NBER WORKING PAPER SERIES

JUDGING NUDGING:
UNDERSTANDING THE WELFARE EFFECTS OF NUDGES VERSUS TAXES

John A. List
Matthias Rodemeier
Sutanuka Roy
Gregory K. Sun

Working Paper 31152
http://www.nber.org/papers/w31152

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2023, Revised September 2023

David Franks went above and beyond with excellent research assistance and project management. Without his efforts the project would have never drawn to completion. We also thank Silvia Amalia Meneghesso and Roberto Fani for additional support. Hunt Allcott, Doug Bernheim, Luca Braghieri, Moritz Drupp, Nicola Gennaioli, Nicola Pavoni, Daniel Reck, and seminar participants at CESifo, the NBER Public Economics meeting, Frankfurt School, and the University of Chicago provided remarks that shaped our work. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Judging Nudging: Understanding the Welfare Effects of Nudges Versus Taxes
John A. List, Matthias Rodemeier, Sutanuka Roy, and Gregory K. Sun
NBER Working Paper No. 31152
April 2023, Revised September 2023
JEL No. C93,D61,D83

ABSTRACT

While behavioral non-price interventions ("nudges") have grown from academic curiosity to a bona fide policy tool, their relative economic efficiency remains under-researched. We develop a unified framework to estimate welfare effects of both nudges and taxes, while allowing for normative ambiguity about how nudges map into utility. We showcase our approach by creating a database of more than 300 carefully hand-coded point estimates of non-price and price interventions in the markets for cigarettes, influenza vaccinations, and household energy. While nudges are effective in changing behavior in all three markets, they are not necessarily the most efficient policy. When nudges are debiasing, they are more efficient in the market for cigarettes, while taxes are more efficient in the vaccine and energy market. Interestingly, these conclusions also often hold when nudges are deceptive rather than debiasing. We identify two key factors that govern the difference in results across markets: i) an elasticity-weighted standard deviation of the behavioral bias, and ii) the magnitude of the average externality. Nudges dominate taxes whenever i) exceeds ii). Finally, we consider cases in which nudges cause direct psychic costs or benefits to consumers.

John A. List
Department of Economics
University of Chicago
1126 East 59th
Chicago, IL 60637
and Australian National University
and also NBER
jlist@uchicago.edu

Matthias Rodemeier
Bocconi University
Via Roentgen 1
Milan 20136
Italy
rodemeier.matthias@gmail.com

Sutanuka Roy
HW Arndt Building 25A
The Australian National University
Kingsley Pl
Acton ACT 2601
Australia
sutanuka.roy5@gmail.com

Gregory K. Sun
Washington University in St Louis
greg.s@wustl.edu

1 Introduction

Non-price interventions motivated by insights from psychology, frequently referred to as "nudges," have received growing attention among academics and policymakers over the last decade. The extraordinary popularity of nudges has led to the implementation of many behavioral interventions across the globe attempting to shape individual behaviors. Nothing has seemingly been off limits, as today's choice architects commonly use social comparison nudges to help consumers conserve energy, warning labels on cigarette packages to deter smoking, and public campaigns to induce vaccinations. Many of these interventions are justified as cost-effective alternatives to traditional price and quantity regulations, such as taxes and subsidies or mandated quotas (Benartzi et al. 2017). The cost advantage over traditional policy tools is that the provision costs of nudges are often low, such that the change in behavior *per dollar spent on the intervention* is large.

While such a cost-based approach provides useful first insights into the comparison across policies, it does not take into account important factors of economic efficiency. Importantly, it does not quantify how the change in behavior caused by the nudge changes welfare to consumers and other members of society. A more complete benefit-cost analysis uses a revealed preference approach to understand how changes in behavior from various policy interventions map into welfare implications. There are few existing studies that include an applied welfare analysis of this type. By contrast, hundreds of empirical studies have estimated the reduced-form effects of nudges in many important markets of interest. We are therefore confronted with a vast amount of information on the efficacy of various nudges but have a far more limited understanding of their welfare impacts.

This challenge represents the motivation of our work. We develop an approach to estimate welfare effects of both nudges and taxes from the reduced-form treatment effects reported in hundreds of prior studies. In doing so, we provide the first comprehensive meta-analysis of welfare effects. Our methodology builds on the reduced-form approach to behavioral public policy evaluation due to Mullainathan, Schwartzstein and Congdon (2012), and combines it with a large database of carefully hand-coded point estimates of non-price and price interventions. To provide a glimpse across key domains in which nudges are ubiquitous and behavioral biases are allegedly important, we focus on data from three distinct markets: cigarettes, influenza vaccinations, and household energy consumption.

In total, we collect 304 point estimates on the effects of nudges and price interventions.[1] Our catalogue

[1] Since it is often difficult to precisely define what a nudge is, our selection of nudge categories is based on Benartzi et al. (2017).

of studies covers a wide array of interventions including social norm comparisons, reminders and feedback, planning prompts and goal setting nudges, defaults, as well as any informational intervention. The intuition of our theory and nature of our empirical results can naturally be viewed within the three distinct markets we empirically explore.

A general limitation in the literature on the welfare effects of nudging is that researchers need to make a number of judgement calls as to how biases affect utility and how nudges may correct these biases. We decide to embrace this ambiguity and test various competing models of nudges. We focus much of our attention on the most common model in the literature in which nudges act as "soft" paternalistic interventions that help consumers make better choices. This specification assumes that the treatment effects of nudges represent at least a partial reduction in a behavioral bias. This optimistic view is as a useful benchmark because it allows us to study whether, even under favorable assumptions about nudging, we find markets in which taxes dominate. We then explore alternative models that allow nudges to deceive consumers or impose direct psychic costs/benefits. Interestingly, a model in which nudges deceive consumers often yields the same policy recommendations to the model in which nudges are debiasing. Results are more sensitive, but often still conclusive, when nudges impose psychic costs or benefits. Our exercise, therefore, illustrates that we can identify efficient policies despite the ambiguity inherent in behavioral welfare economics.

In the market for cigarettes, we estimate that nudges, on average, increase the smoking cessation probability by 7.5% and reduce cigarette demand by 14%. We estimate an average price elasticity of -0.49, suggesting that the average nudge has the same effect on aggregate demand as a tax that increases the price of cigarettes by 28%. Leveraging our theoretical framework, we estimate the implied welfare impacts of nudges, of an optimal cigarette tax, as well as of a policy mix that combines both tools. In the benchmark case in which nudges debiase choices, we find that nudges cause a statistically significant increase in social welfare by $127 per consumer per year. Importantly, nudges tend to outperform cigarette taxes for a wide range of auxiliary parameter values. The optimal cigarette tax amounts to $3.49 per pack and raises welfare by $97 per consumer. Interestingly, a policy mix that combines a nudge with a tax is only slightly superior in terms of welfare gains than the nudge in isolation.

The key intuition driving these results can be found in two important statistics: i) the elasticity-weighted standard deviation of the behavioral bias, and ii) the size of the average externality. While both nudges and taxes can correct the *average* behavioral distortion, they each have a unique comparative advantage in our framework. The comparative advantage of nudges lies in the potential ability to reduce heterogeneity in the bias, while the comparative advantage of taxes is the internalization of marginal externalities. We show that nudges dominate taxes whenever i) is larger than ii). In the market for cigarettes, heterogeneity in the bias turns out to be a more important market distortion than externalities from smoking. This insight explains why nudges are more economically efficient than cigarette taxes.

For the second market we consider, influenza vaccinations, we estimate that a nudge increases the vaccination take-up by, on average, 35%, which corresponds to 13 percentage points. The mean price elasticity is -0.33, which indicates that vaccine subsidies would have to decrease prices by 105% to generate the same effect as nudges. Within our framework, this implies that the behavioral bias alone justifies a subsidy that makes influenza vaccines free. When estimating welfare effects, we find that the benefits of nudges over influenza subsidies are more limited than in the cigarette market. Nudges only dominate subsidies in the unlikely case where price elasticities and nudge treatment effects are extremely negatively correlated. In the more likely scenario in which nudge and price effects are positively correlated, subsidies slightly outperform nudges. In this scenario, the nudge raises welfare by, on average, $29 per person, while the optimal vaccine subsidy increases welfare by $65 per person. The ratio between heterogeneity in bias to the average externality can again explain this result. The average positive externality of getting vaccinated is substantially larger than the standard deviation of the behavioral bias, which makes subsidies more efficient than nudges. Yet, it is important to note that due to large standard errors, we cannot exclude the possibility that both policy tools yield equivalent welfare gains.

Finally, our most robust set of estimates come from the energy market. In terms of reduced-form effects, we find a mean average treatment effect of -4.8% on electricity consumption, and a mean price elasticity of -0.45. Thus, nudges have the same effect as a tax that raises the electricity price by 12.7%. For our baseline parameter values, we find that welfare gains from taxing electricity vastly exceed welfare gains from nudging: nudges increase social welfare by $108 per household per year, while the optimal tax raises welfare by $974 per household per year (the optimal tax equals $0.21 per kWh). For most parameter values, gains from taxation are 7-9 times larger than gains from nudging. Importantly, this conclusion is independent of our specific assumption regarding how nudge and price effects are correlated. Furthermore, a policy mix that adds a nudge to the tax provides virtually no additional benefits over implementing a tax alone. The underlying mechanism for this stark result is that the negative externalities from electricity consumption (in the form of carbon emissions) are dramatically larger than the standard deviation of the behavioral bias. Thus, heterogeneity in biases is a relatively negligible source of friction in the electricity market when put into perspective.

We consider an alternative model in which the treatment effect to a nudge does not necessarily constitute a reduction in the bias. A model in which nudges are deceptive makes taxes even more attractive but does not end up changing our policy recommendation in each market. If nudges impose moral costs, policy recommendations can change in the cigarette and vaccine market, but not in the energy market.

Overall, our results show that, while nudges are *effective* in all three applications, they are not always the most *efficient* intervention. In the benchmark model, the key statistics that predict when nudges dominate taxes are heterogeneity in the behavioral bias and the size of the average externality. As such, these two statistics can guide policymakers in choosing the most efficient policy instrument.

We view our combination of theory and empiricism as contributing to several unique strands of the literature. Our insights contribute to the nascent literature in behavioral public economics that studies optimal regulation in the presence of behavioral biases. We develop a theoretical model that builds on the frameworks in Bernheim and Rangel (2009), Mullainathan, Schwartzstein and Congdon (2012) and Farhi and Gabaix (2020) and provide useful extensions. Our empirical results add to a small set of papers that quantify the welfare effects of behavioral public policies. Prior studies have estimated the welfare effects of nudges and similar non-price interventions (Chetty, Looney and Kroft 2009, DellaVigna, List and Malmendier 2012, DellaVigna et al. 2016, Allcott and Kessler 2019, Rodemeier 2020, Goldin and Reck 2020, Allcott, Cohen, Morrison and Taubinsky 2022, Butera, Metcalfe, Morrison and Taubinsky 2022, Goldin and Reck 2022, Rodemeier and Löschel 2022, Löschel, Rodemeier and Werthschulte 2022, Barahona, Otero, Otero and Kim 2023, Reck and Seibold 2023), as well as of behaviorally-motivated taxes (Allcott, Mullainathan and Taubinsky 2014, Allcott and Taubinsky 2015, Allcott, Lockwood and Taubinsky 2019) and legal mandates (Dubois, Griffith and O'Connell 2018). Our paper makes a unique contribution to this literature by offering the first meta-analysis of welfare effects of both nudges and taxes, while relying only on reduced-form effects from prior studies.

Our structural estimates of behavioral parameters also contribute to the emerging literature in "structural behavioral economics" that identifies structural parameters proposed in theories at the intersection of psychology and economics (DellaVigna 2018). This literature has provided insights across a myriad of topics, including estimating discount functions over the life-cycle (Laibson, Repetto and Tobacman 2007), measuring the nature of risk preferences (Barseghyan et al. 2013) and projection bias (Conlin, O'Donoghue and Vogelsang 2007), as well as exploring gift exchange at work (DellaVigna et al. 2022) and why firms engage in corporate social responsibility (Hedblom, Hickman and List 2019).

Our paper also ties into a recent literature that uses the Marginal Value of Public Funds (MVPF) to evaluate the efficiency of prior interventions (Finkelstein and Hendren 2020, Hendren and Sprung-Keyser 2022). The MVPF method infers welfare from a policy by individuals' willingness to pay for that policy. This approach is enormously useful for applied welfare analysis in settings in which individual decisions are optimal. Since we study settings in which individuals allegedly fail to optimize, we develop a complementary approach that explicitly allows choices to be systematically biased.

From an applied policy perspective, our welfare results speak to an extensive literature on the role of behavioral interventions in health and energy policy. An interdisciplinary field involving medicine, psychology, and economics has studied how nudges can help people to stop smoking (e.g., Armitage and Arden 2008, Hennrikus et al. 2005, Borland, Balmford and Swift 2015), or get vaccinated (e.g., Milkman et al. 2011, Srinivasan et al. 2020, Frank, McMurray and Henderson 1985). Another large number of studies has investigated how nudges can be used to reduce households' energy consumption (e.g., Allcott 2011, Jessoe and Rapson 2014, Allcott and Wozny 2014, Allcott and Rogers 2014, Houde 2018, Andor, Gerster, Peters and Schmidt 2018, Allcott and Knittel 2019, Löschel, Rodemeier and Werthschulte 2022). Our paper complements these studies by offering an insight into the welfare effects of these interventions.

Finally, our paper offers a novel database of reduced-form effects of behavioral interventions and relates to a growing literature using meta analyses on nudging (Benartzi et al. 2017, Antinyan and Asatryan 2019, Hummel and Maedche 2019, DellaVigna and Linos 2022). Different from prior contributions, however, we leverage reduced-form estimates to draw insights on efficiency effects of various policy approaches and highlight how methodologies in this spirit can provide deep implications for the choice of the optimal policy tool.

The remainder of our study is structured as follows. In Section 2, we show theoretically how to evaluate the welfare effects of nudges and taxes based on reduced-form treatment effects. Section 3 describes the empirical estimation of the welfare formulae. We discuss the data collection in Section 4. In Section 5, we discuss both reduced-form estimates and the implied welfare effects. Section 6 concludes.

