren Banks, Paul Ruddle, Erin Kennedy, and Michael G. Planty. Arrest-related deaths program redesign study, 2015-16: Preliminary findings. *U.S. Department of Justice-Bureau of Justice Statistics*, December 2016.

Charles de Secondat Baron de Montesquieu. *Esprit des lois*, volume 1. Firmin Didot frères, fils et cie, 1872.

Alan Barreca, Karen Clay, Olivier Deschenes, Michael Greenstone, and Joseph S Shapiro. Adapting to climate change: The remarkable decline in the us temperature-mortality relationship over the twentieth century. *Journal of Political Economy*, 124(1):105–159, 2016.

Paul A. Bell and Robert A. Baron. Aggression and heat: The mediating role of negative affect. *Journal of Applied Social Psychology*, 6(1):18–30, 1976.

RA Berk. A primer on robust regression. In J. Fox and JS Long, editors, *Modern Methods of Data Analysis*. Sage Publications. London, 1990.

John N Boyd and Philip G Zimbardo. Time perspective, health, and risk taking. In *Understanding behavior in the context of time*, pages 97–119. Psychology Press, 2006.

Domenica Bueti, Bahador Bahrami, and Vincent Walsh. Sensory and association cortex in time perception. *Journal of Cognitive Neuroscience*, 20(6):1054–1062, 2008.

Marshall Burke, Solomon M. Hsiang, and Edward Miguel. Climate and conflict. *Annual Review of Economics*, 7(1):577–617, 2015.

Alexander A Bustamante. Overview of less-lethal force tools and deployment, 2017.

Cristian Candia, C Jara-Figueroa, Carlos Rodriguez-Sickert, Albert-László Barabási, and César A Hidalgo. The universal decay of collective memory and attention. *Nature Human Behaviour*, 3(1):82–91, 2019.

Tamma Carleton, Michael Delgado, Michael Greenstone, Trevor Houser, et al. Valuing the global mortality consequences of climate change accounting for adaptation costs and benefits. 2018.

Aaron Chalfin and Justin McCrary. Are us cities underpoliced? theory and evidence. *Review of Economics and Statistics*, 100(1):167–186, 2018.

Aaron Chalfin, Benjamin Hansen, Jason Lerner, and Lucie Parker. Reducing crime through environmental design: Evidence from a randomized experiment of street lighting in new york city. Technical report, National Bureau of Economic Research, 2019.

Hyunyoung Choi and Hal Varian. Predicting the present with google trends. *Economic Record*, 88:2–9, 2012.

Steve Cicala. Imperfect markets versus imperfect regulation in us electricity generation. Technical report, National Bureau of Economic Research, 2017.

CPD. Chicago police department general order g03-02-04, 2017.

Aurore Delaigle, Peter Hall, Alexander Meister, et al. On deconvolution with repeated measurements. *The Annals of Statistics*, 36(2):665–685, 2008.

Olivier Deschenes and Michael Greenstone. Climate change, mortality, and adaptation: Evidence from annual fluctuations in weather in the us. *American Economic Journal: Applied Economics*, 3(4):152–85, October 2011a.

Olivier Deschenes and Michael Greenstone. Climate change, mortality, and adaptation: Evidence from annual fluctuations in weather in the us. *American Economic Journal: Applied Economics*, 3(4):152–85, 2011b.

Jennifer L Doleac and Nicholas J Sanders. Under the cover of darkness: How ambient light influences criminal activity. *Review of Economics and Statistics*, 97(5):1093–1103, 2015.

Michael S Drake, Darren T Roulstone, and Jacob R Thornock. Investor information demand: Evidence from google searches around earnings announcements. *Journal of Accounting research*, 50(4):1001–1040, 2012.

David M Eagleman, U Tse Peter, Dean Buonomano, Peter Janssen, Anna Christina Nobre, and Alex O Holcombe. Time and the brain: how subjective time relates to neural time. *Journal of Neuroscience*, 25(45):10369–10371, 2005.

Frank Edwards, Michael H. Esposito, and Hedwig Lee. Risk of police-involved death by race/ethnicity and place, united states, 2012-2018. *American Journal of Public Health*, 108(9):1241–1248, 09 2018.

EPA. Inventory of u.s. greenhouse gas emissions and sinks 1990-2017. *United States Environmental Protection Agency 430-P-19-001*, 2019.

Michael Ettredge, John Gerdes, and Gilbert Karuga. Using web-based search data to predict macroeconomic statistics. *Communications of the ACM*, 48(11):87–92, 2005.

Václav Fanta, Miroslav Šálek, and Petr Sklenicka. How long do floods throughout the millennium remain in the collective memory? *Nature communications*, 10(1):1–9, 2019.

T Scott Findley. Hyperbolic memory discounting and the political business cycle. *European Journal of Political Economy*, 40:345–359, 2015.

Irving Fisher. *Theory of interest: as determined by impatience to spend income and opportunity to invest it*. Augustusm Kelly Publishers, Clifton, 1930.

Roland G. Jr. Fryer. An empirical analysis of racial differences in police use of force. *Journal of Political Economy*, Forthcoming 2018a.

Roland G. Jr. Fryer. Reconciling results on racial differences in police shootings. *AEA Papers and Proceedings*, 108:228–33, 2018b.

Ruth García-Gavilanes, Milena Tsvetkova, and Taha Yasseri. Dynamics and biases of online attention: the case of aircraft crashes. *Royal Society open science*, 3(10):160460, 2016.

Sharad Goel, Jake M Hofman, Sébastien Lahaie, David M Pennock, and Duncan J Watts. Predicting consumer behavior with web search. *Proceedings of the National academy of sciences*, 107(41):17486–17490, 2010.

Christian Gollier. *Pricing the planet's future: the economics of discounting in an uncertain world*. Princeton University Press, 2013.

Felipe Goncalves and Steven Mello. A few bad apples? racial bias in policing. *Racial Bias in Policing (June 15, 2020)*, 2020.

Google. Google trends. data retrieved from the Google Trends website (accessed October 2019), `https://www.google.com/trends`.

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

John F Gunn III and David Lester. Using google searches on the internet to monitor suicidal behavior. *Journal of affective disorders*, 148(2-3):411–412, 2013.

Dan Hinkel and Jennifer Smith Richards. Chicago cops get more tasers, but red flags remain, 2017.

Solomon M. Hsiang, Kyle C. Meng, and Mark A. Cane. Civil conflicts are associated with the global climate. *Nature*, 476:438 EP–, 2011.

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

Internet World Stats. Internet users statistics for africa, 2019. data retrieved from the Internet World Stats Indicators (accessed December 26, 2019), `https://www.internetworldstats.com/stats.htm`.

Brian Jacob, Lars Lefgren, and Enrico Moretti. The Dynamics of Criminal Behavior: Evidence from Weather Shocks. *Journal of Human Resources*, 42(3), 2007.

C Jara-Figueroa, Amy Z Yu, and César A Hidalgo. How the medium shapes the message: Printing and the rise of the arts and sciences. *PloS one*, 14(2):e0205771, 2019.

Daniel Kahneman. *Attention and effort*, volume 1063. Prentice-Hall Inc., 1973.

Jacob Kaplan. Uniform crime reporting program data: Offenses known and clearances by arrest, 1960-2016, 2018.

Jan Horst Keppler, Sebastien Phan, and Yannick Le Pen. The impacts of variable renewable production and market coupling on the convergence of french and german electricity prices. *The Energy Journal*, 37(3):343–359, 2016.

Matthew Wilburn King. How brain biases prevent climate action, 2019. (accessed July, 2020), `https://www.bbc.com/future/article/20190304-human-evolution-means-we-can-tackle-climate-change`.

Dean Knox, Will Lowe, and Jonathan Mummolo. The bias is built in: How administrative records mask racially biased policing. *Working Paper*, 2019.

David Laibson. Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics*, 112(2):443–478, 1997.

Benjamin Leard, Kevin Roth, et al. Weather, traffic accidents, and climate change. *Resources for the Future Discussion Paper*, pages 15–19, 2015.

Penny A Lewis and R Chris Miall. Brain activation patterns during measurement of sub-and supra-second intervals. *Neuropsychologia*, 41(12):1583–1592, 2003.

Guoying Li. Robust regression. *Exploring data tables, trends, and shapes*, 281:U340, 1985.

Kristin Linnerud, Torben K Mideksa, and Gunnar S Eskeland. The impact of climate change on nuclear power supply. *The Energy Journal*, pages 149–168, 2011.

Philipp Lorenz-Spreen, Bjarke Mørch Mønsted, Philipp Hövel, and Sune Lehmann. Accelerating dynamics of collective attention. *Nature communications*, 10(1):1–9, 2019.

John M MacDonald and Jeffrey Fagan. Using shifts in deployment and operations to test for racial bias in police stops. In *AEA Papers and Proceedings*, volume 109, pages 148–51, 2019.

Samuel M McClure, David I Laibson, George Loewenstein, and Jonathan D Cohen. Separate neural systems value immediate and delayed monetary rewards. *Science*, 306(5695):503–507, 2004.

Grant R McDermott and Oivind A Nilsen. Electricity prices, river temperatures, and cooling water scarcity. *Land Economics*, 90(1):131–148, 2014.

John McWhorter. Police kill too many people- white and black. *Time Magazine*, 2016.

Torben K Mideksa and Steffen Kallbekken. The impact of climate change on the electricity market: A review. *Energy policy*, 38(7):3579–3585, 2010.

Antony Millner and Geoffrey Heal. Discounting by committee. *Journal of Public Economics*, 167:91–104, 2018.

Craig Morgan and Richard de Dear. Weather, clothing and thermal adaptation to indoor climate. *Climate Research*, 24.3:267–284, 2003.

Anna Nobre and Jennifer Coull. *Attention and time*. Oxford University Press, USA, 2010.

Takao Noguchi, Neil Stewart, Christopher Y Olivola, Helen Susannah Moat, and Tobias Preis. Characterizing the time-perspective of nations with search engine query data. *PloS one*, 9(4):e95209, 2014.

William D Nordhaus. The political business cycle. *The Review of Economic Studies*, 42(2): 169–190, 1975.

Christian NL Olivers and Martijn Meeter. A boost and bounce theory of temporal attention. *Psychological review*, 115(4):836, 2008.

Aurelie Ouss and John Rappaport. Is police behavior getting worse? the importance of data selection in evaluating the police. *University of Chicago Coase-Sandor Institute for Law & Economics Research Paper*, (865), 2019.

Joseph J Paton and Dean V Buonomano. The neural basis of timing: Distributed mechanisms for diverse functions. *Neuron*, 98(4):687–705, 2018.

Tobias Preis, Helen Susannah Moat, H Eugene Stanley, and Steven R Bishop. Quantifying the advantage of looking forward. *Scientific reports*, 2(1):1–2, 2012.

Tobias Preis, Helen Susannah Moat, and H Eugene Stanley. Quantifying trading behavior in financial markets using google trends. *Scientific reports*, 3:1684, 2013.

Yun Qiu and Jinhua Zhao. Too hot to focus: The mean and distributional effects of heat on labor productivity. *Available at SSRN 3421890*, 2019.

Stefan Rahmstorf, Jason E Box, Georg Feulner, Michael E Mann, Alexander Robinson, Scott Rutherford, and Erik J Schaffernicht. Exceptional twentieth-century slowdown in atlantic ocean overturning circulation. *Nature climate change*, 5(5):475, 2015.

Frank Plumpton Ramsey. A mathematical theory of saving. *The Economic Journal*, 38(152): 543–559, 1928.

Matthew Ranson. Crime, weather, and climate change. *Journal of Environmental Economics and Management*, 67(3):274 – 302, 2014a.

Matthew Ranson. Crime, weather, and climate change. *Journal of environmental economics and management*, 67(3):274–302, 2014b.

Henry L Roediger and K Andrew DeSoto. Forgetting the presidents. *Science*, 346(6213): 1106–1109, 2014.

Cody T. Ross. A multi-level bayesian analysis of racial bias in police shootings at the county-level in the united states, 2011-2014. *PLOS ONE*, 10(11):1–34, 11 2015.

Ariel Rubinstein. "economics and psychology"? the case of hyperbolic discounting. *International Economic Review*, 44(4):1207–1216, 2003.

G Rupert Jr et al. *Simultaneous statistical inference*. Springer Science & Business Media, 2012.

Wolfram Schlenker and Michael J. Roberts. Nonlinear temperature effects indicate severe damages to u.s. crop yields under climate change. *Proceedings of the National Academy of Sciences*, 106(37):15594–15598, 2009.

Manuel T. Silvia. Deconvolution. In Douglas F. Elliott, editor, *Handbook of Digital Signal Processing*, pages 741 – 788. Academic Press, San Diego, 1987.

Seth Stephens-Davidowitz. The cost of racial animus on a black candidate: Evidence using google search data. *Journal of Public Economics*, 118:26–40, 2014.

Seth Stephens-Davidowitz and Andrés Pabon. *Everybody lies: Big data, new data, and what the internet can tell us about who we really are*. HarperCollins New York, 2017.

Kevin E Trenberth, John T Fasullo, and Theodore G Shepherd. Attribution of climate extreme events. *Nature Climate Change*, 5(8):725, 2015.

Ludwig Von Mises. *Human action*. Lulu Press, Inc, 2016.

Aldert Vrij, Jaap Van Der Steen, and Leendert Koppelaar. Aggression of police officers as a function of temperature: An experiment with the fire arms training system. *Journal of Community & Applied Social Psychology*, 4(5):365–370, 1994.

Ingmar Weber and Carlos Castillo. The demographics of web search. In *Proceedings of the 33rd international ACM SIGIR conference on Research and development in information retrieval*, pages 523–530, 2010.