2 Theoretical Framework

In this section, we introduce a simple model that quantifies the welfare effects of nudges and taxes. We build on the frameworks in Bernheim and Rangel (2009), Mullainathan, Schwartzstein and Congdon (2012) and Farhi and Gabaix (2020) to provide new insights that link directly to our empirical work.

In order to describe our framework, we begin by specifying a (rational benchmark) demand side of the model. We assume that within a given market, there is a unit mass of potentially heterogeneous consumers indexed by i with incomes M_i and quasi-linear utility over the consumption good q_i and a numeraire good y_i. We assume that the possible choices of quantity consumed q_i is supported on $\mathcal{Q}$. Our empirical applications feature both a binary discrete choice setting, in which case, $\mathcal{Q} = \{0, 1\}$ and a continuous choice scenario with $\mathcal{Q} = \mathbb{R}_+$. Given prices p, each consumer i chooses q_i, y_i to solve

$$U_i(p) = \max_{(q_i, y_i) \in \mathcal{Q} \times \mathbb{R}_+} u_i(q_i) + y_i \quad \text{s.t.} \quad pq_i + y_i \leq M_i$$

Denote by $q_i^*(p, M_i)$ the solution to individual i's optimization problem. We make the usual assumption throughout that the consumer does not spend all of their income on the consumption good in question so $pq_i^*(p, M_i) < M_i$ for all i.[2] Together with the assumption of quasi-linearity, this implies that $q_i^*(p, M_i)$ does not depend on M_i. As a result, demand is also only a function of price and given by $D(p) = \int_i q_i^*(p)\,\mathrm{d}i$. Define the private aggregate benefit function as $R(p) \equiv \int_i U_i(p)\,\mathrm{d}i = \int_i u_i(q_i^*(p)) + M_i - pq_i^*(p)\,\mathrm{d}i$. Because we will eventually model departures from the rational benchmark model, it is helpful to decompose the private aggregate benefit function as

$$R(p) = \int_i u_i(q_i^*(p))\,\mathrm{d}i - pD(p) + \int_i M_i\,\mathrm{d}i$$

The first term in the above decomposition represents the aggregate gross consumption utility from consuming units of q_i (i.e., utility before accounting for benefits derived from the numeraire good) while the second term represents the aggregate disutility from paying for q_i. The last term depends only on income. In what follows, it is useful to express our formulae in terms of quantity consumed, so we exploit duality to define

[2] We view this assumption as empirically innocuous. It is difficult to think of *any* single consumption good for which some households plausibly spend *all* of their money.

the aggregate gross consumption utility in terms of quantity consumed:

$$V(q) \equiv \int_i u_i(q_i^*(D^{-1}(q)))\,\mathrm{d}i, \quad D^{-1}(q) = \inf\{p : D(p) \leq q\} \tag{1}$$

Intuitively, $V(q)$ thus corresponds to the aggregate consumption utility in the market when total level of consumption in the market is q and where these q units of consumption are allocated to their highest marginal utility consumers first. It thus allows us to isolate the consumption-based private *benefits* accruing to a benevolent policymaker who is able to induce aggregate consumption to equal q. Under the rational benchmark model, whether demand is discrete or continuous, aggregating the first-order conditions (FOC) of the individuals' optimization problems implies that $V'(D(q)) = p$. When we add behavioral frictions shortly, it will be possible that due to biases in consumer decision making, actual demand q differs from the rational benchmark level of demand $D(q)$, thus violating the above FOC.

The supply side of the model is characterized by a constant returns production function with marginal cost c. Firms are assumed competitive, so given a tax t, consumers face price $p = c + t$. Given our constant returns assumption, the supply side of the market is perfectly elastic, so firms in equilibrium earn zero profits.

Finally, we add two market frictions to the model that motivate policy making: externalities and internalities. The marginal externality is denoted by ξ and assumed constant. The internality, also referred to as the behavioral bias, is given by b_n, and is a function of a binary nudge $n \in \{0, 1\}$. A value of $b_n \neq 0$ affects the *behavioral* aspects of the model by making consumers systematically misperceive the marginal benefit of a unit of consumption. Formally, denoting by $D(p, b_n)$ the level of aggregate demand when price is p and behavioral bias is b_n, we assume that biased aggregate demand is characterized by the modified FOC that $V'(D(p, b_n)) + b_n = p$, rather than the usual first-order condition $V'(D(p)) = p$. Thus, when taxes are t and nudges are n, market demand is given by $D(c + t - b_n)$. Consumers overvalue benefits of consumption whenever $b_n > 0$, and undervalue them whenever $b_n < 0$. In this common specification, biases do not enter utility directly but rather lead to mistakes in choices, which create a wedge between marginal utility and price. We explore alternative models further below.[3] Conversely, the externality ξ does not affect choice but directly enters the social welfare function.

We add additional layers of heterogeneity by letting the private benefit V be a random function and letting

[3] We restrict our analysis to models in which preferences parameters are stable such that a welfare analysis is, in principle possible. We do not consider cases of context-dependent preferences in which choices can never be fully consistent, no matter the policy intervention. Bernheim (2023) discusses this case and proposes important solutions for future research that go beyond the scope of our paper.

internalities b and externalities ξ be random variables. Preferences and bias may co-vary in any arbitrary way. We model randomness in b and ξ as stemming from two conceptually distinct sources and hence treat the two sources of randomness somewhat differently. We view heterogeneity in b arising due to individual-level heterogeneity of market participants. For example, different consumers may be differentially inattentive to the costs of their energy consumption and hence exhibit different degrees of bias. On the other hand, we view the randomness in the marginal externality, ξ, as reflecting uncertainty about its true value *from the perspective of the policymaker*. For example, the policymaker may be uncertain about the externalities from energy consumption because damages from climate change are unknown. Because randomness in ξ is therefore exogenous from the perspective of individual market participants, throughout our discussion, we assume it to be independent of bias and utility, i.e. $\xi \perp b, V$.[4] This assumption does *not* mean that consumers with different preferences and biases produce the same expected amount of externalities; it simply means that the expected externality *per unit of consumption* is the same across consumers.

We are now ready to characterize the social welfare function. To do so, we need to be explicit about our treatment of taxation. We assume that the government retains a neutral budget and returns taxes to consumers in lump sum. Our quasi-linear framework then implies that these lump sum transfers are valued one-for-one with changes in consumer surplus. For any given realization of market parameters (V, b_n, ξ), social welfare is therefore

$$W(t,n) = \underbrace{V(D(c+t-b_n)) - (c+t)D(c+t-b_n)}_{\text{Consumer Surplus}} + \underbrace{tD(c+t-b_n)}_{\text{Government Revenue}} - \underbrace{\xi D(c+t-b_n)}_{\text{Externality}} \quad (2)$$

$$= V(D(c+t-b_n)) - cD(c+t-b_n) - \xi D(c+t-b_n). \quad (3)$$

Since policy variables (t, n) enter into the social welfare function *only* based on how they affect aggregate demand $D(c+t-b_n)$, we will often abuse notation and define the welfare function in quantity space: $W(q) \equiv V(q) - cq - \xi q$.

Note that, so far, this model implies that nudges affect social welfare *only* through their effects on consumer choice behavior. We consider alternative cases in which nudges also *directly* affect utility below, e.g., because consumers actively dislike warning labels on cigarette packages (Glaeser 2006, Loewenstein and

[4] In some settings, one may alternatively view randomness in ξ as arising from heterogeneity, as well. For instance, in the context of the electricity market, even if the social cost of carbon is a fixed number, some consumers may receive their electricity from solar panels, while others receive it from coal plants, leading to heterogeneity in ξ that could be correlated with b and D'. We abstract from these considerations in this paper.

O'Donoghue 2006).

For a given realization of market parameters (V, b_n, ξ), the first-best allocation of an omniscient social planner satisfies the condition that $V'(q) = c + \xi$. However, equilibrium quantities solve the first order condition $V'(q) = c + t - b_n$. Thus, consumption deviates from the first-best allocation because 1) consumers ignore social costs, and 2) consumers misperceive private benefits. Unless these two market frictions coincidentally cancel each other out for every consumer (i.e., $\xi = -b_n$ with probability 1), there is room for welfare-enhancing policy interventions. Within this simple framework, we next consider how taxes and nudges can alleviate market distortions and how we can measure their welfare effects empirically.

2.1 Quantifying the Welfare Effects of Nudges and Taxes

Throughout this section and the remainder of the paper, we make one final and common assumption, which can alternatively be thought of as an approximation. We assume that demand functions, $D(p)$, are linear. Since $V'(D(q)) = p$, linearity amounts to the assumption that $V''(q) = V''$ is constant. Thus, $V(q)$ is quadratic, as is $W(q)$. A more general interpretation of our results is that the following derivations can be viewed as second-order Taylor approximations of welfare effects under any arbitrary welfare function that can be represented by equation (3).

Let W^* be expected welfare given the first-best allocation described in the previous section, and let q^* solve $V'(q^*) = c + \xi$, so that $W^* = \mathbb{E}[W(q^*)]$. Since the welfare function is quadratic, we must have that $W(q) = W^* + \frac{1}{2}\mathbb{E}[(q - q^*)^2 V'']$. Additionally, our linear demand assumption implies that demand is also linear function of the policy variables t and n and is given by $D(c + b_n - t) = q^* + \frac{t-(b_n+\xi)}{V''}$ and that $D' = (V'')^{-1}$. We can thus alternatively write welfare as a function of taxes and nudges,

$$W(t, n) = W^* + \frac{1}{2}\mathbb{E}[(t - (b_n + \xi))^2 D']. \tag{4}$$

The key insights here are that i) the realized welfare loss from market frictions depends quadratically on the wedge between price and social cost and ii) its magnitude is increasing in the slope of the demand curve. This result is intuitive: where the demand curve has a larger slope, the social welfare function is more curved, so the welfare loss from missing society's first order condition is higher per unit error.

Figure 1 provides some intuition for the case in which there is only a behavioral bias, but no externalities and no pre-existing taxes ($\xi = 0$, $t = 0$). Panel (a) is the case with homogeneous bias, where D is biased demand and D^* is true marginal utility. The bias, $b < 0$, creates a wedge between willingness to pay and marginal utility, resulting in under-consumption by bD' units. The red triangle is the resulting deadweight loss of size $b_n^2 D'/2$ and is reminiscent of the classical "Harberger triangle" of the deadweight loss from taxation (Harberger 1964). It is then straightforward to see that if we introduce another market friction, such as externalities, the deadweight loss analogously becomes $(\xi + b_n)^2 D'/2$.

Panel (b) gives an example of heterogeneity, with $\mathbb{E}[b] = 0$ but $\mathrm{Var}(b) > 0$ so the bias is purely mean-zero noise. There are two levels of bias that realize with equal probability: $b_1 > 0$ causing overconsumption and $b_2 < 0$ causing underconsumption. The respective individual demand curves are D_1 and D_2. Aggregate demand is D and equal to unbiased aggregate demand. A fully de-biasing policy intervention would, therefore, not change aggregate demand and one might falsely conclude that it therefore had no positive effects on consumer surplus. In reality, the policy increased surplus by the sum of the red and blue triangles: $(b_1^2 + b_2^2)D'/2$. This result highlights the importance of variance in the behavioral bias.

Finally, panel (c) shows the case in which demand elasticities are heterogeneous but the bias is homogeneous. Note that a given level of b causes a larger deviation from the optimal quantity if consumers are more price elastic: $bD_1' > bD_2'$. The larger demand slope, D_1', implies a deadweight loss equal to the red triangle, while the smaller demand slope, D_2' implies a deadweight loss equal to the blue triangle. Higher price elasticities, for a given value of b, therefore, imply a larger deadweight loss. An important intuitive takeaway from this panel is that demand elasticities *weight* the import of market frictions on welfare distortions.

Our subsequent discussion will refer to the weighting interpretation of formula (4) often, so we introduce some additional notation to accommodate this insight. Define a set of weights as $W \equiv \frac{D'}{\mathbb{E}[D']}$. Define the weighted mean $\mathbb{E}_W$ by $\mathbb{E}_W[X] = \mathbb{E}[WX]$ for some random variable X. Similarly, define the weighted variance by $\mathrm{Var}_W(X) = \mathbb{E}_W[X^2] - \mathbb{E}_W[X]^2$. The interpretation of these weights is that $E_W[X]$ and $\mathrm{Var}_W(X)$ are the expectation and variance of X of *marginal* consumers, i.e. those consumers who would increase demand if prices fall by some very small amount.

We can now express equation (4) as

$$W(t, n) = W^* + \frac{1}{2}\mathbb{E}[D']\mathbb{E}_W[(t - (b_n + \xi))^2]. \tag{5}$$

We refer to welfare with no tax and no nudge, $W(0, 0)$, as our *baseline*, and then analyze welfare effects of nudges and taxes relative to this baseline. $W(0, 0)$ is given by

$$W(0,0) = W^* + \frac{1}{2}\mathbb{E}[D']\mathbb{E}_W[(b_0 + \xi)^2]. \tag{6}$$

2.1.1 Welfare Effects of a Nudge

The effect of a nudge depends on how it changes the internality and whether it causes direct (dis-)utility to consumers. We distinguish between three models that encompass a large variety of mechanisms. The first and most common specification in the literature is that nudges are (partially) de-biasing, meaning they correct the internality. For instance, reminders may make people more attentive and information provision may reduce biased beliefs. This is arguably how nudges are typically promoted: "soft" policy interventions that help consumers make better choices. The second specification is that nudges may deceive consumers, causing them to over- or undervalue consumption. In the third and final version of the model nudges impose a moral tax on consumers, making them feel guilty (or proud) about their consumption. We discuss each of these models below.

2.1.2 Bias-correcting Nudges

A bias-correcting nudge changes the internality by $\Delta_n b = b_1 - b_0$ with $b_1 = (1-\theta)b_0$ and $\theta \in (0,1]$. A partially de-biasing nudge has $\theta \in (0,1)$ and a fully de-biasing nudge has $\theta = 1$. In both cases, nudges reduce the mean and variance of the behavioral distortion. In our empirical analysis, we allow for arbitrary values of θ.[5] Welfare under nudging is given by

$$W(0,1) = W^* + \frac{1}{2}\mathbb{E}[D']\mathbb{E}_W\left[((1-\theta)b_0 + \xi)^2\right]. \tag{7}$$

The deadweight loss relative to first-best, therefore, depends on the squared sum of the *remaining* distortions that prevail in the market after nudging. These are the part of the bias that the nudge could not eliminate,

[5] In a contemporaneous and important study, Allcott et al. (2022) consider the case in which nudges may *increase* the variance of the bias, for example, because initially unbiased smokers see a cigarette warning label that they misinterpret, which then distorts their consumption. In our model this would amount to letting θ be heterogenous. Since no study among the vast literature on nudges, except for Allcott et al. (2022), provides a measure of this heterogeneity, we do not have sufficient information about the distribution of θ to accomodate this extension. However, if we are willing to make the assumption that θ is heterogenous but independent of b, then it is easy to show that our framework gives an *upper bound* for the welfare effects of nudges. Specifically, the weighted expected bias becomes $\mathbb{E}_W[b] = \mathbb{E}[TE]/\mathbb{E}[\theta][D']$ and the generalized welfare effect of a nudge in equation 9 is now

$$\Delta_n W(0,0) = -\frac{1}{2}\mathbb{E}[D']\left(2\mathbb{E}[\theta]\left(\mathbb{E}_W[b^2 + b\xi]\right) + \mathbb{E}_W[\theta^2 b^2]\right).$$

The important term is $\mathbb{E}_W[\theta^2 b^2] = \mathbb{E}[\theta^2]\mathbb{E}_W[b^2] = \left(\mathbb{E}[\theta]^2 + \mathrm{Var}(\theta)\right)\mathbb{E}[b^2]$. It follows that with $\mathrm{Var}(\theta) > 0$, the welfare gain from nudging is smaller than with $\mathrm{Var}(\theta) = 0$, and our empirical estimates will overstate the benefits of nudging.

$(1-\theta)b_0$, and the marginal externality, ξ, that the nudge cannot internalize. If the nudge is fully de-biasing, welfare simplifies to

$$W(0,1) = W^* + \frac{1}{2}\mathbb{E}[D']\mathbb{E}[\xi^2]. \tag{8}$$

The remaining deadweight loss relative to first best after nudging is therefore proportional to the expected squared externality, again an analogue to the usual Harberger triangle. Using the definition of variance, we can alternatively express equation (8) in terms of the mean and variance of the externality $W(0,1) = W^* + \frac{1}{2}\mathbb{E}[D']\left(\mathbb{E}[\xi]^2 + \text{Var}(\xi)\right)$. This shows that, even if there is no externality in expectation ($\mathbb{E}_W[\xi] = 0$), uncertainty in the externality causes an ex-ante welfare loss that is proportional to its variance.[6]

The welfare effect of a nudge relative to baseline can be written as

$$\Delta_n W(0,0) = -\frac{1}{2}\mathbb{E}[D']\mathbb{E}_W[\theta(2-\theta)b_0^2 + 2\theta b_0\xi], \tag{9}$$

which, in the case of complete de-biasing, simplifies to

$$\Delta_n W(0,0) = -\mathbb{E}[D']\mathbb{E}_W\left[\frac{b_0^2}{2} + b_0\xi\right]. \tag{10}$$

The first term ($-\frac{1}{2}\mathbb{E}[D']\mathbb{E}_W[b_0^2]$) in equation 10 is the behavioral analogue to the Harberger triangle, as discussed in Figure 1. As above, we can further decompose this effect of nudges into mean and variance components, so $-\frac{1}{2}\mathbb{E}[D']\mathbb{E}_W[b_0^2] = -\frac{1}{2}\mathbb{E}[D'](\mathbb{E}_W[b_0]^2 + \text{Var}(b_0))$. As we will show, the comparative advantage of a bias-correcting nudge relative to taxes is precisely in its ability to potentially address the variance term in the decomposition.

The second term in equation (10), ($-\mathbb{E}[D']\mathbb{E}_W[b_0\xi]$), captures the interaction between the effects of bias and externality on welfare. This term reflects the fact that when one distortion already exists in the market, adding a second distortion will have first-order effects on welfare.

2.1.3 Deceptive Nudges

Nudges may also be deceptive and exacerbate biases or move them into the opposite direction. In our framework, these nudges have either $\theta > 1$ or $\theta < 0$, respectively.[7] The welfare effect of these nudges is captured

[6] Diamond (1973) studies optimal taxation in this setting, in which there is no behavioral bias but a heterogeneous externality. Our framework accommodates his model.

[7] Since we specified that a nudge changes the bias by some *proportion* ($\Delta b = -\theta b_0$), we implicitly exclude the case in which a nudge biases fully unbiased consumers, i.e., cases with $b_0 = 0$ and $b_1 \neq 0$. This is merely a technicality. In our empirical analysis

by equation 9, as well, but the effect on consumer surplus is now ambiguous instead of strictly positive. To illustrate this more clearly, we can rewrite equation 9 as

$$\Delta_n W(0,0) = -\frac{1}{2}\mathbb{E}[D']\mathbb{E}_W\left[\underbrace{b_0^2}_{\text{de-biasing}} + \underbrace{b_0^2(\theta(2-\theta)-1)}_{\text{departure from complete de-biasing}} + \underbrace{2\theta b_0 \xi}_{\text{indirect externality effect}}\right]. \tag{11}$$

The first term is the effect of-debiasing on consumer welfare as in equation 10. The second term measures any efficiency departures from full de-biasing on consumers. The second term is strictly negative for $\theta \in (0,1)$ which shows that partial de-biasing will yield lower consumer surplus effects than full de-biasing. For the case of deception with $\theta > 1$ or $\theta < 0$ the term is (weakly) negative but there is an additional term that captures first-order effects of deception on externalities: $2\theta b_0 \xi$.[8] Deception increases social welfare if the reduction in externalities outweighs the reduction in consumer surplus (see also Rodemeier and Löschel 2020). Formally, deception dominates a fully de-biasing nudge if and only if

$$(\theta(2-\theta)-1)\mathbb{E}_W[b_0] > 2(1-\theta)\mathbb{E}[\xi]. \tag{12}$$

If governments can steer the level of θ, optimal deception implies

$$\theta^* = 1 + \frac{\mathbb{E}_W[b_0]\mathbb{E}[\xi]}{\mathrm{Var}_W[b_0] + \mathbb{E}_W[b_0]^2}. \tag{13}$$

The term for optimal deception is obtained by maximizing 7 with respect to θ. Equation 13 illustrates that optimal deception depends on the size of the externality relative to the behavioral bias. This becomes particularly clear if the bias is homogenous, in which case equation 13 would set the bias exactly equal to the average externality, $\theta^* = 1 + \frac{\mathbb{E}[\xi]}{b_0}$. In other words, optimal deception operates like a Pigou tax that internalizes the externality on average. Of course, this is a highly idealized result as it requires that the social planner can induce any arbitrary bias (i.e., set θ to any value) in order to dictate aggregate quantities. We do not consider this view as realistic in most settings.

There are additional arguments why deceptive nudges are unlikely to be part of an optimal policy mix outside of our framework. First, consumers may learn over time which information to trust, such that decep-

we do not find a single case in which the initial bias is *exactly* zero.

[8] The term that captures the departure from complete de-biasing is *weakly* negative because it equals zero in the special case in which $\theta = 2$. In this case, the nudge only changes the sign of the bias but not the absolute magnitude: $b_1 = -b_0$.

tion only works in the short run. Second, if governments use deception, they are likely to lose voters' trust and support, making deception unsustainable. Despite these arguments, we will incorporate deception into our analysis to embrace the inherent ambiguity about how nudges affect welfare.

2.1.4 Nudges as Moral Taxes

Finally, we consider the case in which nudges act as a moral tax. Let κ_n be the "guilt" for every unit of consumption induced by a binary nudge (so negative values of κ_n could be interpreted as "warm glow" or "pride"). Specifically, normative consumer surplus is given by $V(q)-pq-\kappa_n q$, and consumers make choices according to their first-order condition $V'(q) = p - b_n + \kappa_n$. Relative to the first best, $V'(q) = c + \xi$, the allocation chosen by consumers now satisfies $V'(q) = c-b_n+\kappa_n$. A key conceptual difference relative to the previous versions of the model is that nudges now have a *direct* impact on consumer's utility which enters the welfare function. In what follows, we will often group debiasing/deceptive nudges together under the umbrella term of "information nudges", to highlight that they only affect utility indirectly through changing choices and contrast these to the moral tax nudges considered here.

In the interest of highlighting the key conceptual difference, we will work here with a model of moral taxation which mirrors our model of information nudges, except that nudges directly enter into the social welfare function. To do so, we posit that the moral tax κ_n is proportional to the bias b_n with proportionality constant θ, so $\kappa_n = \theta b_n$. This parameterization means that the implications for nudges or taxes in terms of *outcomes* are identical to the information model of nudges. Thus, the only difference in welfare arises due to the direct welfare impact of the guilt/warm glow induced by the nudge. When considering the effect of a nudge in isolation, the direct welfare effect of the nudge is therefore

$$-\mathbb{E}[\underbrace{(D + \theta b_0 D')}_{\text{post-nudge demand}} \times \underbrace{\theta b_0}_{=\kappa_n}].$$

This direct welfare effect is negative whenever biases are positive (i.e., individuals *over-consume* relative to the rational benchmark) and is positive whenever biases are negative (i.e., individuals *under-consume* relative to the rational benchmark). As a result, there are only two possible cases where viewing nudges as moral taxation rather than as information will change the relative attractiveness of nudging vs. taxing. Either (i) nudges originally looked superior to taxes, but viewing them as moral taxes adds an additional cost to nudging or (ii) nudges originally looked inferior to taxes, but viewing them as moral subsidies adds an additional benefit to nudging. In either case, we can compute the welfare effect of a nudge on welfare relative to the status quo policy as

$$\Delta_n W(0,0) = -\frac{1}{2}\mathbb{E}[D']\mathbb{E}_W\left[2\theta b_0^2 + 2\theta b_0 \xi\right] - \mathbb{E}[D\theta b_0]. \tag{14}$$

We will explore empirically if and when the welfare implications of viewing nudges as moral taxes (equation 14) leads to different conclusions compared to viewing nudges as information (equation 11).

2.1.5 Which model of nudges should be preferred?

The previous models illustrate that behavioral welfare analysis is subject to substantially more ambiguity than traditional policy evaluation. The analyst needs to take a stance not only on how much nudges reduce internalities, but also on how they affect consumer welfare directly. This ambiguity cannot be avoided when studying welfare effects of nudges and our paper shares this limitation with the prior literature.

Importantly, our meta-analysis is not aimed at discerning among these distinct models. Instead, our goal is to explore whether—despite the modelling uncertainty—the substantial body of evidence concerning the reduced-form impact of nudges can contribute to advancing our understanding of the comparative effectiveness between nudges and conventional taxes. While we investigate all of the three prior models in this paper, we must ultimately decide on a preferred specification that receives most attention in this paper.

In our main analysis, we focus on the most common case in the literature in which nudges are (partially) bias-correcting. We choose this as our benchmark specification because we want to model nudges the way proponents of nudges typically view them: as "soft" paternalistic interventions that help consumers make better decisions. Many of the nudges we study empirically are in fact reminders that plausibly increase attention and informational interventions that provide objective facts about the consequences of consumption. We recognize that a model in which nudges correct biases could be viewed as a notably optimistic portrayal, but it proves valuable for informing public policy by addressing the subsequent questions: Even if nudges work as intended, in how many markets do taxes still dominate? Which factors determine when nudges dominate?

After we have established answers to these questions, we explore the alternative models of nudging in a robustness analysis. As we discuss later in this paper, these extensions change the magnitude of the *absolute* welfare effect of nudges, but they rarely change whether we want to use nudges rather than taxes in a given market (i.e., the relative welfare effect). The latter is the main question we intend to answer within our framework. Reassuringly, our results imply that we can often offer qualitative policy recommendations independent of the underlying model of nudging. Our paper, therefore, shows that we can identify efficient policies despite the substantial ambiguity inherent in behavioral welfare economics.

2.1.6 Welfare Effects of a Tax

Consider again the welfare function in equation 5. Suppose now that the government instead decides to use a tax rather than a nudge. To optimize the choice of a tax amount, we differentiate the expected welfare criterion with respect to t, which yields:

$$\mathbb{E}\left[(t-(b_0+\xi))D'\right]=0. \tag{15}$$

The optimal t must therefore satisfy $t_0^* = \mathbb{E}_W[(b_0+\xi)]$, which is a weighted average of the sum of the market frictions, where again, the weights are proportional to the slope of the demand curve. The optimal tax puts larger weights on more price elastic consumer groups, because larger demand elasticities imply larger deviations from the optimum for any given value of $b_0+\xi$. This result again highlights the import of our weighting intuition and has previously been established by Allcott and Taubinsky (2015).

Consider now, the welfare effect of the optimal tax. Substituting the optimal tax into $W(t,0)$ shows that welfare under this policy is $W(t_0^*,0) = W^* + \frac{1}{2}\mathrm{Var}_W(b_0+\xi)\mathbb{E}[D']$. This result illustrates that once the government has corrected the expected sum of the market frictions, the remaining deadweight loss stems from heterogeneity, which is determined by the weighted variance of market frictions, $\mathrm{Var}_W(b_0+\xi)$, and the aggregate demand elasticity. Since the tax cannot be targeted—i.e., be set equal to each realization of $b_0+\xi$—it cannot fully correct each consumer's choice, and therefore cannot achieve the first-best allocation.

The welfare effect from the optimal tax relative to a baseline with no tax and no nudge is

$$\Delta_t W(0,0) = -\frac{1}{2}\mathbb{E}[D']\mathbb{E}_W[b_0+\xi]^2 \tag{16}$$

$$= -\frac{1}{2}\frac{\left(\mathbb{E}[D']\mathbb{E}[b_0+\xi]+\mathrm{Cov}(D',b)\right)^2}{\mathbb{E}[D']} \tag{17}$$

The welfare effect increases quadratically in the weighted expectation of the market frictions. This expectation can be decomposed into the unweighted expected sum of the market frictions and the covariance between these frictions and demand elasticities. If the bias is uncorrelated with demand slopes ($\text{Cov}(D', b_0) = 0$) the welfare effect simplifies to $\Delta W_t(0,0) = -\frac{1}{2}\mathbb{E}[D']\mathbb{E}[b_0 + \xi]^2$.