Ingmar Weber and Alejandro Jaimes. Who uses web search for what: and how. In *Proceedings of the fourth ACM international conference on Web search and data mining*, pages 15–24, 2011.

Leonie Wenz, Anders Levermann, and Maximilian Auffhammer. North–south polarization of european electricity consumption under future warming. *Proceedings of the National Academy of Sciences*, 114(38):E7910–E7918, 2017.

Norbert Wiener. *Extrapolation, interpolation, and smoothing of stationary time series: with engineering applications*. MIT press, 1950.

Paul Wohlford. Extension of personal time, affective states, and expectation of personal death. *Journal of Personality and Social Psychology*, 3(5):559, 1966.

Frank A Wolak. Measuring the benefits of greater spatial granularity in short-term pricing in wholesale electricity markets. *American Economic Review*, 101(3):247–52, 2011.

Lynn Wu and Erik Brynjolfsson. The future of prediction: How google searches foreshadow housing prices and sales. In *Economic analysis of the digital economy*, pages 89–118. University of Chicago Press, 2015.

Philip Zimbardo and John Boyd. Putting time in perspective: A valid, reliable individual-differences metric. *Journal of Personality and Social Psychology*, 77:1271–1288, 12 1999. doi: 10.1037/0022-3514.77.6.1271.

Appendix A

A.1 Adaptation Through Market Integration

Figure A.1: Time coverage

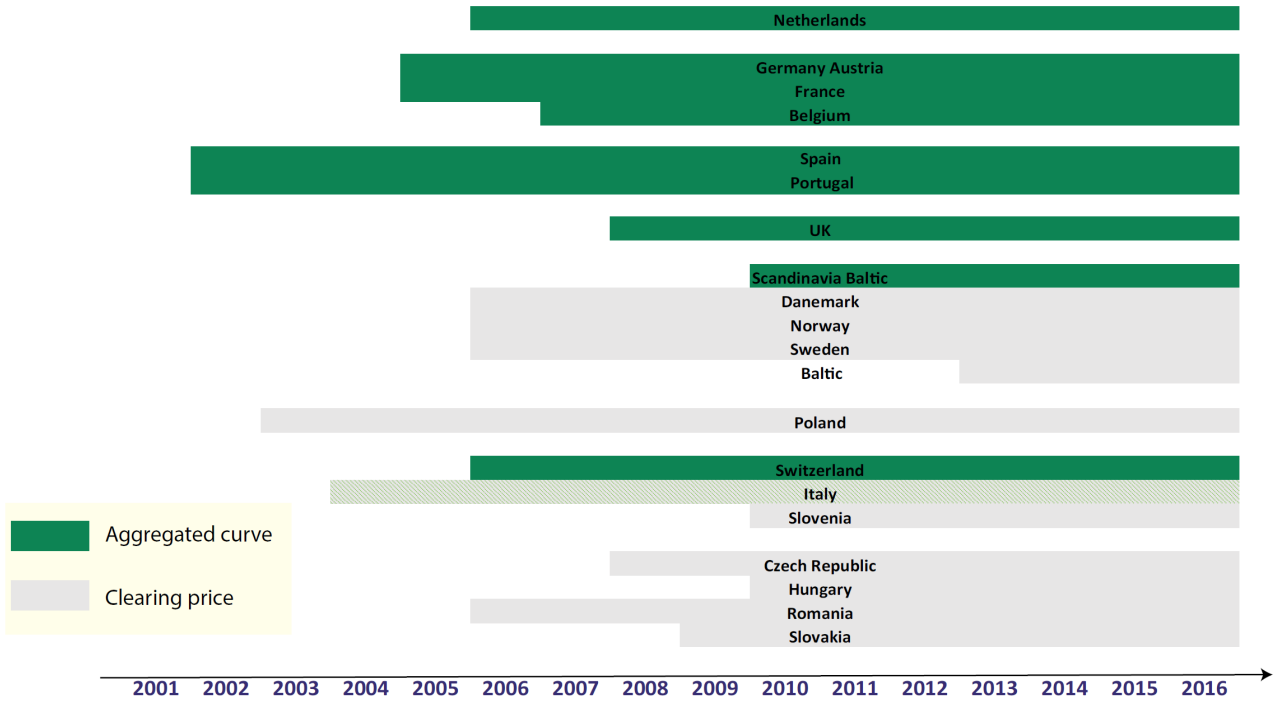

Countries are grouped by market coupling: Central Western Europe (DE, FR, BE), Southern Western Europe (PR,ES), 4M (CZ, HU, RO, SK), and Northern Europe (SE,NO,DK,FI,EE,LV,LT)

Figure A.2: German Demand Curve' normalization - January 1st 2006 at 5:00 pm

1.Raw data

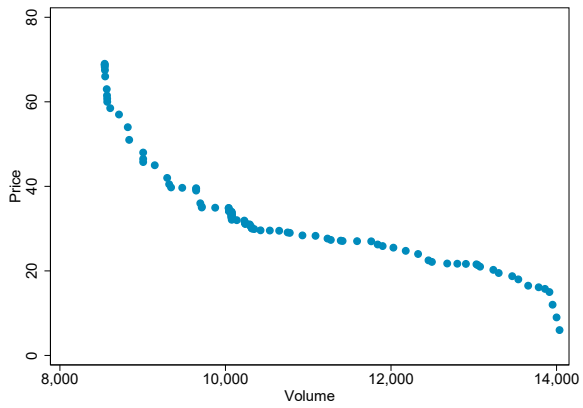

2.Linear interpolation

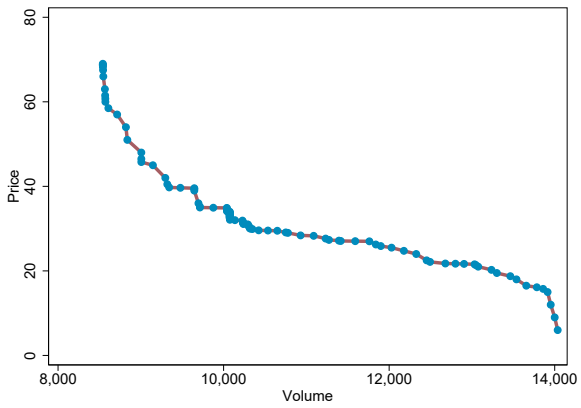

3.Uniform slicing

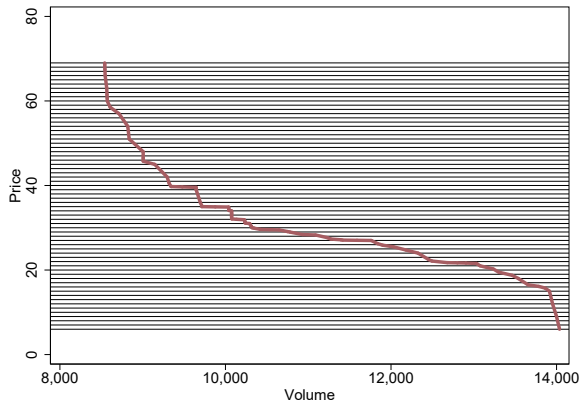

4.Final "input"

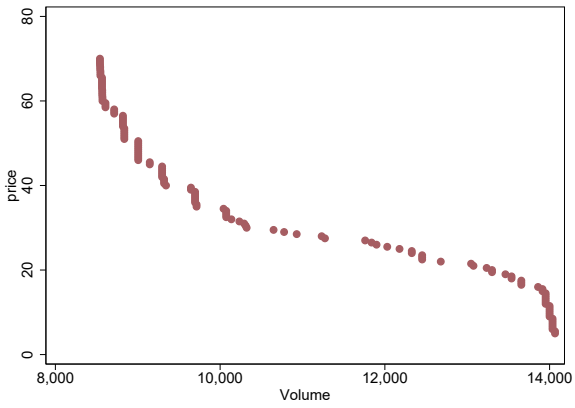

Figure A.3: Market coupling impact on interconnection usage

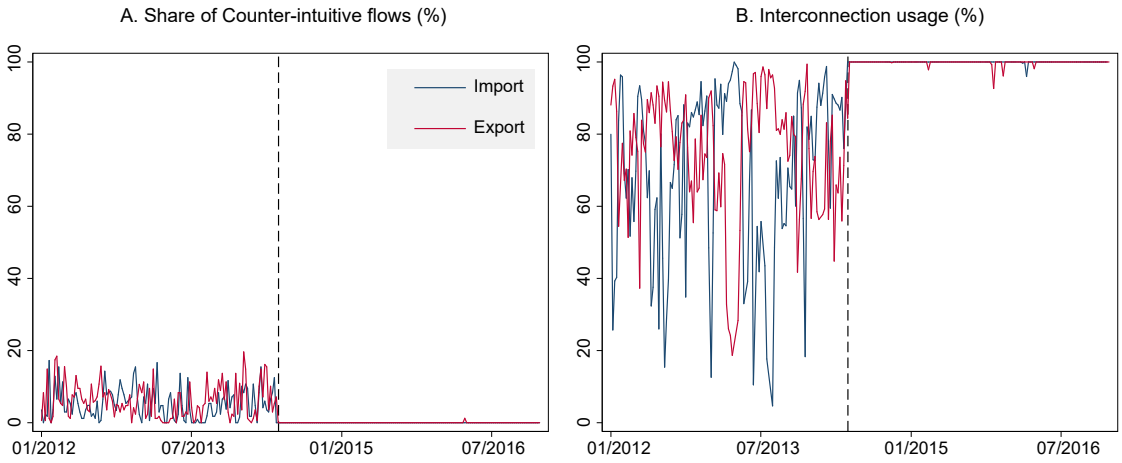

Hourly value aggregated at the weekly level. Panel A plots the percentage of hours with counter-intuitive flows, *ie* when the importer (exporter) has a lower (higher) price. Panel B represents the weekly average usage of the interconnection capacity whenever the price spread is different than zero. The vertical line is at the market coupling date.

Figure A.4: Monthly wind generation - Central Europe

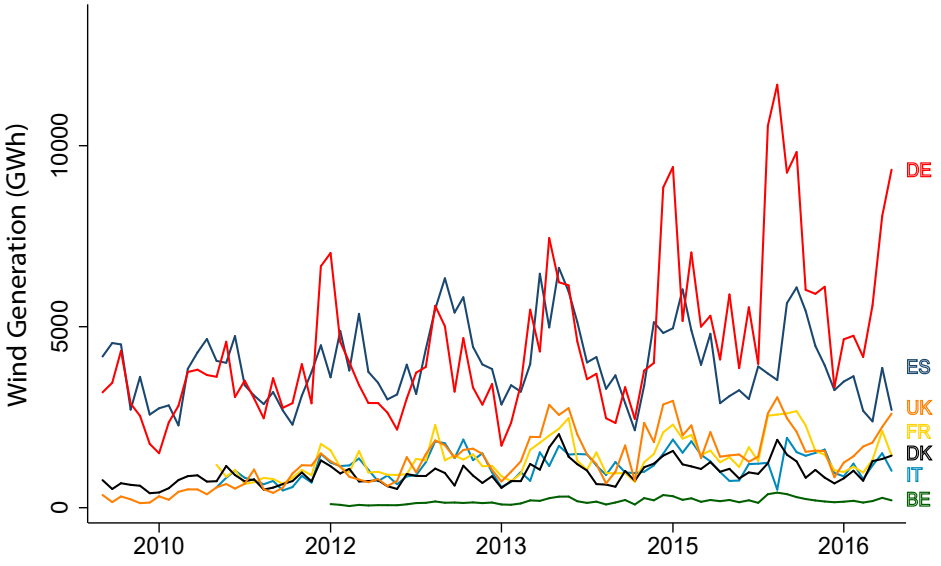

Monthly wind generation in treated countries and selected neighbors. Data source: country specific National Transmission System operators websites.