It is worthwhile to point out the relationship between the above formulae and the prior sufficient statistics literature. Specifically, consider the case where there are no externalities ($\xi = 0$). The optimal tax then becomes $t_0^* = \mathbb{E}_W[b_0]$ while the welfare impact becomes $-\frac{1}{2}\mathbb{E}[D']\mathbb{E}_W[b_0]^2$. Thus, $\mathbb{E}_W[b_0]$ is a key quantity of interest, and we aim to enhance our understanding of this term. Some simple algebraic manipulations show that $\mathbb{E}_W[b_0] = \frac{\mathbb{E}[D'b_0]}{\mathbb{E}[D']}$. Recall from Figure 1 that $D'b_0$ would be the change in demand induced by a fully de-biasing nudge. Thus, letting T be the *treatment effect* of a de-biasing nudge, we see that in fact, $\mathbb{E}_W[b_0] = \frac{\mathbb{E}[T]}{\mathbb{E}[D']}$. The RHS of this final expression is a well-known quantity in the literature and is commonly referred to as the *equivalent price metric* (EPM). This quantity can be interpreted as measuring the average bias of *marginal* consumers. Our framework provides an interpretation of the EPM as a particular sufficient statistic. Specifically, under fairly weak restrictions on heterogeneity (i.e., linear demand functions), it corresponds to the optimal uniform tax when i) markets are perfectly competitive, ii) nudge policies are unavailable, or infeasible at scale, and iii) there is no cost of raising government funds. While these assumptions may be restrictive in general, the discussion nonetheless clarifies that the EPM corresponds to a (very specific) policy relevant parameter.[9]

2.1.7 Welfare Effects of Taxes and Nudges in Combination

We now compare the welfare effects of the optimal tax to the welfare effects of a de-biasing nudge. To isolate the key tradeoffs, it is helpful to consider the best-case scenario for nudges where they are perfectly effective ($\theta = 1$). Using the formulae derived above and in this subsection, we find the welfare difference under these respective policies as

$$W(t_0^*, 0) - W(0, 1) = \frac{1}{2}\mathbb{E}[D']\left(\text{Var}_W(b_0) - \mathbb{E}[\xi]^2\right). \tag{18}$$

Nudges are therefore superior to taxes iff the expression in the parentheses on the RHS of equation (18) is greater than or equal to 0, or likewise if

$$\text{sd}_W(b_0) \geq |\mathbb{E}[\xi]|, \tag{19}$$

[9] Allcott and Taubinsky (2015) show that the EPM statistic may deviate from $\mathbb{E}_W[b]$ when the density of marginal consumers with some realization of b changes along the aggregate demand curve. Our assumption of linear demand for each realization of b rules out this possibility. Despite its potential limitations, linear demand is a standard assumption in the literature on sufficient statistics and structural estimation.

where sd_W is the standard deviation of the behavioral bias among marginal consumers. This result reveals that de-biasing nudges are superior to taxes if the degree of heterogeneity in bias is *larger* than the magnitude of the average externality. We therefore refer to the ratio between the LHS and the RHS of Equation (18), $\frac{\mathrm{sd}_W(b_0)}{|\mathbb{E}[\xi]|}$, as the *targeting ratio*.

This discussion highlights that in our framework de-biasing nudges and taxes fundamentally target two separate sources of market inefficiency: nudges act implicitly in the spirit of a first-degree price discrimination tool, targeting consumers and correcting their individual-specific biases. They therefore tend to work best when biases are heterogeneous, so that the value of this implicit first-degree price discrimination is large. This results goes back to our assumption that nudges reduce variance in bias, which is one of the important judgement calls that one has to make when inferring welfare effects from prior studies. Since there may be cases in which nudges could *increase* variance, our framework provides an optimistic view of how nudges affect choices. Again, our objective is to offer a framework in which nudges work the way proponents of nudges want them to work and study whether even under this favorable set of assumptions, we discover markets in which taxes dominate.

Different from taxes, however, nudges have no ability to internalize the *marginal* externality. While they can affect the *total* level of externalities indirectly through altering consumption, the wedge between social and private benefits at the margin (as measured by ξ) remains unchanged. For example, nudging households to reduce energy consumption decreases the overall level of carbon emissions, but it does not internalize the per-unit external damage of energy consumption (i.e., the social cost of carbon).[10] By contrast, an energy tax imposes a price on every unit of carbon, thereby internalizing the marginal damage of energy consumption.

As such, while both nudges and taxes can correct the *average* bias of marginal consumers, only nudges can reduce variance in bias, and only taxes can internalize the externality.[11]

The insight of equation (19) is that taxes have a comparative advantage over nudges when the average externality is high and the variance in bias of consumers at the margin is low. In light of the relative strengths of nudges and taxes, a policy that uses taxes and nudges in conjunction can obtain the best of both worlds, as they are able to compensate for each other's shortcomings. While this is true by assumption in our framework, it is an empirical question by how much a policy mix adds quantitatively over using one policy tool in isolation.

[10] An interesting extension for future work is to allow ξ to be a function of q, in which case the nudge may indirectly change the marginal externality, as well.

[11] In our alternative models of deception and nudges as moral taxes, nudges may also potentially internalize the marginal externality. We show that given the data, these alternative models do not change many of our qualitative conclusions from our preferred model.

Later we find that empirically the incremental benefits from mixing the two tools can be vanishingly small in some markets.

Before moving to the empirics, however, let us first establish how we can quantify the welfare effect of a policy mix. Note that we can improve upon a nudge with a tax that corrects the remaining market distortions after the nudge has been implemented. This is operationalized by the tax $t_1^* = \mathbb{E}[\xi]$. Welfare under this policy is given by $W(t_1^*, 1) = \frac{1}{2}\mathbb{E}[D']\mathrm{Var}(\xi)$, implying that the remaining deadweight loss relative to first best again stems from heterogeneity in the externality.

The welfare gain from this combination of taxes and nudges relative to no nudge and no tax is

$$\Delta_{tn} W(0,0) = -\frac{1}{2}\mathbb{E}[D'](\mathbb{E}_W[b_0^2] + 2\mathbb{E}_W[b_0\xi] + \mathbb{E}[\xi]^2). \tag{20}$$

The first term, $\mathbb{E}_W[b_0^2]$, reflects the deadweight loss from the behavioral bias, which nudges are assumed to fully address. The second term, $2\mathbb{E}_W[b_0\xi]$, again is an interaction between biases and externalities. Finally, $\mathbb{E}[\xi]^2$ reflects the welfare gain from taxing the externality. As already mentioned in the context of taxes in isolation, it differs from the total distortion arising from the externality, $\mathbb{E}[\xi^2] = \mathbb{E}[\xi]^2 + \mathrm{Var}(\xi)$, due to the fact that taxes cannot address the uncertainty in the externality.

2.1.8 General Welfare Formula of a Policy Mix

If nudges are only partially de-biasing the welfare gain of a policy mix is

$$\Delta_{tn} W(0,0) = -\frac{1}{2}\mathbb{E}[D'](\mathbb{E}_W[\theta(2-\theta)b_0^2 + 2\theta b_0\xi] + \mathbb{E}_W[(1-\theta)b_0 + \xi]^2). \tag{21}$$

Given this more general formula, we can also generalize our analysis of when nudges dominate taxes. Specifically, we derive the following generalization of Equation (18) by comparing Equations (9) and (17):

Nudges dominate taxes if and only if

$$\theta(2-\theta)\mathrm{sd}_W(b_0) \geq |(1-\theta)\mathbb{E}_W[b_0] + \mathbb{E}[\xi]| \tag{22}$$

This inequality generalizes the logic of Equation (18). The LHS continues to be proportional to the weighted standard deviation of biases and again relates to the ability of nudges to (partially) address heterogeneous biases. This heterogeneity is multiplied by a factor of $\theta(2-\theta)$, reflecting the fact that when $\theta < 1$, nudges only partially address biases. Relative to Equation (18), the RHS of Equation (22) has an extra term, $(1-\theta)\mathbb{E}_W[b_0]$. This is the average degree of bias even after nudging. When nudges are fully debiasing ($\theta = 1$), this term vanishes, so again, the presence of this extra term reflects the fact that for general values of θ, nudges may not be fully effective in addressing behavioral biases. In what follows, we often refer to

$$\frac{\theta(2-\theta)\mathrm{sd}_W(b_0)}{|(1-\theta)\mathbb{E}_W[b_0] + \mathbb{E}[\xi]|} \tag{23}$$

as the *generalized targeting ratio*. Our results imply that nudges dominate taxes if and only if the generalized targeting ratio exceeds one.

3 Empirical Implementation

In this section we connect our theoretical framework to empirical observations. In Subsection 3.1, we describe the form of our meta-analysis datasets and discuss the key identification and estimation challenges in linking our theoretical framework to the available data. In Subsection 3.2, we describe our concrete algorithm for estimating the welfare impacts of nudges and taxes following the discussion in Subsection 3.1.

3.1 Identification and Estimation

3.1.1 Meta-analysis data

The discussion in this section should be understood as describing our empirical strategy for a single market (e.g., cigarettes). We implement the formulae derived in this section three times, one for each studied market. For each market, we have access to two meta-analysis datasets, which we will refer to as N and Π, corresponding respectively to *nudge* and *price* treatments. Each study $n \in N$ estimates the effect of some nudge on quantities while each study $\pi \in \Pi$ estimates the effect of price on quantity. For each nudge study $n \in N$, we record an estimated percent treatment effect $\hat{T}_n$ as well as a reported standard error $\hat{\sigma}_n$. Alternatively, each pricing study $\pi \in \Pi$ corresponds to the effect of a price change on quantities. For these studies, we record an estimated demand elasticity $\hat{\varepsilon}_\pi$ as well as a reported standard error, $\hat{\sigma}_\pi$.

We assume that our meta-analysis datasets are generated as follows. Each study $n \in N$ corresponds to a realization of the market, so each study is associated with a draw from the joint distribution of demand curves and nudges: $(D_n, b_n) \sim F_{D,b}$. Given these realizations of market parameters, the "true" percent treatment effect for the study is given by $T_n = \theta \frac{D'_n}{D_n} b_n$. The estimated treatment effects and standard errors are then drawn in such a way that $\hat{T}_n \sim \mathcal{N}(T_n, \hat{\sigma}_n)$. We assume that data from the price study are generated in a similar manner. Each study $\pi \in \Pi$ corresponds to a realization of the market and is thus associated with a draw $(D_\pi, b_\pi) \sim F_{D,b}$. The true elasticity is thus given by $\varepsilon_\pi = \frac{pD'_\pi}{D_\pi}$, while estimated treatment effects and standard errors are drawn so that $\hat{\varepsilon}_\pi \sim \mathcal{N}(\varepsilon_\pi, \hat{\sigma}_\pi)$.

Note that the formulation stated here, which we implement in the main body of the paper, assumes that studies are drawn from the population distribution of markets. One salient reason why this assumption may be flawed is the presence of publication bias: if not all studies get published with equal probability, and non-publication occurs for systematic reasons, the sample of point estimates we have obtained for our meta-analysis may be non-representative. In Appendix D, we correct for publication bias and find that our main qualitative insights are not sensitive to this correction.

3.1.2 Identification of Nudge Treatment Effect and Elasticity Marginal Distributions

Given the data-generating process for N and Π, standard deconvolution techniques can be used to show that the marginal distributions F_T and F_ε of treatment effects T and elasticities ε are non-parametrically identified, assuming a large dataset of studies.[12] In practice, however, we have found the non-parametric approach to be too demanding on data relative to our sample sizes. As a result, we have opted for a parametric approach. Specifically, we assume that $T_n \sim \log\mathcal{N}(\mu_T, \sigma_T^2)$ while $\varepsilon_n \sim \log\mathcal{N}(\mu_\varepsilon, \sigma_\varepsilon^2)$.

In this case, the law of iterated expectations shows that $\mathbb{E}[\hat{T}_n] = \mathbb{E}[T_n]$ while the law of total variance implies that $\mathrm{Var}(\hat{T}_n) = \mathrm{Var}[T_n] + \mathbb{E}[\mathrm{Var}_n(T_n)] = \mathrm{Var}[T_n] + \mathbb{E}[\hat{\sigma}_n^2]$. Thus, the sample mean of estimated treatment effects is a consistent estimate of the average treatment effect while the sample analogue of the ANOVA identity allows us to consistently estimate the variance of treatment effects in the population. The same analysis shows how the mean and variance of the distribution of elasticities can be consistently estimated. Because there is an injective mapping from means and variances into the parameters, μ, σ^2 of a log-normal distribution, the method of moments allows us to translate these estimates of $\mathbb{E}[T_n], \mathbb{E}[\sigma_n^2], \mathbb{E}[\varepsilon_\pi], \mathrm{Var}[\varepsilon_\pi^2]$ into estimates of the parameters of the log-normal distribution, $\mu_T, \sigma_T^2, \mu_\varepsilon, \sigma_\varepsilon^2$.

[12] See, for instance, Stefanski and Carroll (1990).

3.1.3 Non-identification of Welfare Formulae

Above we learned that the marginal distribution of treatment effects, F_T, and the marginal distribution of elasticities, F_ε, are identified using our meta-analysis data. However, looking at welfare formulae (10), (17), and (20), it is clear that the joint distribution of (b_0, ε), along with the nudge effectiveness parameter θ are all needed to identify the welfare effects of our key policies under consideration. This implies that our meta-analysis data alone does not suffice to point identify the effects we aim to study without additional structure. In this section, we therefore detail what information is missing and how we can use auxiliary information, structural assumptions, and sensitivity analyses to overcome the resulting non-identification issues.

First, since we are focusing on performing a meta-analysis of nudge effects, we do not directly analyze data about the marginal distribution of externalities ξ. Instead, we rely on relevant prior literatures to provide us with estimates of the distribution of ξ. Second, the studies in our meta-analysis do not include price and nudge treatments within the same experiment. As a result, we are not able to identify the dependence structure between b_0 and ε. We therefore perform a sensitivity analysis to explore the robustness of our results to assumptions on how b_0 and ε are related. Specifically, in the next subsection, we describe how we parameterize the correlation between nudge treatment effects and price elasticities, which in turn parameterizes the degree of correlation between internalities and price elasticities.

One final consideration is that nudge effectiveness θ is difficult to measure and not addressed in most studies. The standard approach in structural behavioral economics is to assume that the nudge is fully effective ($\theta = 1$), such that its treatment effect can be used to identify the magnitude of the behavioral bias. As with the dependence between b_0 and ε, in our empirical exercises, we compute our welfare estimates under various assumptions about θ and show that our key conclusions do not rely on the particular assumption we make about the value of θ.

3.2 Welfare Formula Implementation Details

Given the framework presented in the previous Subsection, the joint distribution of T, ε, ξ is given by

$$\begin{pmatrix} T \\ |\varepsilon| \\ \xi \end{pmatrix} \sim \log \mathcal{N} \left(\begin{pmatrix} \mu_T \\ \mu_\varepsilon \\ \mu_\xi \end{pmatrix}, \begin{pmatrix} \sigma_T^2 & \rho\sigma_T\sigma_\varepsilon & 0 \\ \rho\sigma_T\sigma_\varepsilon & \sigma_\varepsilon^2 & 0 \\ 0 & 0 & \sigma_\xi^2 \end{pmatrix} \right) \tag{24}$$

where $\rho = [-1, 1]$ is a sensitivity parameter dictating the correlation between the treatment effect and elasticity. Let p be the price of the good and let q be the baseline quantity demanded. Then fixing θ, the bias B is given by $B = \frac{pT}{\theta\varepsilon}$ while slope S is given by $S = \varepsilon\frac{q}{p}$. Applying the definition of the log-normal distribution shows that the joint distribution of B, S, ξ is again a log-normal distribution given by

$$\begin{pmatrix} B \\ |S| \\ \xi \end{pmatrix} \sim \log\mathcal{N}\left(\begin{pmatrix} \mu_B \\ \mu_S \\ \mu_\xi \end{pmatrix}, \begin{pmatrix} \sigma_B^2 & \sigma_{BS} & 0 \\ \sigma_{BS} & \sigma_S^2 & 0 \\ 0 & 0 & \sigma_\xi^2 \end{pmatrix}\right) \tag{25}$$

where $\mu_B = \mu_T - \mu_\varepsilon + \log\frac{p}{\theta}$, $\mu_S = \mu_\varepsilon + \log qp$, $\sigma_B^2 = \sigma_T^2 - 2\rho\sigma_T\sigma_\varepsilon + \sigma_\varepsilon^2$, $\sigma_S^2 = \sigma_\varepsilon^2$, and $\sigma_{BS} = \rho\sigma_T\sigma_\varepsilon - \sigma_\varepsilon^2$.

In Appendix B, we derive closed form expressions for various welfare formulae in terms of the parameters of the log-normal distribution, $\mu_T, \mu_\varepsilon, \mu_\xi, \sigma_T^2, \sigma_\varepsilon^2, \sigma_\xi^2, \rho$ as well as the nudge-effectiveness parameter θ:

$$\begin{aligned}
\Delta_t W(0,0) &= \frac{1}{2}\{\exp\left(2\mu_\xi + \mu_S + \sigma_\xi^2 + \sigma_S^2/2\right) \\
&\quad + \exp(\mu_S + 2\mu_B + [\sigma_S^2 + 4\sigma_{BS} + 2\sigma_B^2]/2) \\
&\quad + \exp(\mu_B + \mu_S + \mu_\xi + [\sigma_B^2 + 2\sigma_{BS} + \sigma_S^2 + \sigma_\xi^2]/2)\} \\
\Delta_n W(0,0) &= \frac{1}{2}\{[1 - (1-\theta)^2]\exp(\mu_S + 2\mu_B + [\sigma_S^2 + 4\sigma_{BS} + 2\sigma_B^2]/2) \\
&\quad + [1 - 2(1-\theta)]\exp(\mu_B + \mu_S + \mu_\xi + [\sigma_B^2 + 2\sigma_{BS} + \sigma_S^2 + \sigma_\xi^2]/2)\} \\
\Delta_{tn} W(0,0) &= \frac{1}{2}\{\exp(2\sigma_\xi + \sigma_S + \sigma_\xi^2 + \sigma_S^2/2) \\
&\quad + [1 - (1-\theta)^2]\exp(\mu_S + 2\mu_B + [\sigma_S^2 + 4\sigma_{BS} + 4\sigma_B^2]/2) \\
&\quad + \exp(\mu_S + 2\mu_B + [\sigma_S^2 + 4\sigma_{BS} + \sigma_B^2]/2)\}
\end{aligned} \tag{26}$$

In the case where biases act as moral taxes rather than as pure debiasing/deception, the welfare effects of nudges instead become

$$\begin{aligned}
\Delta_n W(0,0) &= \frac{1}{2}\{[2\theta]\exp(\mu_S + 2\mu_B + [\sigma_S^2 + 4\sigma_{BS} + 2\sigma_B^2]/2) \\
&\quad + [1 - 2(1-\theta)]\exp(\mu_B + \mu_S + \mu_\xi + [\sigma_B^2 + 2\sigma_{BS} + \sigma_S^2 + \sigma_\xi^2]/2)\} \\
&\quad + \theta\mathbb{E}[D]\exp(\mu_B + \sigma_B^2/2).
\end{aligned} \tag{27}$$

We thus have closed form expressions for the various welfare formulae in terms of the parameters of the log-normal distribution, $\mu_T, \mu_\varepsilon, \mu_\xi, \sigma_T^2, \sigma_\varepsilon^2, \sigma_\xi^2, \rho$ as well as the nudge-effectiveness parameter θ. Recall that while parameters $\mu_T, \mu_\varepsilon, \sigma_T^2, \sigma_\varepsilon^2$ come from our meta-analysis, and parameters μ_ξ, σ_ξ^2 are drawn from the literature estimating externalities, the final two parameters ρ, θ are unknown and hence subjected to a sensitivity analysis. For each value of ρ, θ we consider, the entire vector of parameters are plugged into the above welfare formulae to obtain a single estimate of the effects of the various policies under consideration.

4 Inclusion Criteria and Data Collection

4.1 Inclusion Criteria

Our empirical focus relies on using results from consumer markets wherein our theoretical framework can be easily applied. We therefore refrain from studying topics related to human capital formation, labor supply, and savings behavior, as these applications cannot be captured by our simple framework without additional extensions. Given these considerations, we chose three important markets in which behavioral public policies are ubiquitous and for which a number of nudge and price studies are available: the markets for cigarettes, vaccines, and household electricity.

In addition to the selection of markets, we must define what constitutes a "price intervention" and what constitutes a "nudge". Our definition of price intervention is relatively straightforward and includes any policy that changes the sales price of the underlying good (i.e., of cigarettes, vaccines, and electricity). By contrast, it is more difficult to define precisely a nudge. For instance, one might argue that any non-price intervention could be defined as a nudge. However, such a definition would also imply that medication or therapy sessions with a psychologist (e.g., to reduce nicotine addiction) would be included in our analysis. Such an example would be inconsistent with how the literature typically understands nudges.

To remain consistent with the literature, we follow the spirit of Benartzi et al. (2017) in choosing which interventions are nudges. Thus, we exclude financial incentives, legal mandates, medication, and therapy sessions but include "informational" nudges (e.g., giving consumers information about consumption of peers, effects on health, or billing), "planning" nudges (e.g., goal-setting prompts and offering consumers the ability to pre-commitment to choices), and "streamlining" nudges that simplify tasks (e.g., setting defaults and pre-selecting a choice set). We deviate from Benartzi et al. (2017) in that we count "educational interventions" as nudges. The reason we include such interventions is that there is no clear distinction between what constitutes information provision (a form of nudge in Benartzi et al. 2017) and an educational intervention. For instance, a leaflet that informs about the health risks of smoking may be defined as both an informational nudge and an educational intervention. To avoid this ambiguity, we therefore classified any informational intervention as a nudge.

4.2 Outcome Variables

To connect our meta-analysis to our theory, we must choose outcome variables consistent with the theoretical framework in Section 2. To do so, we needed to make a number of choices about variable definitions. For smoking, we included only papers that addressed cigarette usage. The most commonly used outcome of interest in studies is the probability to *stop* smoking, i.e. a form of extensive margin. For the purposes of our meta-analysis, we therefore focus primarily on this outcome. However, we also have a number of studies that document both the cessation probability *and* changes in number of cigarettes smoked (i.e., the intensive margin) for each subject. We use this information to estimate how the cessation elasticity maps into changes in aggregate cigarette demand. For influenza vaccinations, our outcome of interest is the vaccination probability. We include only papers that address seasonal flu vaccination, not other epidemic strains such as H1N1. For electricity consumption, we include all papers that measure residential electricity use.

For comparison purposes, we also needed to convert point estimates and standard errors of absolute treatment effects into relative percent changes. This conversion required papers to report a number of pieces of information. For the cigarette and vaccine markets, this was straightforward since the outcome variables were binary, so means and sample sizes within different treatment groups suffice to characterize the asymptotic distribution of estimates. For energy nudges and price elasticities, this transformation required additional information, and we excluded papers that did not report sufficient information. For example, we excluded papers that report absolute treatment effects but no control group consumption. More generally, estimates of nudge effectiveness and price elasticity had to include either a standard error or sufficient information to impute a standard error, and we excluded any paper where such imputation was not possible.

4.3 Data Collection

4.3.1 Scraping

Our data on nudges were obtained primarily by scraping Microsoft Academic search results using various keywords.[13] We chose Microsoft Academic because it was easy to scrape and had similarly good coverage as

[13]Unfortunately, Microsoft Academic was discontinued in May 2021.

others potential sources (Martín-Martín et al. 2021). Our scraping procedure yielded titles, reference information, abstracts, and (intermittently) lists of citations and references. After scraping Microsoft Academic, we assigned two research assistants, working independently, to check the first 500 results for each search term, first by title, then by abstract. We reconciled the works of the research assistants by checking them against each other, breaking the tie in the event of a disagreement.

This process left us with an abundance of studies, especially for vaccine and cigarette nudges. Two research assistants then read each paper and, if it was deemed fit for inclusion, independently collected the relevant information —- treatment effects, sample sizes, p-values and other statistics, as well as information on the details of the intervention. Afterwards we reconciled their work, checking them against each other and the paper if there was disagreement.

4.3.2 Leveraging Other Meta-Analyses

The overwhelming majority of papers related to nudging influenza and smoking are from the medical literature. The digitization rate of the medical literature, including older literature is fairly high, so we could be reasonably confident that the scraping procedure described above yielded the relevant sample of papers in these respective literatures. For the literature on energy nudges, by contrast we found that this scraping approach missed relevant papers. We therefore consulted an additional meta-analysis by Delmas, Fischlein and Asensio (2013) as well as another survey of the literature by Darby et al. (2006). We searched through the studies that were cited or mentioned in these two additional sources, and added any that were missed from scraping. Through this procedure, we were able to find 10 additional papers. Of these additional papers, most were old (dating from the first wave of research into behavioral energy reduction following the second oil embargo), grey literature, or had been improperly rejected in an earlier stage of data collection.

There are a variety of high-quality meta-analyses estimating price elasticities. For cigarette price elasticities, we obtained the underyling data of the meta analysis by Gallet and List (2003). Similarly for energy price elasticities, we collected all the papers mentioned in Zhu et al. (2018) and excluded all the estimates of price elasticity that did not meet our inclusion criteria. We were unable to recover a meta-analysis for price elasticities of influenza vaccines. To obtain estimates of such price elasticity estimates, we searched on Microsoft Academic and Google Scholar for all papers estimating the effects of monetary incentives on vaccination take-up. We also estimated price elasticities from policies that made flu shots free to subjects. Because of the difficulties in finding studies on price interventions in the market for vaccines, the sample of studies we were able to obtain in this market is smaller than in the other two markets.

5 Empirical Results

In this section we report welfare effects of nudges and taxes under various scenarios. We start by using our benchmark model in which nudges reduce biases. In Section 5.5, we then discuss how these results change under alternative assumptions where nudges are deceptive or cause direct moral costs or benefits.

5.1 Cigarette Consumption

We begin by plotting the distribution of nudge and price elasticities for cigarettes. Panel A of Figure 2 plots a histogram of point estimates of nudge effects together with the estimated log-normal density. Positive values indicate by how many percent the nudge *increased* the smoking cessation probability. The underlying data includes 53 point estimates, with each point estimate representing a different study. The mean of the point estimates is a decrease in the probability of smoking by 7.5% with a standard deviation of 1.6%. Relative to control, this corresponds to an average increase in the likelihood to quit smoking of 5.6 percentage points. The effect sizes are fairly representative of typical nudge intervention studies in other contexts. As a comparison, the meta-analysis by DellaVigna and Linos (2022) finds an average take-up effect of 8.7 percentage points for academic studies and 1.4 percentage points for studies implemented by "nudge units."

The empirical distribution is left-skewed and has a fairly large standard deviation of 7.5%. The estimated log-normal density provides a good fit for these data. By construction, it has the same mean and standard deviation as noted above, a median of 5.3%, and a mode of 2.7%. Recall that the estimated log-normal accounts for standard errors of the point estimates by deconvolving the empirical distribution from within-study noise. As such, the log-normal distribution does not have to provide a good fit of parts of the empirical distribution where point estimates are very imprecisely estimated. While most treatment effects are below 10%, the distribution has a long right tail. Since its support is restricted to positive values, it cannot predict the few exceptions in which the treatment effect goes into the "unintended" direction—i.e. nudges that increase smoking. None of the negative point estimates are statistically different from zero. We provide complementary plots of ranked treatment effects with standard errors in Appendix C.

The distribution of price elasticities of cigarette demand is plotted in Panel B of Figure 2 and based on 94 point estimates. Positive values indicate a negative price elasticity. The empirical distribution of point estimates has a mean of 0.49 with a standard deviation of 0.14. The standard error of the mean is 0.03, i.e. the average price elasticity is statistically different from zero at conventional levels. The mean implies that, on average, a 10% increase in price reduces cigarette demand by 4.9%. The log-normal population density, again, does an excellent job fitting the data. Mean and standard deviation are identical to the distribution of point estimates, with a median of 0.47 and a mode of 0.43.