Figure A.5: Multiple dose-response

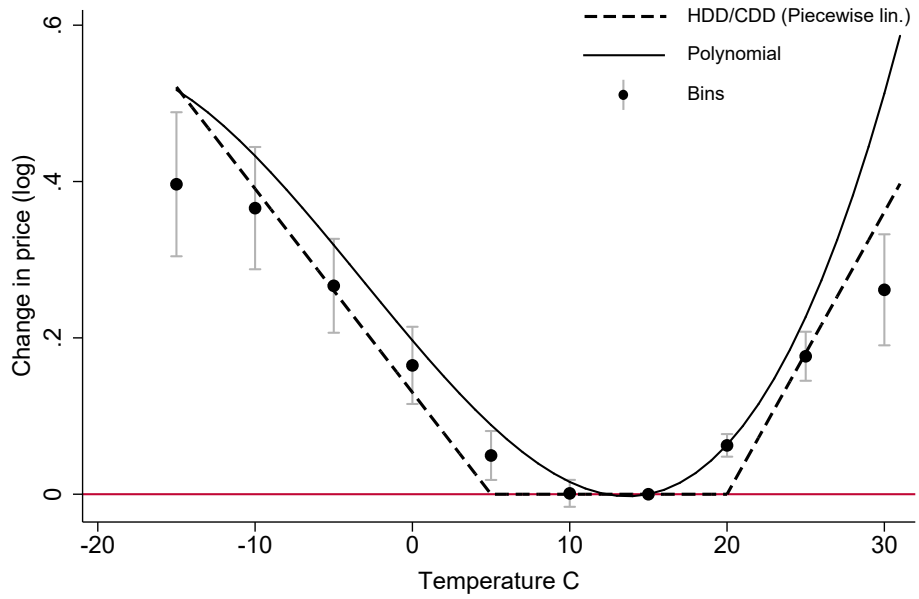

Figure A.6: Trends parameter along the dose-response

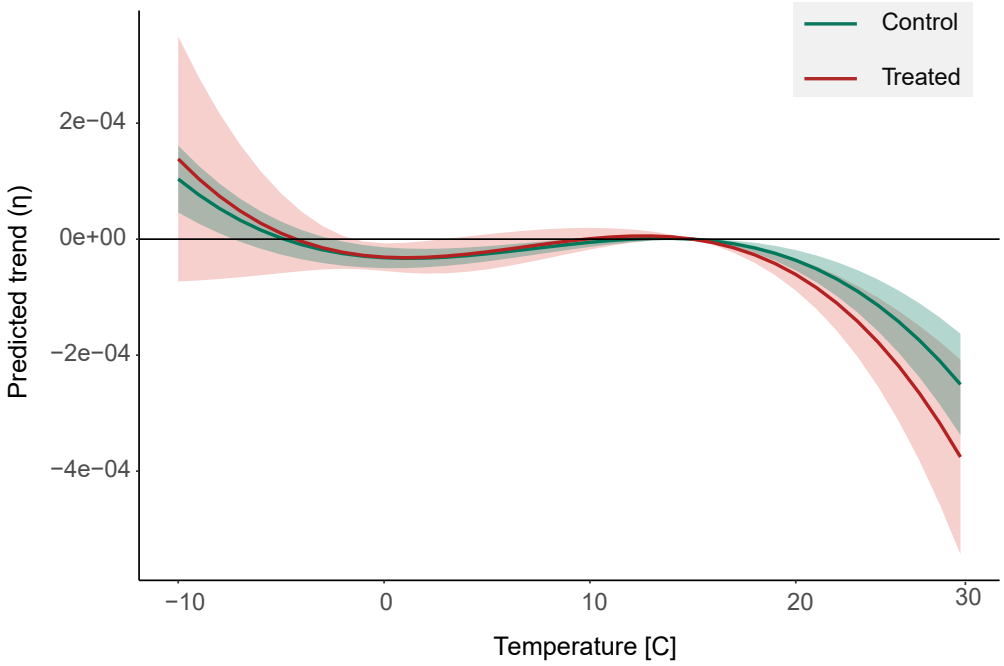

A.2 Hot Temperature, Aggression, and Death at the hands of the Police - Appendix

Data Appendix

European Centre for Medium-Range Weather Forecasts Data Climate variables are based on reanalysis data from the European Centre for Medium-Range Weather Forecasts (ERA-Interim), which is based on a climate model combined with observational data. We used their 0.25 x 0.25 gridded data on daily temperature and precipitation to generate aggregated daily temperature at the county level using population weights.[1] Population weights ensure to report an average of the temperature for places that matter for our study.

FBI Uniform Crime Reporting Data The crime data are obtained from the FBI Uniform Crime Reporting (UCR) data for offenses known and clearances by arrest from 1960 to 2016. The data are available monthly on a police department-level basis for index crimes (violent and property crimes) and assault on officers. For each county, we compute the monthly number of violent crimes (aggravated assault, murder, robbery, and rape), property crimes (burglary, larceny, motor vehicle theft), and officers by assault. We exclude agencies that report negative crimes in a given year. We also drop agencies than reported fewer than 12 months of data for a given year.

Robustness

Results from warm regions

We revisit our analysis from the previous section using U.S. counties where the annual average temperature is at least 19^oC. We do this for two reasons. First, because CEWs are less effective when used on individuals wearing thicker clothing (common during cold weather), it is important to account for the fact that officers are less likely to use their CEWs during cold weather. After 19^oC, people are less likely to wear thicker clothes (Morgan and de Dear (2003)). The behavior from civilians with respect to temperature is more likely to be held constant throughout the year. This helps mitigate the problem that officers are less likely to use CEWs because they anticipate that it would be less effective because subjects are more likely to wear thick clothing. Secondly, this section is intended to disentangle the effect of warm days from the effect of 'very hot' days. Figure A.9 presents the coefficients from estimating equation 2.3 on counties with an average annual temperature of at least 19^oC.

We found analogous results to the previous section. For warm regions in the U.S., we confirmed that temperature has no effect on civilian deaths caused by police shooting. The results remain precise. The number of civilian deaths by CEWs increased by about 6.8%

[1]For our polynomial precipitation form, it is important to take such nonlinear transformations at the cell level before spatially aggregating the data at the county level. This transformation preserves the tails of the distribution within the administrative region and ensure consistent results.

during extremely warm days, whereas the effect of temperature was close to zero and statistically non-significant for days where the temperature is less than 32^oC.

For physical restraints or less-than-lethal force (except CEWs), we found that the effect of temperature on civilian deaths is non-monotonic and very imprecise. The results for physical restraints during 'extremely warm' days were not robust using this subsample.

Effect of temperature on civilian deaths by symptom of mental illness or substance abuse

In order to assess the sensitivity of our results to abnormal behavior of civilians in warmer conditions, we exploited the fact that Fatal Encounters data provides information on whether the civilian exhibited symptoms of mental illness or substance abuse (drug or alcohol). Intuitively, a civilian who exhibits symptoms of a kind associated with mental illness or substance abuse would present behavior that officers find unpredictable, and therefore more threatening. One can also argue that officers might perceive individuals with these symptoms as more combative and/or less cooperative. In this case, substance use becomes an additional compounding factor, making it difficult to pin down the physiological impact of temperature on death by CEWs.

Figure 2.4 presents the coefficients from estimating equation 2.2 by symptom status. For both groups, there was a precise null impact of temperature on the number of deaths by firearm. This indicates that even if a civilian exhibits less predictable behavior, he or she is not more likely to die from a police shooting.

For CEWs, we confirmed the results from previous sections and obtained similar point estimates for the two groups. We show that the monthly number of civilian deaths by CEW increases by about 7.3% for any additional hot day compared to a day in the $12 - 17^oC$ range. The impact of temperature on civilian deaths by CEWs is null and non-significant for days with temperature lower than 32^oC. Despite the fact that substances might be an important factor that affect the number of deaths, it is surprising that both groups have similar point estimates for high temperatures. We view this finding as suggestive evidence that use of CEWs in high temperatures increases the odds of unintended death through some combination of physiological factors.

The significant impact of extremely warm days on the number of deaths by physical restraints seemed to be driven by civilians who exhibit symptom of mental illness or substance abuse. However, the confidence interval is fairly wide. For subjects that do not present symptoms, the temperature is null and not significant for very hot days. However, there does seem to be a marginally significant positive effect ($p<.1$) on the number of deaths of subjects without symptoms for temperatures between 22 and 27^oC.

As discussed and documented in previous sections, results for other less-than-lethal uses of force remain very imprecise. One cannot draw any conclusions about the effect of temperature on the number of civilian deaths for these types of police action.

Effect of temperature on the number of civilian deaths by vehicle

This section briefly presents the effect of temperature on the number of fatalities by vehicle when interacting with the police. Results in Table A.5 show that the number of civilian deaths by vehicle is not statistically influenced by temperature for days with a temperature of at least 2^oC. The coefficients are close to zero when the temperature is between 2 and 32^oC. The incident rate ratio is larger for warm days (1.02) but the effect is not statistically significant. The number of deaths by vehicle seems to be statistically smaller when the temperature is less than 2^oC. The increase during warmest days is consistent with previous studies on the impact of temperature on fatal traffic accident ((Leard et al., 2015)).

Additional tables and figures

Figure A.7: Police Fatal Encounters vs. Murders in the U.S.

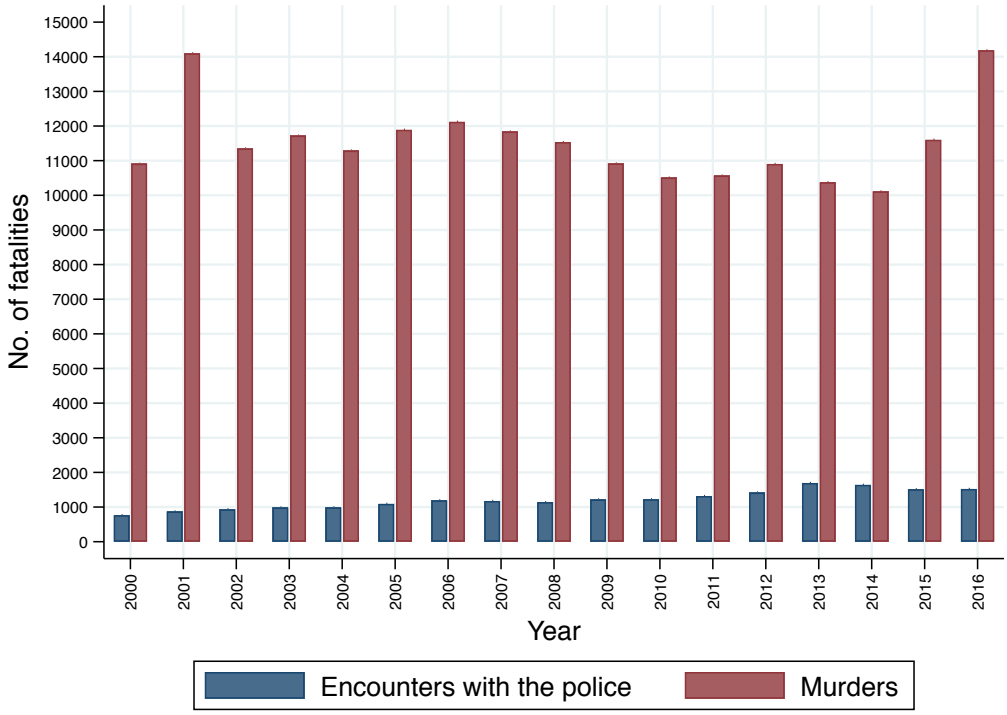

Notes: This figure presents the number of fatal encounters with the police and murders in the U.S. from 2000 to 2016. Sections 2.3 and A.2 provide details about the sample selection.

Figure A.8: Effect of temperature on the number of civilian deaths by firearm and CEWs

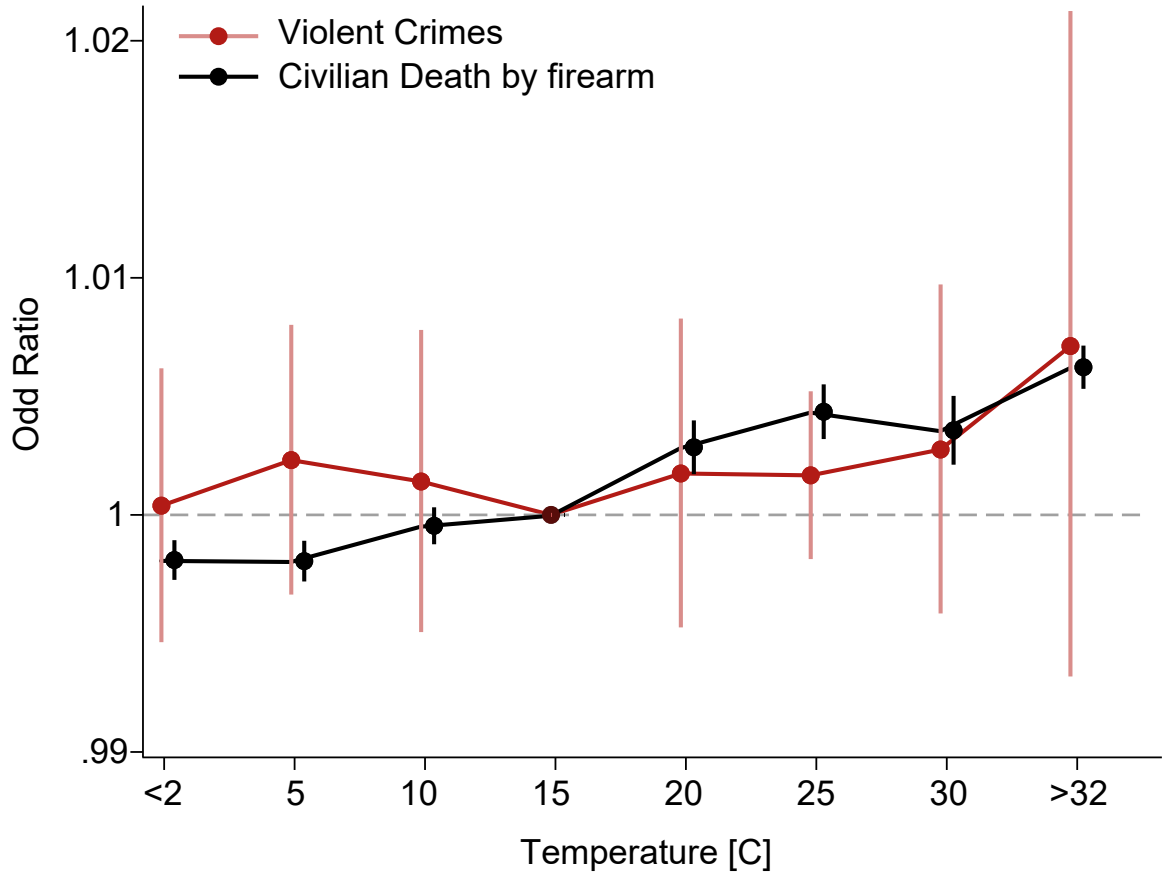

Notes: This figure presents the estimated effect of temperature on the number of civilian fatalities by CEWs (Taser) and firearm. The specification considers a 4th polynomial of temperature. The specification controls for precipitation, state-by-season fixed effect, and county-by-year fixed effect. For interpretation, we report the incident rate ratio on the y-axis. Standard errors are clustered at the year-by-type of death. We report the 95% CI. We report the midpoints for bins with temperature between 2 and 32^oC, with 5 degree increments.

Figure A.9: Effect of temperature on CEW and firearm fatal encounters in warm regions

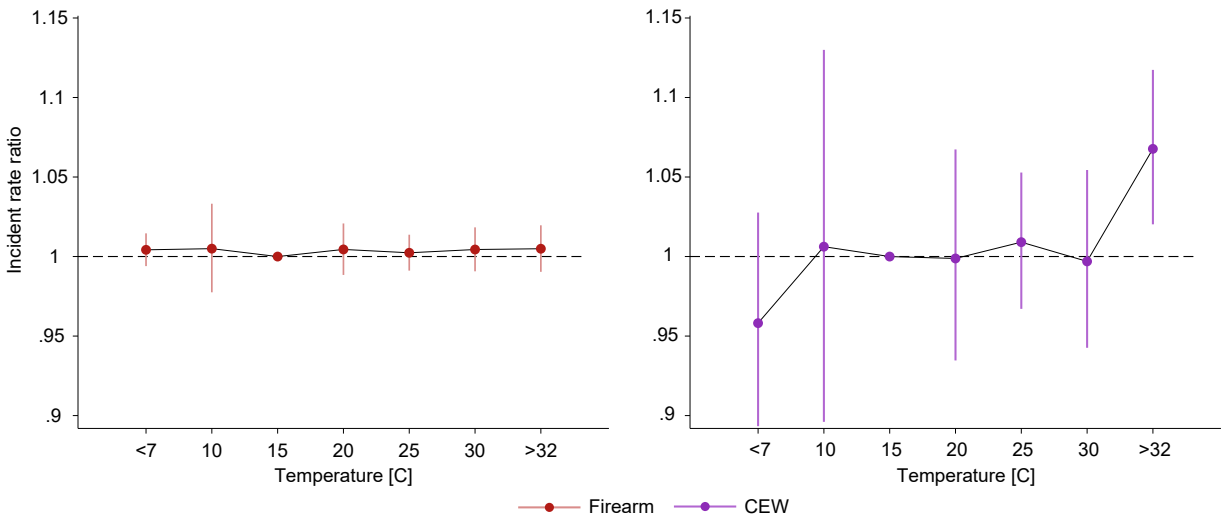

Notes: This figure presents the estimated effect of temperature on the number of civilian fatalities by firearm (left panel) and CEWs (right panel). for warm regions (average annual temperature of at least 19^oC). Section 2.2 provides details about the different causes of death. The specification controls for precipitation, state-by-season fixed effect, and county-by-year fixed effect. For interpretation, we report the incident rate ratio on the y-axis. Standard errors are clustered at the year-by-type of death. We report the 95% CI. We report the midpoints for bins with temperature between 7 and 32^oC, with 5 degree increments.

A.3 A distribution of Human Attention - Appendix

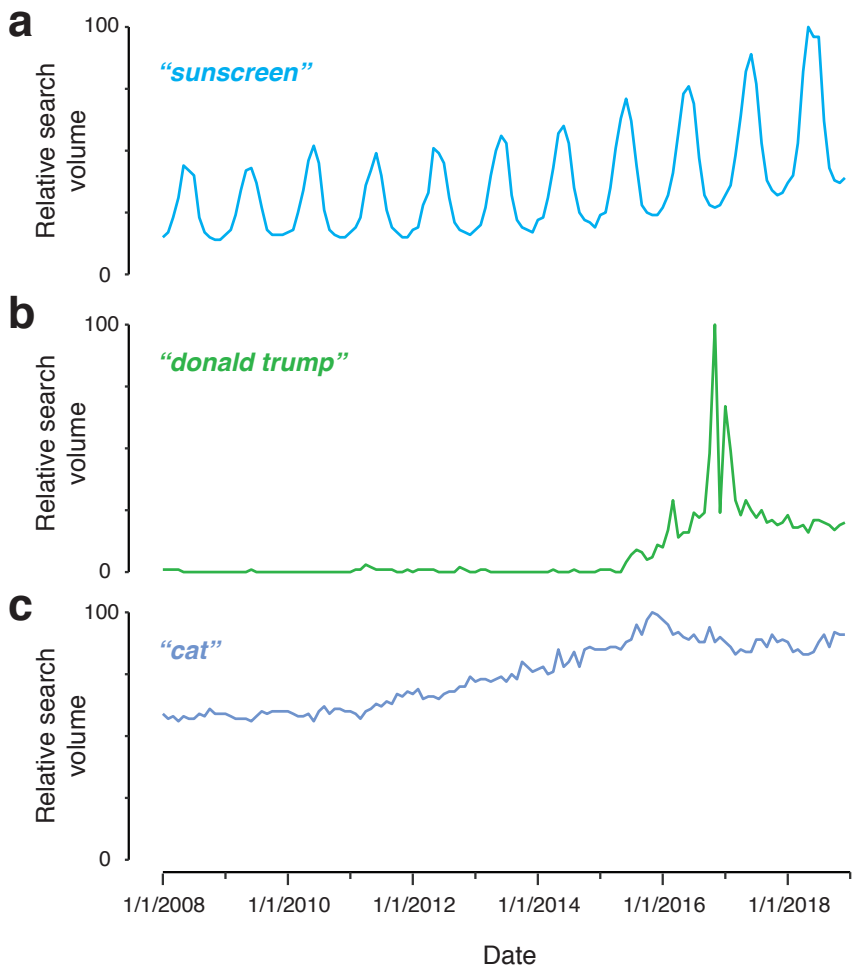

Figure A.10: ***Google Search* volume reflects changes in relative attention over time.** *Google Search* data records relative changes in the total number of queries that contain a specific target string. **(a)** Global queries for the topic *"sunscreen"* have a cyclical temporal structure that peaks during summer in the Northern Hemisphere. **(b)**. Global queries for the topic *"donald trump"* increased abruptly during the 2016 presidential campaign, which began in 2015. **(c)** Global queries for the topic *"cat"* are consistently high with little temporal structure. Raw data are rescaled by the data provider, indicating relative volumes of queries with a maximum value of 100. In 2019, Google reported receiving roughly 1.8B queries per day.

Table T1: Weather summary statistics per bidding zone

	Temperature				Precipitation		
	Mean	Std. dev.	Minimum	Maximum	Mean	Std. dev.	Maximum
Belgium	10.56	6.33	−8.82	27.06	2.17	3.23	33.29
Czech Rep	8.69	8.07	−16.65	27.51	1.82	2.72	31.49
Estonia	6.40	8.66	−22.36	26.43	1.86	2.64	21.68
Finland	4.96	9.00	−24.93	24.96	1.91	2.48	23.52
France	11.40	6.21	−6.09	27.71	2.13	2.53	18.19
Ger Aus Lux	9.51	7.20	−12.16	27.15	2.17	2.39	22.16
Hungary	11.08	8.79	−13.01	30.89	1.58	2.93	35.30
Latvia	7.11	8.72	−23.70	26.34	1.94	2.76	34.07
Lithuania	7.53	8.76	−23.52	26.84	1.93	2.81	37.51
Netherlands	10.60	6.19	−9.79	26.57	2.21	3.14	27.09
Norway	5.10	6.93	−17.54	20.51	3.22	3.13	24.51
Poland	8.84	8.26	−20.33	27.86	1.79	2.28	26.15
Portugal	14.99	5.12	2.28	29.39	2.11	4.40	55.23
Romania	10.33	8.92	−15.17	28.83	1.84	2.68	26.46
Slovakia	9.15	8.48	−15.30	28.63	1.78	2.91	39.95
Slovenia	9.91	8.07	−12.73	28.27	2.72	5.00	46.55
Spain	14.38	6.24	−0.17	28.55	1.42	2.06	20.74
Switzerland	7.37	7.37	−15.26	24.77	3.25	4.87	48.99
UK	9.88	4.93	−4.70	22.61	2.21	2.71	22.85
Sweden							
SE1	2.66	9.14	−27.52	22.95	1.80	2.85	30.78
SE2	4.52	8.13	−21.36	22.94	1.79	2.60	31.51
SE3	7.15	7.54	−15.84	23.01	1.97	2.74	25.63
SE4	8.07	6.95	−13.76	23.90	2.07	3.00	40.83
Italy							
Center South	14.60	6.60	−0.86	28.51	2.07	3.74	47.33
Center North	13.66	7.03	−4.10	28.17	2.36	3.98	49.47
North	10.88	7.45	−7.92	26.89	2.61	4.55	46.11
Sardaignia	16.62	5.60	2.88	29.49	1.27	3.00	41.48
Sicilia	17.29	5.69	2.58	31.12	1.39	3.02	43.37
South	16.35	6.27	1.22	31.21	1.79	3.54	51.65
Denmark							
East	8.92	6.62	−12.18	24.10	2.00	3.16	38.61
West	9.10	6.20	−8.64	22.75	2.09	3.11	35.15

Temperature is defined as the daily mean surface temperature in Celsius, Precipitation is expressed in daily total value (millimeter)

Table A.1: Effect of temperature on the number of civilian deaths by physical restraints

	(1)	(2)	(3)	(4)	(5)	(6)
$< 2^o$	0.00788	0.0156	0.0155	0.0212	0.0329*	0.0329*
	(0.0177)	(0.0171)	(0.0171)	(0.0170)	(0.0175)	(0.0175)
$[2^o, 7^o)$	0.0321*	0.0374**	0.0369**	0.0380**	0.0354*	0.0352*
	(0.0165)	(0.0164)	(0.0164)	(0.0167)	(0.0190)	(0.0190)
$[7^o, 12^o)$	0.0215	0.0217	0.0215	0.0222	0.0273*	0.0272*
	(0.0171)	(0.0168)	(0.0168)	(0.0166)	(0.0159)	(0.0159)
$[17^o, 22^o)$	0.0289	0.0380*	0.0379*	0.0360*	0.0371*	0.0371*
	(0.0231)	(0.0216)	(0.0216)	(0.0214)	(0.0213)	(0.0213)
$[22^o, 27^o)$	0.0267*	0.0235	0.0224	0.0206	0.0230	0.0226
	(0.0154)	(0.0165)	(0.0164)	(0.0161)	(0.0160)	(0.0160)
$[27^o, 32^o)$	0.0339	0.0484**	0.0473**	0.0460**	0.0472**	0.0469**
	(0.0219)	(0.0207)	(0.0208)	(0.0207)	(0.0210)	(0.0210)
$[> 32^o$	0.156**	0.154**	0.151**	0.153**	0.141*	0.141*
	(0.0678)	(0.0723)	(0.0726)	(0.0732)	(0.0734)	(0.0734)
Precipitation	0.000299	0.000225	0.000235	0.000221	0.000306	0.000315
	(0.000331)	(0.000366)	(0.000363)	(0.000364)	(0.000362)	(0.000363)
Precipitation^2	-1.46e-05	-1.48e-05	-1.50e-05	-1.63e-05	-1.72e-05*	-1.74e-05*
	(8.93e-06)	(1.03e-05)	(1.02e-05)	(1.03e-05)	(1.00e-05)	(1.01e-05)
Violent Crimes		-5.59e-05	-3.55e-05			
		(4.94e-05)	(5.18e-05)			
Property Crimes		4.77e-05	-0.000163**			
		(8.46e-05)	(7.90e-05)			
No. Officers assulted/killed			0.00536***			0.00164
			(0.000611)			(0.00223)
Constant	-1.427***	-1.373***	-1.392***	-8.476***	-7.175***	-7.206***
	(0.0550)	(0.0651)	(0.0593)	(0.0611)	(0.0640)	(0.0680)
Exposure	-	-	-	total crime	total arrest	total arrest
Fixed-effects:						
Season-State	YES	YES	YES	YES	YES	YES
County-Year-Type of death	YES	YES	YES	YES	YES	YES
Observations	128,808	112,036	112,036	111,286	106,226	106,226

Notes: This Table presents the estimated effect of temperature on the number of civilian fatalities by physical restraints using equation 2.2. Section 2.2 provides details about the different causes of death. Standard errors are clustered at the year-by-type of death.

Table A.2: Effect of temperature on the number of civilian deaths by less-than-lethal force

	(1)	(2)	(3)	(4)	(5)	(6)
$< 2^o$	-0.00702	0.00421	0.00401	0.00751	0.00746	0.00737
	(0.0168)	(0.0159)	(0.0159)	(0.0154)	(0.0176)	(0.0176)
$[2^o, 7^o)$	-0.0145	-0.0120	-0.0126	-0.0154	-0.0104	-0.0106
	(0.0231)	(0.0272)	(0.0273)	(0.0281)	(0.0264)	(0.0264)
$[7^o, 12^o)$	0.0472*	0.0591**	0.0584**	0.0565**	0.0605**	0.0602**
	(0.0243)	(0.0247)	(0.0246)	(0.0235)	(0.0264)	(0.0264)
$[17^o, 22^o)$	0.0272	0.0447**	0.0446**	0.0407*	0.0481**	0.0480**
	(0.0224)	(0.0214)	(0.0213)	(0.0209)	(0.0229)	(0.0228)
$[22^o, 27^o)$	0.0208	0.0241	0.0227	0.0198	0.0236	0.0232
	(0.0140)	(0.0150)	(0.0150)	(0.0146)	(0.0155)	(0.0156)
$[27^o, 32^o)$	0.00433	0.0177	0.0163	0.0149	0.0203	0.0200
	(0.0225)	(0.0219)	(0.0221)	(0.0221)	(0.0229)	(0.0229)
$[> 32^o$	-0.00486	0.0140	0.0154	0.00972	0.00808	0.00844
	(0.0617)	(0.0573)	(0.0567)	(0.0541)	(0.0549)	(0.0547)
Precipitation	0.000299	0.000225	0.000235	0.000221	0.000306	0.000315
	(0.000331)	(0.000366)	(0.000363)	(0.000364)	(0.000362)	(0.000363)
Precipitation^2	-1.46e-05	-1.48e-05	-1.50e-05	-1.63e-05	-1.72e-05*	-1.74e-05*
	(8.93e-06)	(1.03e-05)	(1.02e-05)	(1.03e-05)	(1.00e-05)	(1.01e-05)
Violent Crimes		-5.59e-05	-3.55e-05			
		(4.94e-05)	(5.18e-05)			
Property Crimes		4.77e-05	-0.000163**			
		(8.46e-05)	(7.90e-05)			
No. Officers assulted/killed			0.00536***			0.00164
			(0.000611)			(0.00223)
Constant	-1.427***	-1.373***	-1.392***	-8.476***	-7.175***	-7.206***
	(0.0550)	(0.0651)	(0.0593)	(0.0611)	(0.0640)	(0.0680)
Exposure	-	-	-	total crime	total arrest	total arrest
Fixed-effects:						
Season-State	YES	YES	YES	YES	YES	YES
County-Year-Type of death	YES	YES	YES	YES	YES	YES
Observations	128,808	112,036	112,036	111,286	106,226	106,226

Notes: This Table presents the estimated effect of temperature on the number of civilian fatalities by less-than-lethal force using equation 2.2. Section 2.2 provides details about the different causes of death. Standard errors are clustered at the year-by-type of death.