A first step in implementing our welfare formulae on the data is understanding how much we would have to tax cigarettes to generate the same effect on aggregate demand as the average nudge. The challenge in answering this question is that most nudge studies measure the intervention effect on a binary outcome, i.e. cessation, while price studies measure the effect on cigarettes demanded. As a consequence, we must transform the effect on the cessation probability into cigarette demand. We do this by exploiting eight nudge studies that measure the effect on both the cessation probability and cigarette consumption. We find that a 1% increase in the cessation probability is associated with a 1.8% decrease in cigarette demand. Based on these data, we assume that the 7.5% reduction in the cessation probability induced by a nudge implies a roughly 13.7% reduction in cigarette demand. Given our elasticity data, this means that prices must increase by roughly 28% to induce the same demand reduction as the average nudge. In the US, this corresponds to a tax of $2.25 per pack. This value is our estimate of the theoretically-derived EPM. It implies that marginal consumers *overvalue* the utility from smoking by $2.25 per pack because they do not fully take into account the associated health consequences.

The collection of these reduced-form effects is of interest alone and offers new insights into the literature on nudges and taxes. However, they are not informative about the efficiency effects of our policy tools. Our theoretical framework allows us to move from reduced-form effects to the quantification of welfare effects. Using Equation (26), we estimate the welfare effects of a nudge, an optimal tax, and a policy mix that uses both the nudge and an optimal tax in combination. We obtain a value for the external damage of smoking by taking the average of externality estimates reported across a number of sources (Sloan et al. 2004, Viscusi 1995, Gruber 2001). This yields an externality value of $0.68 per pack of cigarettes.[14] Panel A of Figure 3 plots these effects for different values of the correlation, ρ, between price elasticities and nudge treatment effects. Under quasi-linear utility, this correlation effectively measures how marginal utility of consumption varies with the nudge effect. We further assume that the nudge eliminates 80% of the bias, i.e., $\theta = 0.8$. As we show in the next graph, our qualitative results are not sensitive to this assumption. The y-axis measures

[14] We assume that one pack has 26 cigarettes.

gains from each policy in USD per consumer per year. Solid lines represent welfare effects while dashed lines are confidence intervals. Table 1 complements the figure (and all following figures) by showing welfare effects for our leading example in which $\rho = 0.5$ and $\theta = 0.8$.

If preferences and nudge effects are independent, $\rho = 0$, then the nudge increases welfare by roughly \$95 USD per consumer annually. This effect is statistically different from zero. With the same correlation, the point estimate of the welfare gain from an optimal tax is roughly 75% of this effect. In fact, we find that for all but the highest values of ρ, the point estimates suggest that nudges dominate taxes. A caveat to this conclusion is that there is large uncertainty in the point estimates, as reflected by the 95% confidence bands.

Another important result to note is that the gains from nudging are decreasing in ρ. To understand this result, first recall that the nudge treatment effect measures by how much consumption deviates from the privately optimal level of consumption. If the treatment effect is larger for more price-elastic consumers, then this means that the deviation is larger for consumers with a low marginal utility of consumption (recall Figure 1). Thus, the behavioral bias distorts consumption where it causes the lowest reduction in consumer surplus. Consumers who are price-inelastic, on the other hand, have smaller treatment effects. While their welfare-loss per unit of consumption is large, the deviation in consumption from the private optimum is relatively low, such that the welfare loss is small. Conversely, if treatment effect and price elasticity are negatively correlated, the deviation in consumption from the private optimum is larger for consumers with high marginal utility of consumption. The bias then distorts behavior where a unit of consumption is most valued. Since we assume that the nudge can mitigate this heterogeneity, the benefits from nudging are larger (smaller) when the correlation is negative (positive). This result is quantitatively important. The gains from nudging increase by over 100% as we change the correlation from 1 to -1.

Gains from cigarette taxation are flat in the correlation. This directly follows from our discussion in Section 2.1.7. The behavioral bias only affects gains from taxation through the average marginal bias, $\mathbb{E}_W[b_0]$ and we showed that this statistic is independent of the correlation between treatment effects and price elasticities: $\mathbb{E}_W[b_0] = \frac{TE}{\mathbb{E}[D']}$. While gains from taxation are, in fact, dependent on the correlation between D' and b, as shown in Equation (17), this does not mean that they are dependent on the correlation between TE and ϵ. Technically, the reason is that varying $\text{Cov}(TE, \epsilon)$ also mechanically varies the *unweighted* average bias in the population, $\mathbb{E}[b_0]$. These two effects exactly offset each other, such that the average marginal bias, $\mathbb{E}_W[b_0]$, in the population remains the same. As a result, the benefits of taxation are completely independent of the correlation between nudge effects and price elasticities.

One key takeaway from this result is that taxes offer an informational benefit over nudges for policy-makers. The evaluation of optimal taxes only requires information on average treatment effects of nudges and prices, while the evaluation of nudges also requires information on the correlation between the two. We are unaware of this point being made in the literature.

The policy mix that combines a nudge and a tax yields slightly larger welfare gains than the nudge in isolation. While the policy mix dominates other policies by theoretical construction, the magnitudes of the welfare effects are not predetermined by theory. It is, therefore, interesting to find that a policy mix is only minimally better than an isolated nudge. Which factors drive this remarkable result? As shown in Equation (19), a nudge is particularly powerful relative to a tax if the standard deviation of the bias is larger than the expected externality of consumption. In economic terms, a nudge dominates the tax when targeting of biases is more important than correcting externalities. In the case of smoking, the expected externality is roughly 0.68 USD per cigarette pack, while the (weighted) standard deviation of the bias is 0.83 USD.

Thus far we have assumed a nudge effectiveness of 80%. We study the sensitivity of our results to this assumption in Panel B of Figure 3. The graph plots welfare effects of the policies for different values of nudge effectiveness, holding the correlation parameter fixed at 0.5. We choose a positive correlation because when behavioral biases and price elasticities are independent, we expect that *treatment effects* of nudges, which in our framework are the product of elasticity with bias, will tend to be positively correlated with elasticity.[15]

As discussed in Section 2, the key factors determining the relative efficiency of nudges compared to taxes are the ability of the nudge to address heterogeneity in biases, the ability of taxes to address externalities, and the effectiveness of the nudge. The impact of these factors factors is summarized by the generalized targeting ratio, defined in Equation (23). In the market for cigarettes, assuming our focal values of $\theta = 0.8$ and $\rho = 0.5$, this ratio is 1.9 and statistically significant, with a bootstrapped standard error of 0.92. Thus, the dominating market failure in the context of cigarette consumption is heterogeneity in bias rather than external damages of consumption. This explains why the nudge dominates the tax and why adding a tax to a nudge only provides a small incremental increase in welfare.

Empirical results in Panel B of Figure 3 illustrate that the gains from any policy are decreasing in nudge effectiveness. To understand this pattern, recall how the behavioral bias is identified in our model (and in

[15] In the case where consumers are rationally inattentive, e.g., following Gabaix (2014), one might expect that biases and elasticities are negatively correlated, which in turn will tend to drive down the correlation between elasticities and treatment effects. However, even in this case, as long as elasticities and biases are not too negatively correlated, we still should expect a positive correlation between elasticities and treatment effects, which is what ρ paramterizes.

the literature more generally). The magnitude of the behavioral bias is increasing (linearly) in the nudge treatment effect. This implies that the bias is larger if the nudge only imperfectly de-biases consumers than if it were fully de-biasing. For example, if a nudge reduces the deviation in consumption from the private optimum by 50%, then the distortion caused by the behavioral bias is twice as large as the treatment effect. Lower values of nudge effectiveness, therefore, imply that the magnitude of the behavioral bias is larger.

In the case of taxation, this mechanically implies that the gains from a corrective tax are larger when nudge effectiveness is lower. In the case of nudging, the directional effect of a change in nudge effectiveness is ambiguous. One the one hand, lower nudge effectiveness implies a larger bias, which increases the benefits of nudging. Alternatively, lower nudge effectiveness implies that the nudge is less powerful at de-biasing, which decreases its benefits. Interestingly, our results suggest that the former effect dominates the latter in the case of cigarette consumption.

We find that for most values of nudge effectiveness, the point estimate of the nudge is above the tax. Only when the nudge becomes very ineffective, at $\theta \approx 0.43$, does the tax slightly dominate the nudge. At this level, the nudge and the tax both increase welfare by around 275 USD per consumer annually. The optimal policy mix raises welfare by more than 381 USD. As we increase the assumed nudge effectiveness, the differences in benefits between policies shrinks. If we use the usual assumption in the behavioral public economics literature that the nudge is fully-debiasing, the isolated nudge generates roughly 93 USD of surplus while the policy mix generates a slightly higher level, 96 USD, of surplus. The isolated tax has a smaller welfare gain of around 68 USD of surplus.

In sum, we find that the potential for welfare gains in the market for cigarettes is larger for nudging than for taxation. This result is robust to a wide range of parameter values for the correlation between preferences and biases, as well as for most levels of nudge effectiveness. The underlying mechanism that generates our results is that heterogeneity in the behavioral bias is more important—from an efficiency perspective—than the external damages from smoking.

5.2 Influenza Vaccinations

Next, we analyze the effects of nudges and taxes on the take-up of influenza vaccinations. In total, we have 36 studies on the effects of nudges on vaccination take-up. Panel A of Figure 4 plots the histogram of nudge treatment effects and the population density functions. The mean point estimate is large and corresponds to a 34.8% increase in the likelihood of vaccination, with a standard error of 6.6%. The standard deviation of the distribution is 32.7%. With an average baseline vaccination probability of 39% in the study samples, the mean effect corresponds to an increase by 13 percentage points. This estimate is statistically significant at conventional levels. Further, the population density provides a good fit for the mean and standard deviation. The median and mode are respectively 22% and 10%.

As mentioned earlier, while we were able to collect a large number of point estimates on nudge effects, we found surprisingly few studies on the effects of price incentives. Most studies focus on behavioral interventions, such as information provision, but rarely implement subsidies that incentivize take-up. One reason for this lack of price studies might be that vaccines are a highly emotionalized topic. Offering money to induce vaccinations might be considered morally reprehensible and could eventually backfire. However, among the point estimates we recovered, there was not a single study that suggested monetary incentives to backfire in the aggregate. Price reductions always increased vaccination take-up. This conclusion is based on a small number of point estimates from 9 studies. The histogram of price elasticities in Panel B of Figure 4 features a mean absolute elasticity of 0.33 and a standard deviation of 0.35. The mean treatment effect is statistically different from zero with a standard error of 0.12.

Given the scant number of point estimates and a wide dispersion, the population density can only imperfectly capture the data. The density has the same mean and standard deviation. The median elasticity is 0.23, the mode is 0.11. These elasticities suggest that to generate the same demand response as the average nudge (+34.8%), subsidies must reduce the price of influenza vaccines by just under 105%. With a typical price of 41 USD for a standard influenza vaccination, this would imply a subsidy of 43 USD. Again, this is the estimated EPM and implies that consumers at the margin undervalue the utility from getting vaccinated by on average 43 USD. From a policy perspective, this result justifies offering flu shots for free *even if vaccines had no positive externality, at all*. Of course, the EPM is not equal to the optimal subsidy level as it does not yet account for the large positive externalities of flu shots.

For our welfare calculations, we assume a positive marginal externality of 153 USD per flu shot. This number is based on the study by White (2021) who estimates an interval of 63 and 243 USD per vaccination. We chose to use the midpoint. One limitation of our model is that it assumes a constant marginal externality, while the marginal externality of a vaccine is likely to fall with the level of vaccinated people in the population. The optimal Pigouvian subsidy would therefore need to be nonlinear, which would complicate the analysis and goes beyond the scope of our exercise. We note, however that most real-world subsidies on influenza vaccines take a linear form.

Panel A of Figure 5 reports welfare effects for a given nudge effectiveness of 80%. For most values of the correlation parameter, subsidies dominate nudges. At $\rho = 1$, the nudge delivers 25 USD of surplus per consumer. Interestingly, at this extreme, rounded to the nearest dollar, both the optimal subsidy and the policy mix deliver 65 USD of surplus, so there are few gains to nudging on top of the optimal subsidy. It is not until $\rho \approx -0.71$ that nudges begin to dominate taxes. Finally, in the extreme case where $\rho = -1$, gains from nudging are estimated to be roughly 87 USD per consumer. The policy mix in this case yields a nontrivial improvement over either nudges or taxes in isolation with a surplus of 126 USD per consumer.

For the more likely case in which nudge effects and price elasticities are positively correlated, subsidies dominate nudges. Panel B of Figure 5 illustrates that this statement is true independent of the nudge effectiveness. As in the case of smoking, we plot welfare effects as a function of nudge effectiveness, while fixing the correlation at 0.5. Estimated benefits from the optimal subsidy decrease from 104 USD to 58 USD as we vary nudge effectiveness from 40% to 100%. The effect is unambiguously decreasing because a larger nudge effectiveness implies a lower behavioral bias, which reduces the benefits of a corrective subsidy. The gains from nudging are always strictly below the gains from subsidizing. With lower benefits of nudging, the advantage of having a policy mix over an isolated subsidy falls, as well.

In conclusion, we obtain a starkly different picture than in the market for cigarettes. First, we find that even with a limited sample on price elasticities, we are able to obtain some clarity about the optimal policy instrument in the market for influenza vaccines. For most values of the correlation parameter, we find that the optimal subsidy is more effective than the nudge at increasing welfare. In fact, in our focal case where $\theta = 0.8$ and $\rho = 0.5$, the optimal policy mix yields a surplus gain of only 4 USD relative to an optimal tax in isolation. Most of the welfare gains in this market can, therefore, be achieved without nudges. Returning to one of our main insights, the generalized targeting ratio within the vaccine market is 0.46 with a standard error of 0.74. This result suggests that the average externality dominates the heterogeneity in the behavioral bias such that the subsidy is the preferred instrument, although this result is subject to some uncertainty due to the wide confidence intervals. Finally, we note that while we view the case where ρ is highly negative as unlikely, the empirical evidence collected in this paper does not allow us to rule out this possibility. This suggests that future studies should aim at documenting this correlation to increase the relevance of their empirical results for optimal policy.