Table A.3: Effect of temperature on the number of civilian deaths by firearm

	(1)	(2)	(3)	(4)	(5)	(6)
$< 2^o$	0.000390	-0.000189	-0.000334	0.00417	0.00296	0.00288
	(0.00295)	(0.00313)	(0.00313)	(0.00312)	(0.00336)	(0.00338)
$[2^o, 7^o)$	0.00231	0.00284	0.00247	0.00217	0.00306	0.00292
	(0.00290)	(0.00312)	(0.00310)	(0.00324)	(0.00324)	(0.00325)
$[7^o, 12^o)$	0.00141	0.000664	0.000562	-2.08e-05	-0.000404	-0.000513
	(0.00325)	(0.00341)	(0.00346)	(0.00350)	(0.00370)	(0.00374)
$[17^o, 22^o)$	0.00174	0.00151	0.00121	-0.00176	-0.000943	-0.00109
	(0.00332)	(0.00353)	(0.00346)	(0.00370)	(0.00395)	(0.00400)
$[22^o, 27^o)$	0.00167	0.00179	0.000985	-0.00215	-0.00192	-0.00215
	(0.00180)	(0.00198)	(0.00201)	(0.00194)	(0.00193)	(0.00190)
$[27^o, 32^o)$	0.00275	0.00336	0.00243	0.000496	0.000982	0.000745
	(0.00353)	(0.00427)	(0.00418)	(0.00475)	(0.00503)	(0.00506)
$[> 32^o$	0.00710	0.00571	0.00451	0.00230	0.00181	0.00150
	(0.00711)	(0.00673)	(0.00664)	(0.00701)	(0.00655)	(0.00655)
Precipitation	0.000299	0.000225	0.000235	0.000221	0.000306	0.000315
	(0.000331)	(0.000366)	(0.000363)	(0.000364)	(0.000362)	(0.000363)
Precipitation^2	-1.46e-05	-1.48e-05	-1.50e-05	-1.63e-05	-1.72e-05*	-1.74e-05*
	(8.93e-06)	(1.03e-05)	(1.02e-05)	(1.03e-05)	(1.00e-05)	(1.01e-05)
Violent Crimes		-5.59e-05	-3.55e-05			
		(4.94e-05)	(5.18e-05)			
Property Crimes		4.77e-05	-0.000163**			
		(8.46e-05)	(7.90e-05)			
No. Officers assulted/killed			0.00536***			0.00164
			(0.000611)			(0.00223)
Constant	-1.427***	-1.373***	-1.392***	-8.476***	-7.175***	-7.206***
	(0.0550)	(0.0651)	(0.0593)	(0.0611)	(0.0640)	(0.0680)
Exposure	-	-	-	total crime	total arrest	total arrest
Fixed-effects:						
Season-State	YES	YES	YES	YES	YES	YES
County-Year-Type of death	YES	YES	YES	YES	YES	YES
Observations	128,808	112,036	112,036	111,286	106,226	106,226

Notes: This Table presents the estimated effect of temperature on the number of civilian fatalities by firearm using equation 2.2. Section 2.2 provides details about the different causes of death. Standard errors are clustered at the year-by-type of death.

Table A.4: Effect of temperature on the number of civilian deaths by CEW

	(1)	(2)	(3)	(4)	(5)	(6)
$< 2^o$	-0.0227*	-0.0275*	-0.0277*	-0.0225	-0.0275*	-0.0276*
	(0.0126)	(0.0146)	(0.0146)	(0.0142)	(0.0151)	(0.0151)
$[2^o, 7^o)$	-0.0103	-0.0108	-0.0110	-0.0124	-0.0107	-0.0109
	(0.0117)	(0.0127)	(0.0127)	(0.0132)	(0.0137)	(0.0137)
$[7^o, 12^o)$	0.00725	0.00561	0.00531	0.00680	0.00794	0.00777
	(0.0151)	(0.0153)	(0.0153)	(0.0151)	(0.0150)	(0.0150)
$[17^o, 22^o)$	0.00183	-0.000188	-0.000488	-0.00350	0.000236	4.05e-05
	(0.0129)	(0.0140)	(0.0139)	(0.0142)	(0.0141)	(0.0141)
$[22^o, 27^o)$	0.00963	0.00847	0.00784	0.00560	0.00513	0.00495
	(0.00632)	(0.00774)	(0.00773)	(0.00778)	(0.00853)	(0.00852)
$[27^o, 32^o)$	0.00113	0.00593	0.00505	0.00431	0.00532	0.00506
	(0.0111)	(0.0119)	(0.0118)	(0.0120)	(0.0123)	(0.0123)
$[> 32^o$	0.0563***	0.0529***	0.0512***	0.0499***	0.0517***	0.0512***
	(0.0105)	(0.0117)	(0.0117)	(0.0120)	(0.0124)	(0.0124)
Precipitation	0.000299	0.000225	0.000235	0.000221	0.000306	0.000315
	(0.000331)	(0.000366)	(0.000363)	(0.000364)	(0.000362)	(0.000363)
Precipitation^2	-1.46e-05	-1.48e-05	-1.50e-05	-1.63e-05	-1.72e-05*	-1.74e-05*
	(8.93e-06)	(1.03e-05)	(1.02e-05)	(1.03e-05)	(1.00e-05)	(1.01e-05)
Violent Crimes		-5.59e-05	-3.55e-05			
		(4.94e-05)	(5.18e-05)			
Property Crimes		4.77e-05	-0.000163**			
		(8.46e-05)	(7.90e-05)			
No. Officers assulted/killed			0.00536***			0.00164
			(0.000611)			(0.00223)
Constant	-1.427***	-1.373***	-1.392***	-8.476***	-7.175***	-7.206***
	(0.0550)	(0.0651)	(0.0593)	(0.0611)	(0.0640)	(0.0680)
Exposure	-	-	-	total crime	total arrest	total arrest
Fixed-effects:						
Season-State	YES	YES	YES	YES	YES	YES
County-Year-Type of death	YES	YES	YES	YES	YES	YES
Observations	128,808	112,036	112,036	111,286	106,226	106,226

Notes: This Table presents the estimated effect of temperature on the number of civilian fatalities by CEW using equation 2.2. Section 2.2 provides details about the different causes of death. Standard errors are clustered at the year-by-type of death.

Table A.5: Effect of temperature on the number of civilian deaths by vehicle

	(1)	(2)	(3)	(4)	(5)	(6)
$< 2^o$	-0.0150***	-0.0165***	-0.0166***	-0.0128**	-0.0152***	-0.0153***
	(0.00532)	(0.00553)	(0.00551)	(0.00574)	(0.00576)	(0.00576)
$[2^o, 7^o)$	0.0108*	0.0122**	0.0120**	0.00999*	0.0126**	0.0125**
	(0.00621)	(0.00575)	(0.00572)	(0.00597)	(0.00630)	(0.00630)
$[7^o, 12^o)$	-0.00267	-0.00368	-0.00381	-0.00420	-0.00345	-0.00355
	(0.00737)	(0.00785)	(0.00786)	(0.00810)	(0.00781)	(0.00784)
$[17^o, 22^o)$	-0.00398	-0.00540	-0.00557	-0.0101	-0.00804	-0.00815
	(0.00592)	(0.00645)	(0.00652)	(0.00677)	(0.00683)	(0.00686)
$[22^o, 27^o)$	0.000521	0.00231	0.00164	-0.00124	-0.00111	-0.00130
	(0.00455)	(0.00446)	(0.00445)	(0.00430)	(0.00457)	(0.00456)
$[27^o, 32^o)$	-0.00764	-0.00834	-0.00911	-0.0124	-0.0104	-0.0106
	(0.00854)	(0.00913)	(0.00915)	(0.00967)	(0.00921)	(0.00922)
$[> 32^o$	0.0313	0.0318	0.0312	0.0287	0.0170	0.0168
	(0.0194)	(0.0195)	(0.0194)	(0.0189)	(0.0169)	(0.0169)
Precipitation	0.000299	0.000225	0.000235	0.000221	0.000306	0.000315
	(0.000331)	(0.000366)	(0.000363)	(0.000364)	(0.000362)	(0.000363)
Precipitation^2	-1.46e-05	-1.48e-05	-1.50e-05	-1.63e-05	-1.72e-05*	-1.74e-05*
	(8.93e-06)	(1.03e-05)	(1.02e-05)	(1.03e-05)	(1.00e-05)	(1.01e-05)
Violent Crimes		-5.59e-05	-3.55e-05			
		(4.94e-05)	(5.18e-05)			
Property Crimes		4.77e-05	-0.000163**			
		(8.46e-05)	(7.90e-05)			
No. Officers assulted/killed			0.00536***			0.00164
			(0.000611)			(0.00223)
Constant	-1.427***	-1.373***	-1.392***	-8.476***	-7.175***	-7.206***
	(0.0550)	(0.0651)	(0.0593)	(0.0611)	(0.0640)	(0.0680)
Exposure	-	-	-	total crime	total arrest	total arrest
Fixed-effects:						
Season-State	YES	YES	YES	YES	YES	YES
County-Year-Type of death	YES	YES	YES	YES	YES	YES
Observations	128,808	112,036	112,036	111,286	106,226	106,226

Notes: This Table presents the estimated effect of temperature on the number of civilian fatalities by vehicle using equation 2.2. Section 2.2 provides details about the different causes of death. Standard errors are clustered at the year-by-type of death.

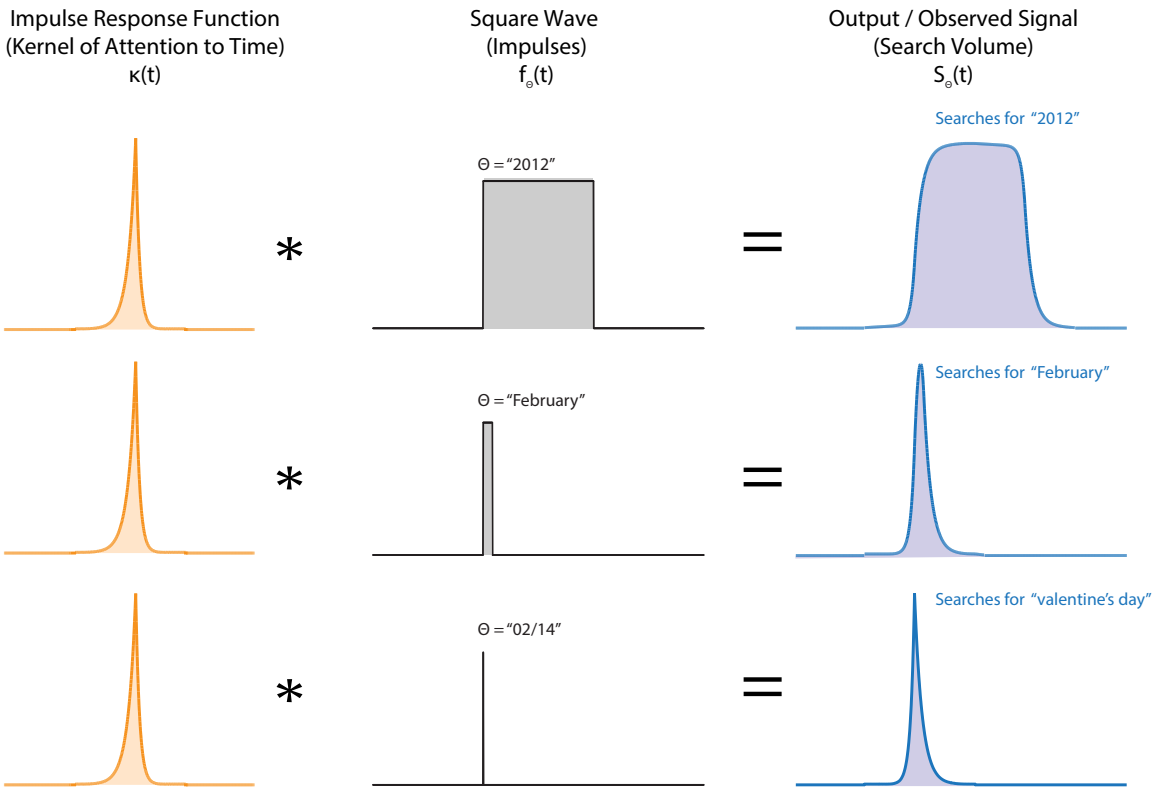

Figure A.11: **Prediction that *Google Search* volume reflects a convolution of the KAT with a square wave**. The KAT-based model of attention to time predicts that *Google Search* volume ($S_\Theta(t)$, blue) should reflect the convolution of the KAT ($\kappa(t)$, orange) with a sequence of impulses (black), where moments of time (θ) contained within the target interval (Θ) — e.g. days within the target year "2016" — are each represented by a separate impulse (see Equation 3.6). Because moments in time contained within the target interval are sequential, the discreet form of the input function describing the sequence of impulses ($f_\Theta(t)$) appears as a square wave. Figure 3.1**d-i** in the main text illustrate how the KAT-based model generates this result. In this analysis, we observe *Google Search* volume ($S_\Theta(t)$) and construct input functions ($f_\Theta(t)$) that reflect intervals of time described by search query targets. We then implement a deconvolution (the inverse of the convolution) to recover the estimated KAT ($\hat{\kappa}(t)$).