5.3 Energy Consumption

To study behaviorally-motivated policies in the energy market, we collect data on energy conservation nudges and energy price elasticities. We focus on nudges that directly affect the end-use of energy in the household. Examples include social comparison nudges that provide households with information on the energy consumption of their neighbors, as well as real-time feedback on people's energy use. We restrict our analysis to electricity consumption because this is the domain for which we found the largest number of nudge- and price-intervention studies.

In total, we collect point estimates from 35 nudge studies whose distribution we plot in Panel A of Figure 6. The mean treatment effect is a reduction in electricity use by 4.8% with a standard error of 1.3%. The standard deviation of point estimates is 0.9%. There are 3 noisy point estimates that have the unintended sign, i.e. suggest an increase in electricity consumption. None of these estimates is statistically different from zero at conventional levels. We find that most nudge interventions have treatment effects below 10% and the median treatment effect is around 5%. The log-normal population density features the same mean, standard deviation and median, and, overall, provides a good fit for the data.

The distribution of price elasticities is shown in Panel B of Figure 6 and based on 68 studies. The mean elasticity is 0.45 with a tight standard error of 0.05. Thus, a 10% increase in the electricity price causes an average reduction in electricity consumption by 4.5%. This estimate is close to estimates from other meta analyses on electricity price elasticities, such as Labandeira, Labeaga and López-Otero (2017), who estimate a (absolute) long-term elasticity of 0.53. Our estimates suggest that nudges have the same average effect on demand as a tax of 12.7% on the electricity price. This corresponds to a tax of 0.02$/kWh (rounded to the nearest cent) and is, again, our measure of the money-metric bias among consumers at the margin.

Interestingly, almost all price elasticities are smaller than 1 in absolute terms. The distribution of point estimates is notably well represented by the estimated population density function. For our welfare estimation, we assume that one kWh of electricity produces 0.95kg of CO_2 and that the social cost of carbon is 181 USD per ton of CO_2. The externality value is estimated by Hänsel et al. (2020), who use a version of the Dynamic Integrated Climate–Economy (DICE) integrated assessment model together with a wide range of expert views on intergenerational fairness.[16] Our qualitative conclusions remain unchanged if we use substantially lower values, including the 51 USD/tCO_2 that is currently used by the Biden administration for

[16]Based on the same data, we obtain a standard deviation of the social cost of carbon of 186 USD/tCO_2.

cost-benefit analyses (Interagency Working Group 2021).

Panel A of Figure 7 shows welfare effects for our baseline assumption that the nudge eliminates 80% of the behavioral bias. It becomes immediately apparent that regulation in the energy market is fundamentally different from the previous examples. For any correlation between price elasticities and nudge effects, benefits from taxation vastly exceed benefits from nudging. The nudge increases welfare by between $105 and $134 per household annually. The optimal tax raises social welfare by close to $1,000, i.e., roughly 7-9 times more than the nudge. The difference in welfare gains between the two tools is statistically significant at conventional levels. A notable result is that a policy mix of nudges and taxes can barely beat the isolated tax. The only case in which the gains from a policy mix exceed the gains from a isolated tax is in the fairly unlikely situation in which the correlation between preferences and nudge effects is extremely negative. Even when $\rho = -1$, the difference in gains between the policy mix and the isolated tax is only slightly larger than than 1%.

Again, our results can be summarized by computing the generalized targeting ratio in the energy markets. For our focal case of $\theta = 0.8$ and $\rho = 0.5$, we compute a ratio of 0.069 and a standard error of 0.018. Thus, the benefits from externality correction vastly exceed the benefits from targeting the behavioral bias.

Our main results remain virtually unchanged for different values of nudge effectiveness and a given value of the correlation of 0.5. As shown in Panel B of Figure 7, the benefits of taxation range between roughly 947 to 1,115 USD per household as we vary nudge effectiveness. As in the prior examples, lower nudge effectiveness implies that the behavioral bias becomes larger, which increases the benefits of a corrective tax. For very low levels of nudge effectiveness, gains from taxation is around 1,000 USD per household per year.

Nevertheless, gains from nudging remain relatively low, ranging from $105 to $122 per consumer. Interestingly, in the energy market, we find that estimated gains are slightly increasing in nudge effectiveness. Thus, while there is less need for policy intervention for larger values of nudge effectiveness, gains from nudging are still increasing because the policy tool becomes more effective at correcting the bias. However, even for favorable parameter values of nudge effectiveness, we find that benefits from taxation are 802% larger than benefits from nudging. The policy mix in all cases has virtually no additional benefit over the optimal tax.

In the market for electricity, we therefore obtain a fairly unambiguous result that policymakers and academics should focus on the implementation of optimal price regulations rather than on behavioral interventions. Since the benefits of the two policies are extremely different in terms of magnitudes, the opportunity cost of studying and implementing nudges in the energy market appears large. This insight is particularly relevant for current policy debates about optimal environmental policy. Most governments have yet to place a price on carbon, either through a carbon tax or a cap-and-trade system. At the same time, energy conservation nudges have been implemented at an extremely large scale by many governments around the globe. While our results do not imply that these behavioral policies provide no benefits, they do suggest that substantially more attention should be paid to optimal carbon prices.[17]

5.4 Heterogeneity by Nudge Type

Thus far, we have estimated average welfare effects of nudges across all behavioral interventions. This approach is particularly useful to increase statistical power but it may neglect important nuances about different designs of nudges. To investigate heterogeneity, we explore welfare effects for different categories of behavioral interventions. We classify interventions into five groups: i) social norm interventions designed to change subjects' perceived norms about consumption, ii) reminders and feedback about consumption, iii) planning prompts and goal setting nudges that ask subjects to set consumption targets, iv) any other type of information provision, and v) other nudges that cannot be classified as i)-iv). Many studies also design nudges that constitute a combination of i)-v), such as smart meters that provide feedback about energy consumption but also information about the consumption of peer households. As a consequence, we create a sixth category for interventions that use combinations of nudges.

For each category of nudges, we repeat our welfare analysis using point estimates only from the nudge studies corresponding to that particular intervention type. Since we are interested in how nudges compare with taxes, we focus on the generalized targeting ratio from Equation (23). Table 3 reports this statistic for each nudge category in each market, using our baseline case of $\theta = 0.8$ and a correlation of 0.5. Additional columns show the average treatment effect of the each nudge type, its standard deviation, and the number of studies.

In the market for cigarettes, all nudge interventions have targeting ratios above 1, indicating that every type of nudge dominates taxes. However, there is substantial heterogeneity in welfare effects. Within the class of social norm interventions, we find a targeting ratio of close to 2. These interventions come in a number of different forms. One example are web-based tools that provide information to subjects about their

[17]In fact, behavioral interventions may prove useful in reducing the prevalent opposition against a carbon tax in many countries. Rodemeier (2023) finds that informational interventions can raise people's willingness to pay for carbon mitigation substantially.

peer's smoking rate. Another example are video messages sent to mobile phones where community members talk about the process of quitting smoking. A third example are text messages that inform smokers about the public perception of smoking. The average treatment effect of these interventions constitutes an increase in cigarette cessation by 7.4%.

Planning prompts and goal setting nudges reduce demand by 6% and feature a targeting ratio close to 2. Reminders to quit smoking, other informational interventions, and combinations of nudges have similar treatment effects but targeting ratios closer to 1. There are a few other interventions that cannot be categorized and that have average treatment effects of 9%, with an average targeting ratio of 2.5. Overall, it appears that larger average treatment effects broadly predict larger welfare effects in this market.

We obtain a different picture for vaccination nudges, where all targeting ratios are below 1. The most frequent intervention are reminders to get vaccinated and informational campaigns about the benefits of flu shots. Despite causing very different treatment effects (2.6% vs. 23.6%), the targeting ratios of these two nudge types are both around 0.3 and statistically significantly different from 1. Other nudges and multi-nudge interventions cause substantially larger increases in the vaccination probability of around 50% and have larger targeting ratios. The point estimates suggest that even these very effective nudges are dominated by optimal vaccine subsidies. However, we cannot reject that the targeting ratios are equal to 1 statistically, emphasizing the policy uncertainty mentioned earlier in this market.

In the market for electricity consumption, nudges continue to be inferior to taxes across all intervention types. Most nudges are either consumption feedback or multi-nudge interventions that combine social norm comparisons with energy savings tips. Reminders and feedback nudges are often implemented through smart meters, in-home displays, and mobile apps that visualize consumption or savings. The impact of these interventions constitutes a 5.2% reduction in electricity consumption. The targeting ratio is 0.089, implying that energy taxes strongly outperform reminders and feedback nudges.

A common social norm nudge are letters sent to households that provide information about electricity consumption of comparable households. Despite the popularity of these interventions, most of them are not "pure" social norm comparisons but rather use other types of nudges in combination, which is why we classify them as category vi). These multi-nudge interventions reduce electricity consumption by, on average, 4.9%. The targeting ratio is 0.08, i.e., very close zero, and statistically different from 1. Generally, energy taxes continue to dominate nudges irrespective of the nudge intervention type.

To conclude, accounting for heterogeneity provides useful insights into the differential effectiveness of various nudge interventions. Social norm interventions are particularly efficient as a corrective policy to curb smoking. Multi-nudge interventions rank among the most efficient in promoting flu shots and electricity conservation. Our qualitative results as to when nudges dominate taxes, seem to be fairly unaffected by this heterogeneity. Nudges continue to dominate in the cigarette market while the opposite is true for the electricity market. For influenza vaccines results are more mixed but tend to favor vaccine subsidies.

5.5 Alternative Models of Nudges

We consider how our main conclusions change under the alternative specifications in Section 2 in which nudges may be deceptive or impose a moral tax on consumers. Many of the qualitative conclusions from the debiasing model turn out to remain unchanged under these alternative specifications.

5.5.1 Deceptive Nudges

As previously specified, we operationalize nudges that deceive consumers by setting $\theta > 1$ or $\theta < 0$. The case in which $\theta < 0$ would make nudges perform strictly worse than in our main specification with de-biasing. This is because $\theta < 0$ implies that nudges exacerbate the existing bias in the market and additionally increase externalities. Thus, both consumers and society as a whole would be worse off. The more relevant case is $\theta > 1$, meaning the nudge changes the sign of the bias. Consumers are then persuaded to deviate from their private optimum but they are moving consumption closer to the social optimum. For instance, smokers would now overestimate the health costs from smoking. While this induces a loss in consumer surplus, it reduces externalities from passive smoking to the benefit of society as a whole. Deceptive nudges could, therefore, increase welfare beyond the effects of de-biasing nudges.

In Figure 8, we study the welfare implications of this alternative model by varying θ above one. The panels are a continuation of the B-panels in Figures 3, 5, and 7, respectively.

In the cigarette market a number of important results stand out. First, welfare effects of a policy mix intersect with those from the optimal tax at $\theta = 2$. This is a mechanical result that holds true for all three markets because when $\theta = 2$ the absolute value of the bias does not change, nor does the variance of the bias. The implication is that the nudge does not have the comparative advantage of reducing variance in bias in this special case. It simply changes the sign of the bias, such that consumers that initially overvalue cigarettes now undervalue it by the exact same amount. This reduces the size of the optimal tax in a policy mix because deception already internalizes part of the externality. However, such a nudge does not change welfare in a policy mix because the tax alone can always correct the average distortions.

The gains from cigarette nudges are lower when nudges become more deceptive because the variance in the bias increases. Yet, cigarette nudges turn out to dominate taxes for most reasonable values of deception. Only at $\theta \geq 1.83$ do taxes start dominating nudges. At this level, nudges would have to be so deceptive that they almost completely change the bias into the opposite direction. Our qualitative conclusion that nudges dominate taxes in the cigarette market, therefore, holds for a reasonable range of θ in the case of deception.

Finally, note that when nudges become very deceptive, a tax outperforms the policy mix. This follows directly from our prior discussion: a policy mix that involves highly deceptive nudges increase the variance in the bias which causes a large deadweight loss. It is then more efficient to use the tax in isolation instead of combining it with a disortative nudge.

In the market for influenza vaccines, a model of deception makes nudges look worse relative to taxes. While we had some uncertainty in the model where nudges are de-biasing, under deception taxes always outperform nudges. In the energy market, taxes continue to strongly outperform nudges for all plotted parameter values. This is because the treatment effects of nudges are simply too small to internalize a significant part of the large externality from carbon emissions—even for extremely large levels of deception such as $\theta = 3$. Due to the low impact of nudges in this market, the welfare gains from taxation in isolation lie almost perfectly on top of the gains from the policy mix.

Summarizing, a model of deceptive nudges is unlikely to change the qualitative conclusions of the main specification where nudges are debiasing. In fact, deception makes taxes even more attractive than in our baseline model. For very large levels of deception, taxes would dominate nudges in all three markets, including the cigarette market.

5.5.2 Nudges as Moral Taxes

Finally, we consider the model in which nudges impose guilt (or warm glow) about consumption. We find that with the exception of the energy market, our welfare results are fairly sensitive to the normative weight we place on the moral taxation mechanism. Our results, summarized below, can be found in Figure 9.

In the market for cigarettes, we find that while nudges slightly dominate taxes when interpreted as a form of debiasing, they become unambiguously unattractive as a policy tool if they are interpreted entirely as levying a moral tax on consumers. In fact, nudges as moral taxes even imply welfare *losses* relative to laissez faire without any policy intervention. To get the intuition behind this stark difference between the moral tax and the debiasing framework it is helpful to consider a case where there are no externalities. In this case, when $b_0 > 0$ (as it is in the cigarette market), moral taxes move consumption in the "right" direction (a second-order welfare gain), but only at a large psychic cost to consumers (a first-order welfare loss), with no corresponding increase in government revenues, as would have been the case with actual taxes. As a result, especially for small to moderately sized biases, the psychic costs of moral taxes are likely to dominate.

Thinking about nudges as moral subsidy in the market for influenza vaccines leads to a converse effect relative to the market for cigarettes. Now, because biases cause people to under-consume vaccines ($b_0 < 0$), we think of nudges as providing positive "warm glow" utility to consumers. While debiasing nudges could not compete with vaccine subsidies, nudges that induce warm glow are quite competitive, especially for large values of warm glow ($\theta > 1$). However, it is also noteworthy that interpreting nudges through the moral taxation lens substantially increases the standard errors of our estimates, thus creating more policy uncertainty than in the model with debiasing nudges.