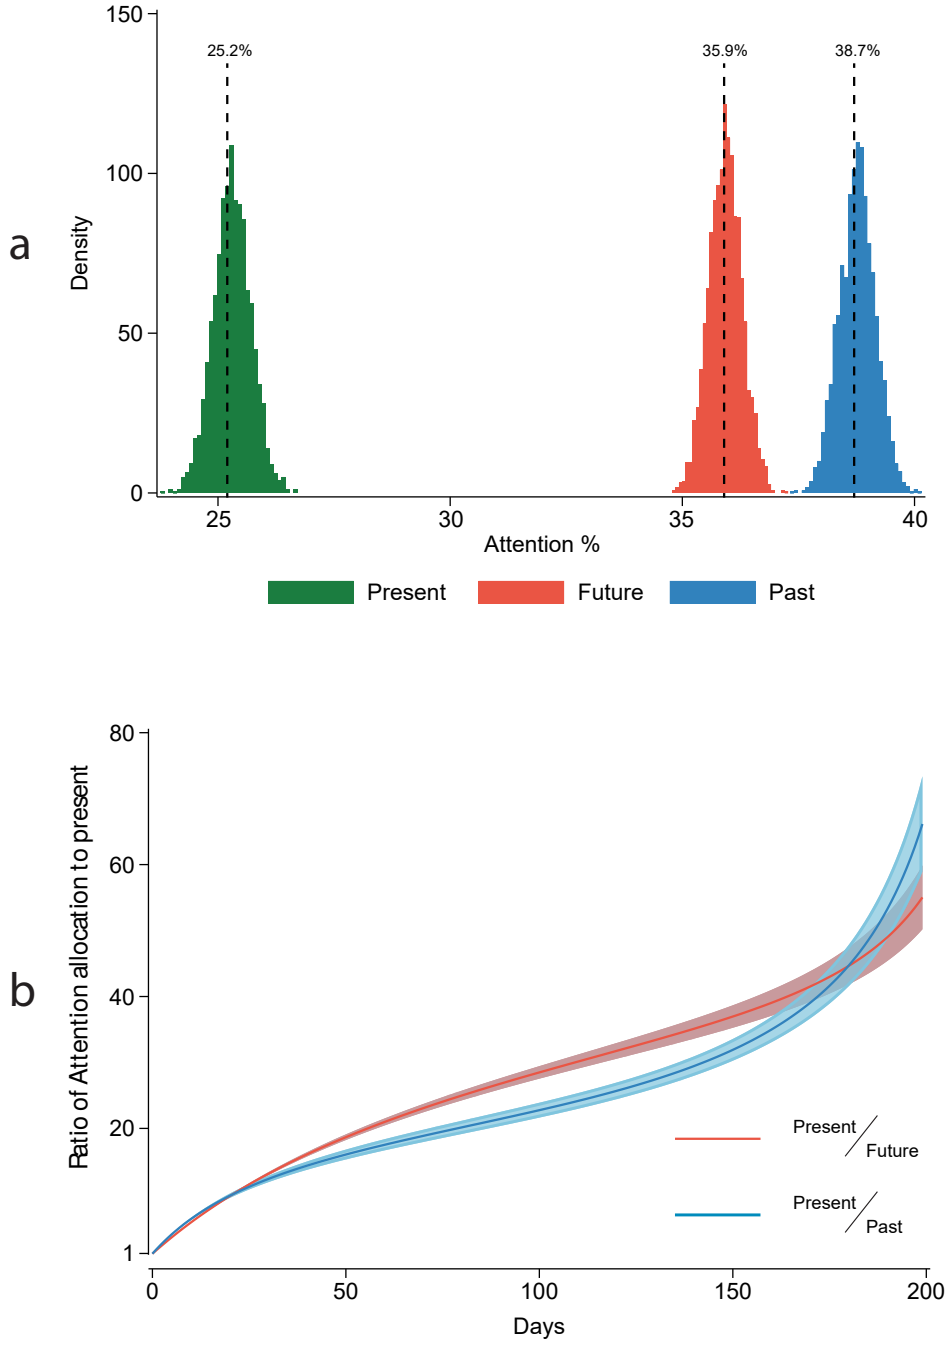

Figure A.12: **Uncertainty in the estimated structure of the KAT.** We re-estimate the KAT (using annual targets) multiple times using 3,000 block-resampled versions of the data (with replacement), blocking at the query-by-location level. **(a)** Histograms indicate the distribution of estimates for the fraction of total attention allocated to the past (blue), present (green) and future (red). **(b)** The ratio of total attention allocated to the present vs each day in the relative past (blue curve) and each day in the relative future (red curve). Attention to one day in the past or future is most similar to attention to the present, so values are near one. For more distant days in relative time, much more attention is allocated to the present than to a distant day (e.g. 40-70× more attention). Lines are central estimates for these ratios, bands are 95% confidence intervals using block-boostrapped estimates.

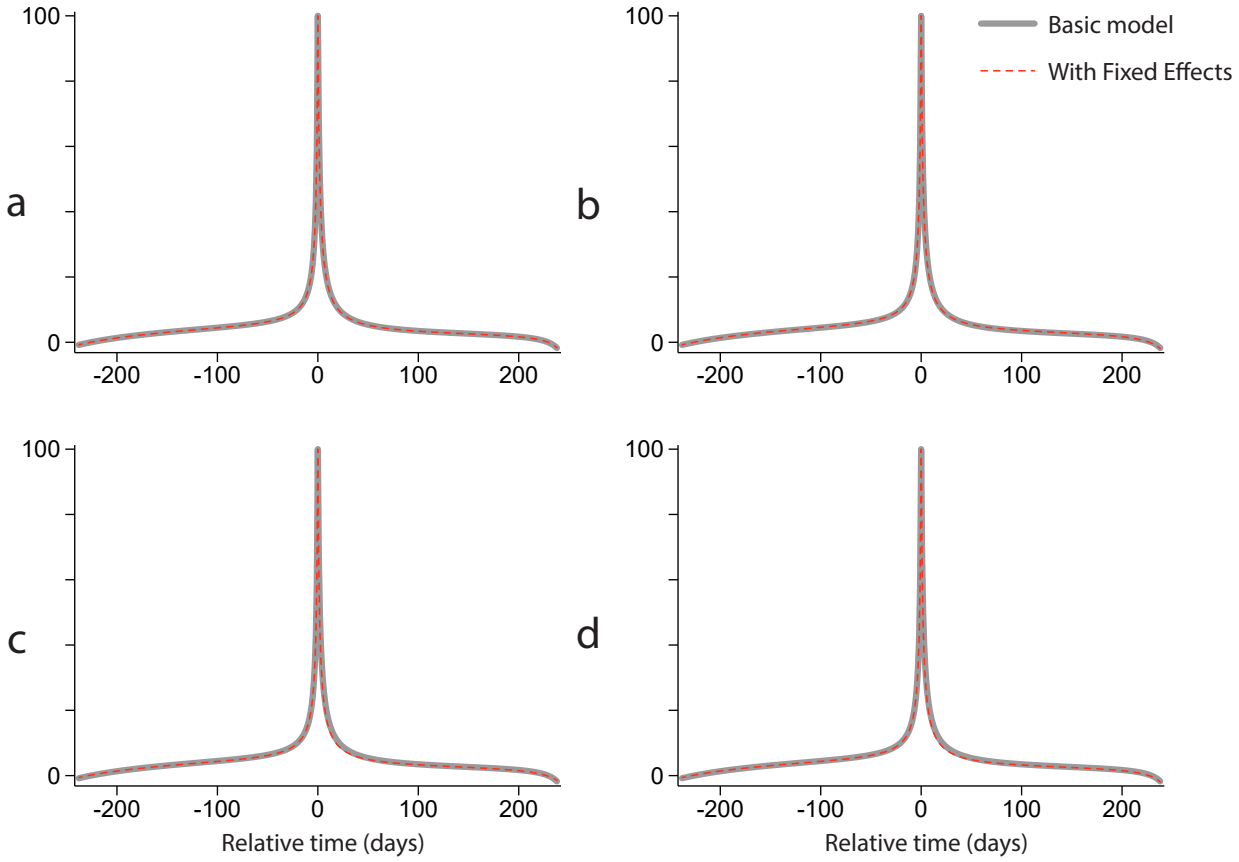

Figure A.13: **Sensitivity analysis for the estimated KAT using alternative regression model specifications.** The red dashed line in each panel represents a KAT estimated using a different set of fixed-effects (i.e. intercepts) that account for possible omitted variable bias during estimation of Equation 3.7. The solid gray line in each panel is the main estimate plotted in Figure 3.3, which is the most parsimonious model without fixed-effects. In each panel, the dashed red line is an alternative specification presented for comparison against this benchmark. **(a)** Including country-level fixed-effects that account for country-specific characteristics that remain constant over time (e.g. higher average query volume). **(b)** Including country-level and year-level fixed-effects to account for country-specific characteristics that remain constant over time and for flexible non-parametric time trends that affect all countries (e.g. increasing search volume over time). **(c)** Including country, year, and query-level fixed effects to account for constant country-specific unobserved characteristics, time trends, and query-specific unobserved characteristics (e.g. higher search volume for years associated with the Olympic games). **(d)** Including country-by-year level fixed effects to non-parametrically account for unobserved variation at the country level within any given year. None of these adjustments substantively alter the fit or structure of the results relative to the main model.

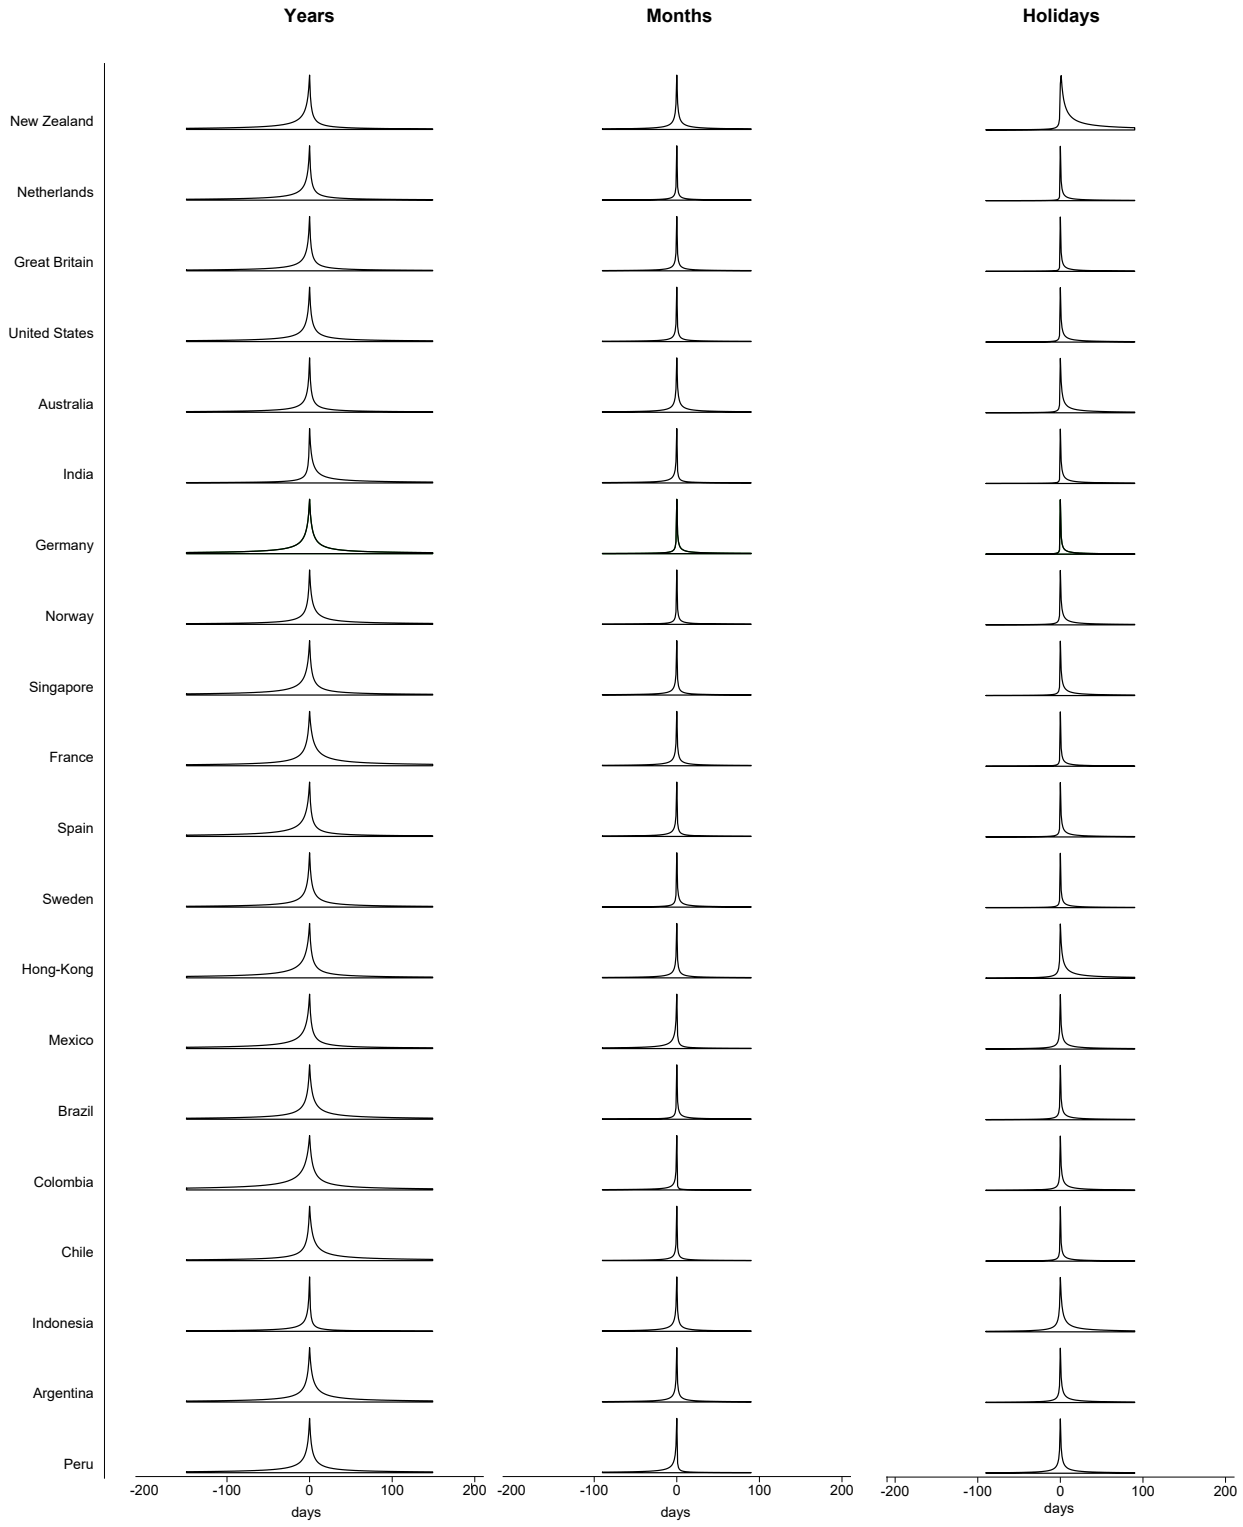

Figure A.14: **Generalizability and reproducibility of the KAT.** Estimated form of the KAT using different classes of target for query and different samples. Left, center and right columns show KAT estimates using annual, monthly, and day targets, respectively. Rows indicate the country used to generate each estimate. Comparable pooled estimates are shown in Figure 3.3d-f.

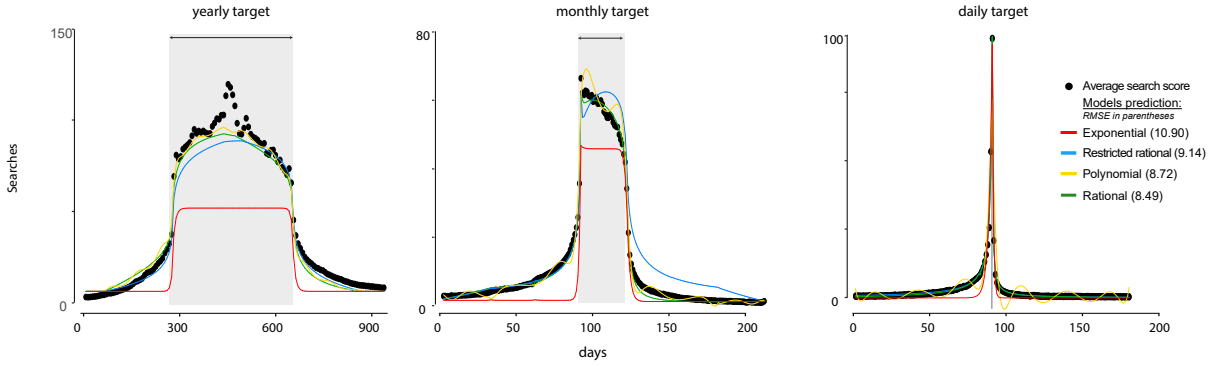

Figure A.15: **Predicted search volume for annual, monthly, and day targets using estimated KATs.** Panels show the predicted search volume generated by convolving estimated KAT functions with corresponding input functions $f_\Theta(t)$ (the process depicted in Figure 3.1d-i and ED Figure A.11). Panels reflect yearly (left), monthly (center), and daily (right) targets and utilized smooth KAT estimates shown in Figure 3.3**c-e**: exponential (red), hyperbolic-discounting-like ($\frac{a}{1+b(\theta-t)}$, blue), eighth-order polynomial (yellow), and rational functions (green). The black dots correspond to the average query volume observed in actual *Google Search* data. The RMSE (weighted by target) associated with each functional form is shown in parentheses.

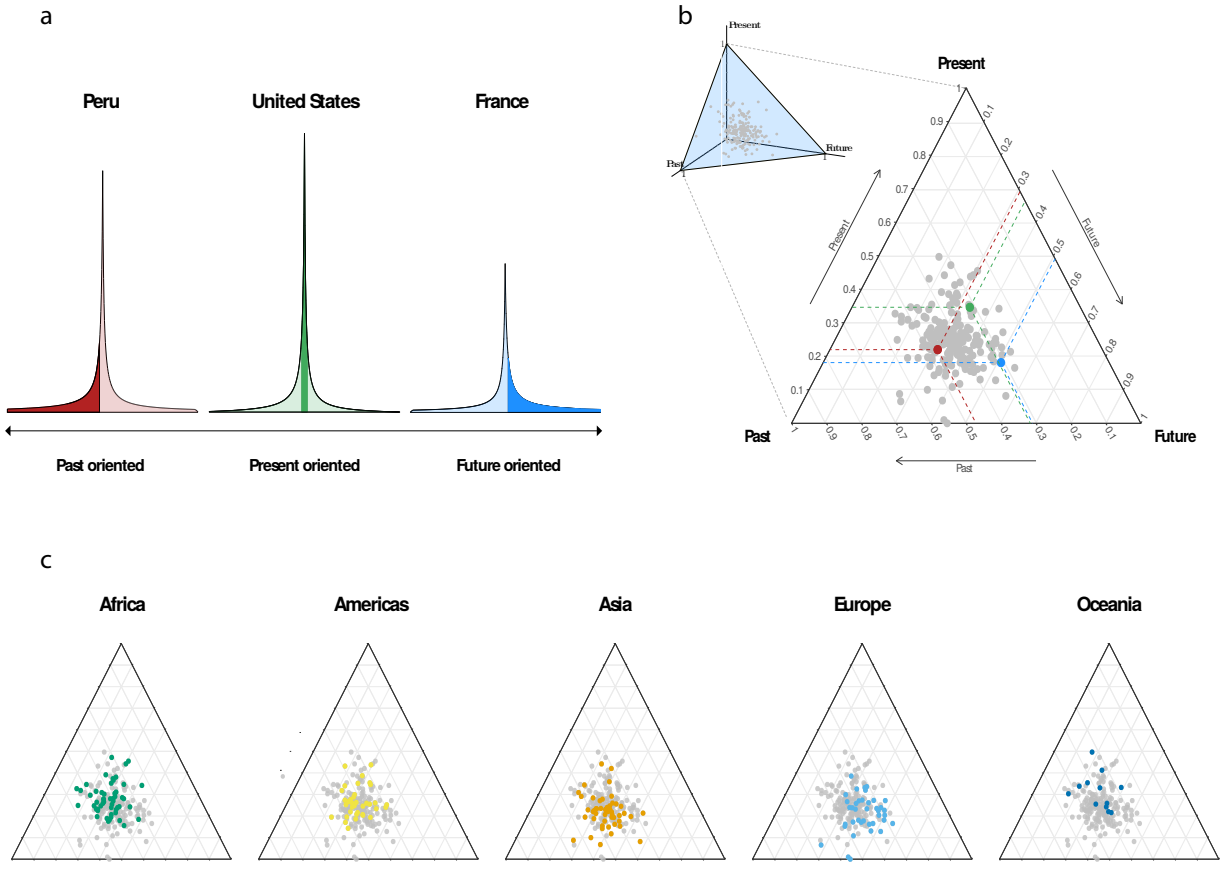

Figure A.16: **The relative attention to the past, present and future for individual countries.** **(a)** Example KAT for countries that are either past-oriented (Peru), present-oriented (United States), and future-oriented (France). Dark shaded areas indicate the portion of the KAT that is integrated to compute the total fraction of attention allocated each period in relative time. **(b)** Ternary plot that locates 181 countries in a three-dimensional space based on the fraction of attention it allocates to the past, present, and future. These fractions sum to one. A country will an equal attentiveness to the past, present, and future would be located in the center of the triangle. The more focused it is toward a direction in time, the closer it would be to the corresponding vertex. The three countries from (a) are highlighted using the same colors. All other countries are gray. Figure 3.1 depicts this distribution using a kernel-density-based heat map. **(c)** Same as (b), but highlighting five geographical regions using colored markers. Figure 3.4a-b in the main text maps these values.

Table A.6: **Country-specific holidays used to estimate the KAT for daily targets**.

Country	First holiday	Second holiday	Third holiday
Argentina	Day of Remembrance for Truth and Justice	International Workers' Day	New Year's Eve
Australia	Australia Day	Easter Monday	National New Year's Eve
Brazil	Tiradentes Day	Independence Day	All Souls' Day
Chile	Valentine's Day	Christmas Day Day	New Year's Eve
Colombia	Maundy Thursday	Teacher's Day	New Year's Eve
France	Victory in Europe Day	Bastille Day	New Year's Eve
Germany	International Women's Day	German Unity Day	New Year's Eve
Hong-Kong	Valentine's Day	Halloween	Christmas Day
India	Republic Day	Independence Day	New Year's Eve
Indonesia	Valentine's Day	Independence Day	The Prophet Muhammad's Birthday
Mexico	Cinco de Mayo	Cry of Dolores	Columbus Day
Netherlands	Valentine's Day	Liberation Day	Ascension Day
New-Zealand	Waitangi Day	Queen's Birthday	Halloween
Norway	17 May Constitution Day	Whit Sunday	New Year's Eve
Peru	Valentine's Day	International Workers' Day	New Year's Eve
Singapore	Valentine's Day	Vesak Day	National Day
Spain	Spanish National Holiday	Constitution Day	New Year's Eve
Sweden	Valentine's Day	Walpurgis Night	National Day
United Kingdom	Good Friday	Saint Patrick's Day	Boxing Day
United States	Valentine's Day	Independence Day	Mother's Day

Note: Holiday names are displayed in English, but data includes queries in local languages. These values are used to construct KAT estimates using day targets (e.g. Figure 3.3**f**).

Table A.7: Shares of attention to the past, present and future for alphabetically ranked countries 1 to 50 (standard errors in parentheses)

Country	Share of attention to the past	Share of attention to the present	Share of attention to the future
Afghanistan	0.4 (0.03)	0.25 (0.03)	0.35 (0.03)
Albania	0.5 (0.02)	0.16 (0.02)	0.34 (0.02)
Algeria	0.39 (0.02)	0.24 (0.02)	0.37 (0.02)
Angola	0.44 (0.02)	0.28 (0.03)	0.28 (0.02)
Argentina	0.35 (0.01)	0.23 (0.01)	0.42 (0.01)
Armenia	0.65 (0.02)	0.09 (0.03)	0.27 (0.02)
Aruba	0.34 (0.02)	0.26 (0.03)	0.4 (0.02)
Australia	0.28 (0.01)	0.33 (0.01)	0.39 (0.01)
Austria	0.3 (0.01)	0.2 (0.01)	0.5 (0.01)
Azerbaijan	0.41 (0.02)	0.21 (0.02)	0.38 (0.02)
Bahamas	0.37 (0.02)	0.35 (0.02)	0.28 (0.02)
Bahrain	0.41 (0.01)	0.22 (0.02)	0.37 (0.01)
Bangladesh	0.46 (0.02)	0.28 (0.03)	0.26 (0.02)
Barbados	0.41 (0.02)	0.3 (0.02)	0.29 (0.02)
Belarus	0.5 (0.01)	0.15 (0.02)	0.34 (0.01)
Belgium	0.31 (0.01)	0.22 (0.01)	0.46 (0.01)
Belize	0.42 (0.03)	0.29 (0.03)	0.29 (0.03)
Benin	0.37 (0.03)	0.31 (0.03)	0.32 (0.03)
Bhutan	0.35 (0.05)	0.36 (0.06)	0.29 (0.05)
Bolivia	0.46 (0.01)	0.19 (0.01)	0.35 (0.01)
Bosnia and Herz.	0.46 (0.02)	0.17 (0.02)	0.37 (0.02)
Botswana	0.49 (0.03)	0.23 (0.03)	0.28 (0.03)
Brazil	0.3 (0.01)	0.24 (0.01)	0.46 (0.01)
Brunei	0.43 (0.02)	0.27 (0.02)	0.29 (0.02)
Bulgaria	0.36 (0.01)	0.27 (0.01)	0.36 (0.01)
Burkina Faso	0.38 (0.03)	0.29 (0.03)	0.33 (0.03)
Burundi	0.35 (0.05)	0.28 (0.06)	0.37 (0.05)
Cabo Verde	0.46 (0.03)	0.18 (0.04)	0.36 (0.03)
Cambodia	0.47 (0.03)	0.23 (0.03)	0.3 (0.03)
Cameroon	0.47 (0.02)	0.18 (0.03)	0.35 (0.02)
Canada	0.32 (0.01)	0.35 (0.01)	0.33 (0.01)
Central African Rep.	0.24 (0.09)	0.34 (0.11)	0.42 (0.09)
Chad	0.24 (0.04)	0.45 (0.05)	0.31 (0.04)
Chile	0.29 (0.01)	0.25 (0.02)	0.46 (0.01)
China	0.42 (0.02)	0.19 (0.03)	0.39 (0.02)
Colombia	0.45 (0.01)	0.21 (0.01)	0.34 (0.01)
Comoros	0.38 (0.09)	0.39 (0.11)	0.23 (0.09)
Congo	0.41 (0.04)	0.28 (0.05)	0.31 (0.04)
Costa Rica	0.42 (0.01)	0.21 (0.02)	0.37 (0.01)
Croatia	0.42 (0.01)	0.17 (0.01)	0.41 (0.01)
Cuba	0.45 (0.02)	0.24 (0.02)	0.31 (0.02)
Curaçao	0.37 (0.03)	0.26 (0.03)	0.38 (0.03)
Cyprus	0.48 (0.01)	0.21 (0.02)	0.31 (0.01)
Czechia	0.36 (0.01)	0.23 (0.01)	0.4 (0.01)
Côte d'Ivoire	0.45 (0.02)	0.19 (0.02)	0.36 (0.02)
Dem. Rep. Congo	0.34 (0.03)	0.24 (0.04)	0.42 (0.03)
Denmark	0.34 (0.01)	0.26 (0.02)	0.4 (0.01)
Djibouti	0.43 (0.03)	0.24 (0.04)	0.33 (0.03)
Dominican Rep.	0.57 (0.01)	0.2 (0.01)	0.23 (0.01)
Ecuador	0.44 (0.01)	0.26 (0.01)	0.29 (0.01)

Table A.8: Shares of attention to the past, present and future for alphabetically ranked countries 51 to 100 (standard errors in parentheses)

Country	Share of attention to the past	Share of attention to the present	Share of attention to the future
Egypt	0.42 (0.02)	0.18 (0.02)	0.4 (0.02)
El Salvador	0.44 (0.02)	0.22 (0.02)	0.34 (0.02)
Eq. Guinea	0.35 (0.05)	0.29 (0.06)	0.36 (0.05)
Estonia	0.41 (0.02)	0.23 (0.02)	0.35 (0.02)
Eswatini	0.26 (0.06)	0.44 (0.07)	0.3 (0.06)
Ethiopia	0.57 (0.06)	0.25 (0.06)	0.18 (0.06)
Fiji	0.47 (0.02)	0.34 (0.02)	0.19 (0.02)
Finland	0.28 (0.01)	0.17 (0.01)	0.56 (0.01)
Fr. Polynesia	0.38 (0.02)	0.22 (0.02)	0.4 (0.02)
France	0.31 (0.01)	0.18 (0.01)	0.51 (0.01)
Gabon	0.51 (0.03)	0.2 (0.03)	0.29 (0.03)
Gambia	0.42 (0.04)	0.25 (0.05)	0.33 (0.04)
Georgia	0.56 (0.02)	0.1 (0.02)	0.34 (0.02)
Germany	0.3 (0.01)	0.24 (0.01)	0.47 (0.01)
Ghana	0.45 (0.02)	0.26 (0.02)	0.29 (0.02)
Greece	0.41 (0.01)	0.18 (0.02)	0.42 (0.01)
Grenada	0.46 (0.05)	0.32 (0.05)	0.22 (0.05)
Guam	0.39 (0.02)	0.24 (0.02)	0.37 (0.02)
Guatemala	0.46 (0.01)	0.27 (0.02)	0.28 (0.01)
Guinea	0.31 (0.06)	0.35 (0.07)	0.34 (0.06)
Guyana	0.39 (0.03)	0.43 (0.03)	0.18 (0.03)
Haiti	0.52 (0.02)	0.14 (0.03)	0.34 (0.03)
Honduras	0.5 (0.02)	0.2 (0.02)	0.3 (0.02)
Hong Kong	0.41 (0.01)	0.23 (0.01)	0.36 (0.01)
Hungary	0.36 (0.01)	0.39 (0.01)	0.26 (0.01)
Iceland	0.44 (0.02)	0.27 (0.02)	0.3 (0.02)
India	0.37 (0.01)	0.29 (0.02)	0.34 (0.01)
Indonesia	0.31 (0.01)	0.42 (0.01)	0.27 (0.01)
Iran	0.46 (0.03)	0.15 (0.03)	0.4 (0.03)
Iraq	0.27 (0.02)	0.21 (0.03)	0.52 (0.02)
Ireland	0.36 (0.01)	0.33 (0.01)	0.31 (0.01)
Israel	0.35 (0.01)	0.25 (0.01)	0.4 (0.01)
Italy	0.32 (0.01)	0.2 (0.01)	0.47 (0.01)
Jamaica	0.44 (0.01)	0.34 (0.02)	0.22 (0.02)
Japan	0.33 (0.01)	0.2 (0.01)	0.47 (0.01)
Jordan	0.53 (0.02)	0.07 (0.02)	0.4 (0.02)
Kazakhstan	0.47 (0.02)	0.23 (0.02)	0.3 (0.02)
Kenya	0.5 (0.02)	0.23 (0.02)	0.27 (0.02)
Kuwait	0.44 (0.01)	0.21 (0.02)	0.35 (0.02)
Kyrgyzstan	0.42 (0.03)	0.28 (0.03)	0.3 (0.03)
Laos	0.43 (0.02)	0.23 (0.03)	0.34 (0.02)
Latvia	0.43 (0.02)	0.18 (0.02)	0.39 (0.02)
Lesotho	0.33 (0.06)	0.38 (0.07)	0.29 (0.06)
Liberia	0.49 (0.05)	0.35 (0.06)	0.17 (0.05)
Libya	0.33 (0.02)	0.18 (0.03)	0.49 (0.02)
Lithuania	0.39 (0.01)	0.16 (0.01)	0.45 (0.01)
Luxembourg	0.29 (0.01)	0.23 (0.01)	0.49 (0.01)
Macao	0.42 (0.02)	0.23 (0.02)	0.35 (0.02)
Macedonia	0.49 (0.02)	0.11 (0.02)	0.41 (0.02)
Madagascar	0.41 (0.02)	0.23 (0.02)	0.36 (0.02)

Table A.9: Shares of attention to the past, present and future for alphabetically ranked countries 101 to 150 (standard errors in parentheses)

Country	Share of attention to the past	Share of attention to the present	Share of attention to the future
Malawi	0.46 (0.04)	0.31 (0.04)	0.23 (0.04)
Malaysia	0.39 (0.01)	0.18 (0.01)	0.43 (0.01)
Maldives	0.5 (0.03)	0.19 (0.03)	0.31 (0.03)
Mali	0.44 (0.03)	0.21 (0.03)	0.35 (0.03)
Malta	0.38 (0.02)	0.27 (0.02)	0.35 (0.02)
Mauritania	0.49 (0.03)	0.23 (0.03)	0.28 (0.03)
Mauritius	0.47 (0.02)	0.25 (0.02)	0.28 (0.02)
Mexico	0.42 (0.01)	0.26 (0.01)	0.31 (0.01)
Micronesia	0.33 (0.09)	0.41 (0.1)	0.26 (0.09)
Moldova	0.65 (0.02)	0.07 (0.02)	0.28 (0.02)
Mongolia	0.54 (0.03)	0.2 (0.03)	0.26 (0.03)
Montenegro	0.39 (0.02)	0.13 (0.02)	0.47 (0.02)
Morocco	0.47 (0.02)	0.19 (0.02)	0.34 (0.02)
Mozambique	0.4 (0.02)	0.28 (0.03)	0.32 (0.02)
Myanmar	0.51 (0.04)	0.29 (0.05)	0.2 (0.04)
Nepal	0.51 (0.03)	0.31 (0.03)	0.18 (0.03)
Netherlands	0.31 (0.01)	0.34 (0.01)	0.35 (0.01)
New Caledonia	0.39 (0.02)	0.22 (0.03)	0.39 (0.02)
New Zealand	0.38 (0.01)	0.26 (0.01)	0.36 (0.01)
Nicaragua	0.52 (0.02)	0.17 (0.02)	0.31 (0.02)
Niger	0.35 (0.04)	0.37 (0.05)	0.28 (0.04)
Nigeria	0.4 (0.02)	0.27 (0.02)	0.33 (0.02)
Norway	0.31 (0.01)	0.27 (0.01)	0.42 (0.01)
Oman	0.46 (0.02)	0.18 (0.02)	0.36 (0.02)
Pakistan	0.47 (0.01)	0.34 (0.02)	0.19 (0.02)
Palestine	0.46 (0.03)	0.12 (0.03)	0.42 (0.03)
Panama	0.42 (0.01)	0.2 (0.02)	0.37 (0.01)
Papua New Guinea	0.38 (0.05)	0.33 (0.05)	0.29 (0.05)
Paraguay	0.38 (0.02)	0.26 (0.02)	0.36 (0.02)
Peru	0.47 (0.01)	0.22 (0.01)	0.31 (0.01)
Philippines	0.51 (0.01)	0.23 (0.01)	0.26 (0.01)
Poland	0.43 (0.01)	0.24 (0.01)	0.34 (0.01)
Portugal	0.32 (0.01)	0.24 (0.01)	0.45 (0.01)
Puerto Rico	0.4 (0.01)	0.25 (0.01)	0.35 (0.01)
Qatar	0.42 (0.01)	0.24 (0.02)	0.34 (0.01)
Romania	0.43 (0.01)	0.3 (0.01)	0.28 (0.01)
Rwanda	0.43 (0.03)	0.3 (0.04)	0.27 (0.04)
S. Sudan	0.31 (0.04)	0.47 (0.04)	0.22 (0.04)
Saint Lucia	0.33 (0.04)	0.39 (0.04)	0.28 (0.04)
Samoa	0.42 (0.06)	0.35 (0.07)	0.22 (0.06)
Saudi Arabia	0.41 (0.02)	0.19 (0.02)	0.4 (0.02)
Senegal	0.5 (0.02)	0.25 (0.02)	0.25 (0.02)
Serbia	0.35 (0.01)	0.13 (0.01)	0.52 (0.01)
Sierra Leone	0.48 (0.04)	0.26 (0.05)	0.26 (0.05)
Singapore	0.38 (0.01)	0.26 (0.01)	0.36 (0.01)
Slovakia	0.4 (0.01)	0.26 (0.01)	0.34 (0.01)
Slovenia	0.37 (0.01)	0.17 (0.01)	0.47 (0.01)
Somalia	0.35 (0.05)	0.37 (0.06)	0.28 (0.05)
South Africa	0.45 (0.01)	0.27 (0.02)	0.29 (0.01)
South Korea	0.42 (0.02)	0.16 (0.02)	0.42 (0.02)

Table A.10: Shares of attention to the past, present and future for alphabetically ranked countries 151 to 181 (standard errors in parentheses)

Country	Share of attention to the past	Share of attention to the present	Share of attention to the future
Spain	0.44 (0.01)	0.24 (0.01)	0.32 (0.01)
Sri Lanka	0.49 (0.02)	0.22 (0.02)	0.29 (0.02)
St. Vin. and Gren.	0.47 (0.03)	0.27 (0.04)	0.26 (0.03)
Sudan	0.4 (0.02)	0.24 (0.03)	0.36 (0.02)
Suriname	0.44 (0.02)	0.28 (0.03)	0.28 (0.02)
Sweden	0.32 (0.01)	0.33 (0.01)	0.35 (0.01)
Switzerland	0.26 (0.01)	0.27 (0.01)	0.47 (0.01)
Syria	0.35 (0.03)	0.08 (0.03)	0.57 (0.02)
Taiwan	0.4 (0.01)	0.15 (0.01)	0.45 (0.01)
Tajikistan	0.48 (0.03)	0.23 (0.03)	0.29 (0.03)
Tanzania	0.43 (0.03)	0.31 (0.03)	0.26 (0.03)
Thailand	0.45 (0.02)	0.24 (0.02)	0.32 (0.02)
Timor-Leste	0.34 (0.03)	0.44 (0.04)	0.22 (0.03)
Togo	0.38 (0.03)	0.3 (0.04)	0.32 (0.03)
Trinidad and Tobago	0.44 (0.02)	0.25 (0.02)	0.31 (0.02)
Tunisia	0.42 (0.01)	0.22 (0.02)	0.37 (0.01)
Turkey	0.37 (0.01)	0.26 (0.02)	0.38 (0.01)
Turkmenistan	0.38 (0.03)	0.2 (0.03)	0.42 (0.03)
U.S. Virgin Is.	0.27 (0.04)	0.44 (0.05)	0.29 (0.04)
Uganda	0.44 (0.02)	0.28 (0.03)	0.28 (0.02)
United Arab Emirates	0.42 (0.01)	0.22 (0.01)	0.36 (0.01)
United Kingdom	0.33 (0.01)	0.32 (0.01)	0.35 (0.01)
United States of America	0.32 (0.01)	0.35 (0.01)	0.34 (0.01)
Uruguay	0.39 (0.01)	0.19 (0.01)	0.41 (0.01)
Uzbekistan	0.49 (0.03)	0.14 (0.03)	0.37 (0.03)
Vanuatu	0.43 (0.06)	0.26 (0.07)	0.31 (0.06)
Venezuela	0.48 (0.01)	0.25 (0.01)	0.27 (0.01)
Vietnam	0.48 (0.02)	0.15 (0.02)	0.36 (0.02)
Yemen	0.35 (0.03)	0.21 (0.03)	0.43 (0.03)
Zambia	0.37 (0.03)	0.35 (0.04)	0.28 (0.03)
Zimbabwe	0.49 (0.02)	0.3 (0.03)	0.21 (0.02)