Finally, in the market for household electricity, taxes continue to strongly outperform nudges. This result is straightforward since the internality causes an overconsumption of electricity, $b_0 > 0$, and moral taxes then impose an additional feeling of guilt on consumers. Since taxes already dominated both debiasing and deceptive nudges, consumers can only be worse off if nudges impose moral costs. We even find that for any plotted value of θ, moral taxes *decrease* aggregate welfare relative to laissez faire.

In sum, interpreting nudges as moral taxes makes them unambiguously less attractive relative to taxes in the cigarette and energy market, and even implies that they decrease welfare. In the vaccine market, moral subsidy nudges can slightly outperform financial subsidies but also involve substantial policy uncertainty.

6 Conclusion

What can hundreds of studies on the reduced-form effects of nudges teach us about their welfare effects? How do traditional tax policies compare to nudge interventions in terms of economic efficiency? We answer these, and related questions, by linking theory with a meta-analysis to estimate the welfare effects of behavioral public policies. We believe our approach is novel in that we are not aware of any previous studies that leverage meta-analysis in this manner. To operationalize our idea, we derive sufficient statistics of the welfare effects of nudges, taxes, and a policy mix that uses both tools. We show that point estimates of nudge and price effects suffice to quantify efficiency effects for given values of the correlation between preferences and behavioral biases.

To showcase the utility of our approach, we apply our framework to three widely regulated settings in which behavioral biases are allegedly ubiquitous. In our benchmark specification in which nudges debias choices, we find that nudges are more socially efficient in the market for cigarettes, but are far less efficient than taxes in the electricity market. For flu shots, the effects are too noisy to draw definitive conclusions, but subsidies have a slight advantage for reasonable parameter values. A key insight is that two factors govern the difference in results across markets: i) the weighted standard deviation of the behavioral bias and ii) the magnitude of the average externality. Nudges have the unique advantage over taxes in that they potentially reduce the heterogeneity in the behavioral distortion. Taxes, on the other hand, have the advantage of internalizing the marginal externality. Whenever the heterogeneity in bias is large relative to the size of the externality, nudges dominate taxes. This insight highlights a call to researchers to estimate these statistics in their empirical work. Providing such policy-based evidence yields wisdom that usefully guides the optimal design of public policies.

While a combination of nudges and taxes outperforms each policy in isolation, there is large variation in how much combining policy tools add to social welfare. Under certain parameter values, adding a tax/subsidy to a nudge can provide important incremental efficiency gains in the market for cigarettes and influenza vaccines. However, in the electricity market, the additional benefit of adding a nudge to a tax is vanishingly small under virtually all parameter values. This exercise highlights the importance of the empirical quantification of welfare effects.

Our policy recommendations in each market are robust to an alternative model in which nudges are deceptive rather than debiasing. Results are more senstive to a third model in which nudges cause psychic costs or benefits to consumers but often favor taxes over nudges. While ambiguity about the true welfare function has often discouraged research in this area, we hope to show that insigthful welfare evaluations are possible. This requires a more tolerant approach that allows for competing models of normative welfare.

Finally, our empirical analyses relied on assumptions about the correlation between price and nudge effects. While most of our qualitative conclusions are insensitive to these assumptions, quantitative results may change. Our urgent call in this area is for more empirical work to i) distinguish between competing models of nudges, and ii) include both nudge and price interventions in the same sample to estimate the covariance of these effects. Future research can make important contributions in quantifying these important policy parameters.

Tables and Figures

Figure 1: Consumer Surplus Loss from Behavioral Biases

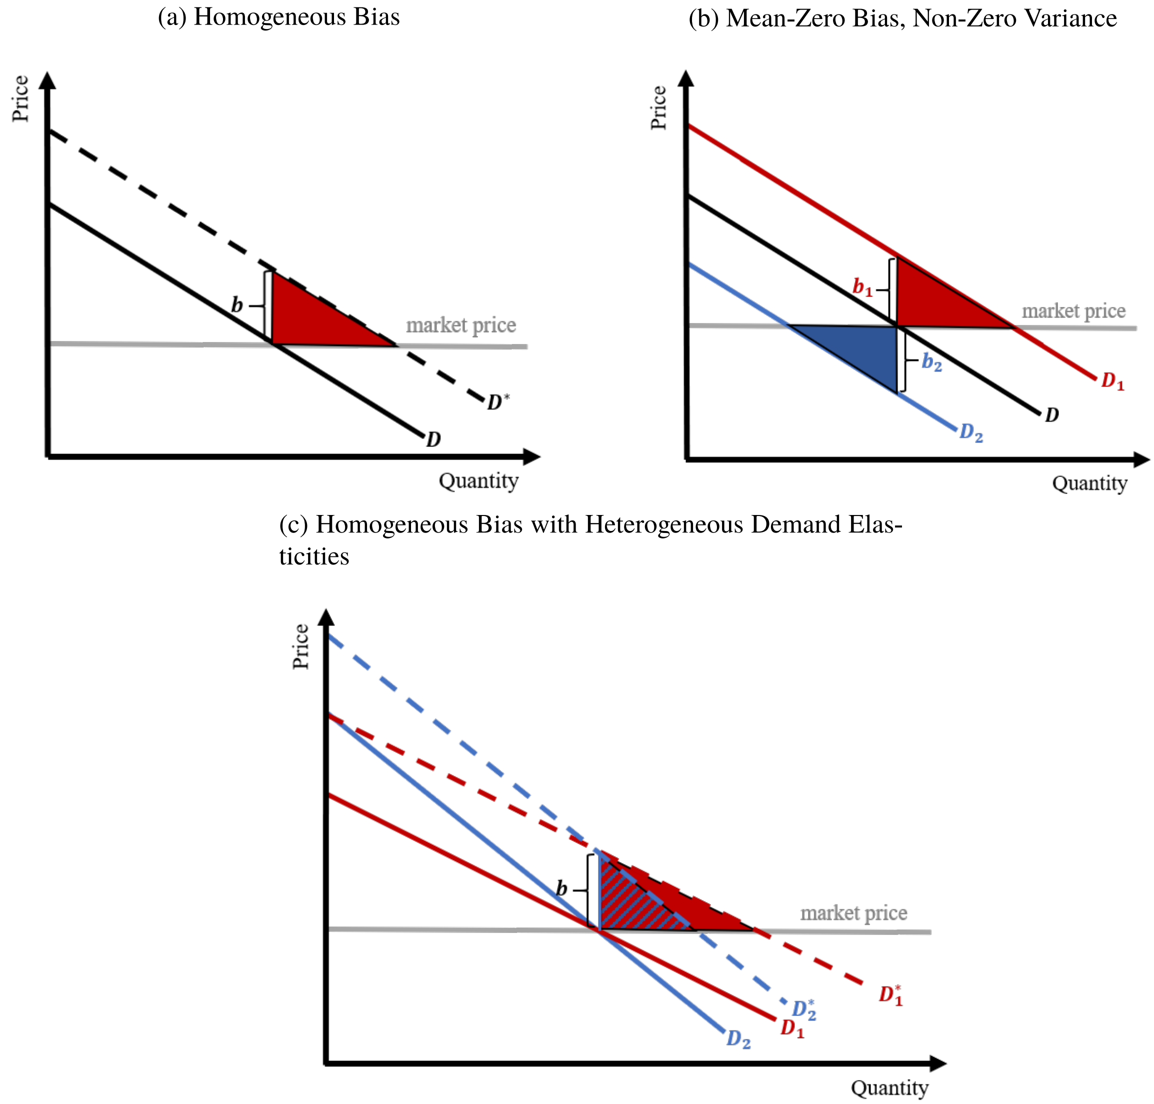

(a) Homogeneous Bias

(b) Mean-Zero Bias, Non-Zero Variance

(c) Homogeneous Bias with Heterogeneous Demand Elasticities

Notes: This figure illustrates examples of the deadweight loss from a behavioral bias. Panel a) shows the deadweight loss from a homogenous behavioral bias of size $b < 0$. D is demand subject to a behavioral bias, while D^* is unbiased demand. The colored triangle measures deadweight loss. Panel b) is a scenario with an average bias of zero but positive variance, coming from two consumer groups: one with a positive bias, $b_1 > 0$, and another with a negative bias, $b_2 < 0$. Panel c) shows the case in which the bias is homogenous but demand elasticities are heterogeneous, resulting in a larger deadweight loss for more price-elastic consumers.

Figure 2: Distribution of Nudge Effects and Price Elasticities for Cigarette Demand

(a) Nudge Effects

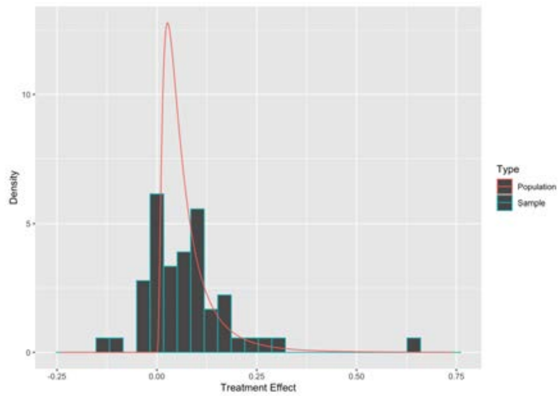

(b) Price Elasticities

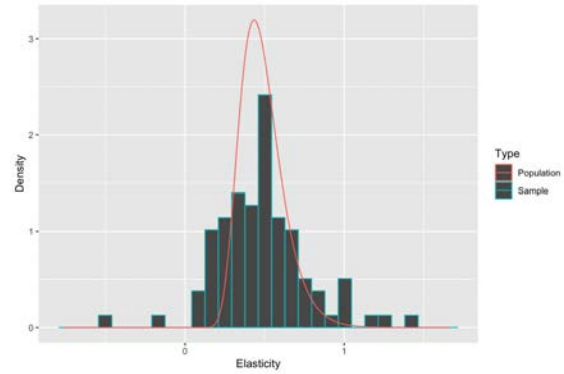

Notes: The figures illustrates the empirical distributions of nudge treatment effects (panel a) and of price elasticities (panel b) in the market for cigarettes. Positive values indicate by how much the intervention *decreased* cigarette consumption. The red line is the estimated log-normal distribution.

Figure 3: Welfare Effects in The Cigarette Market

(a) Fixing $\theta = 0.8$

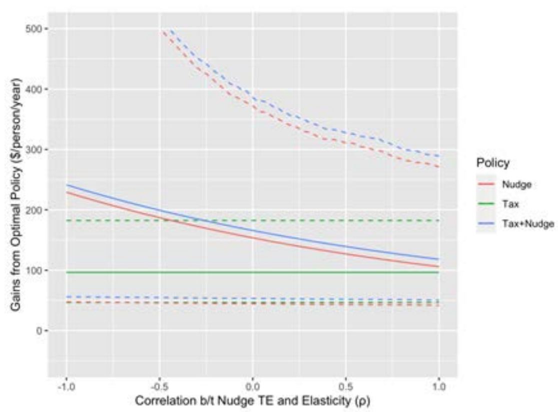

(b) Fixing Correlation $= 0.5$

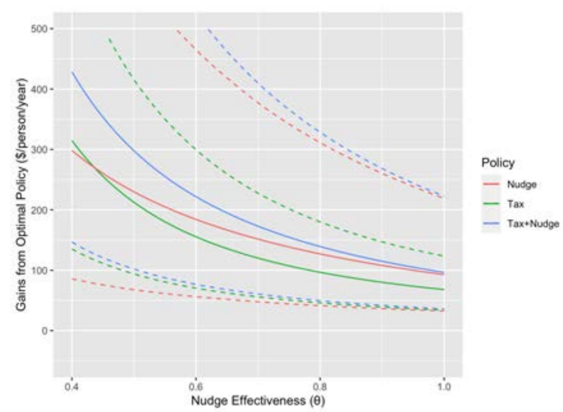

Notes: The figures illustrates welfare effects of nudges, optimal taxes, and a combination of the two policies in the market for cigarettes. Panel a) reports welfare effects for different correlations between nudge treatment effect and price elasticity, while assuming that the nudge is 80% effective in reducing the behavioral bias. Panel b) reports welfare effects for different values of nudge effectiveness, while assuming a correlation between nudge treatment effect and price elasticity of 50%. Dashed lines indicate confidence intervals.

Figure 4: Distribution of Nudge Effects and Price Elasticities for Vaccination Take-Up

(a) Nudge Effects

(b) Price Elasticities

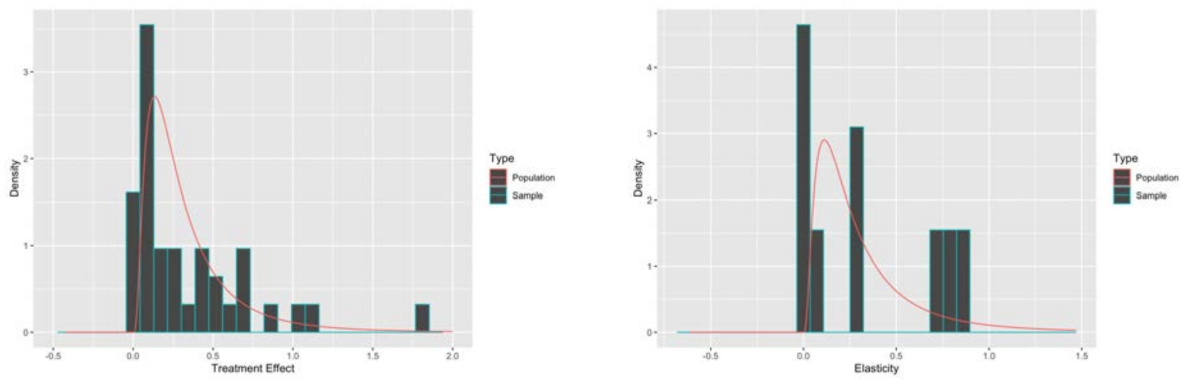

Notes: The figures illustrates the empirical distributions of nudge treatment effects (panel a) and of price elasticities (panel b) in the market for influenza vaccines. Positive values indicate by how much the intervention increased the probability to get vaccinated. The red line is the estimated log-normal distribution.

Figure 5: Welfare Effects in the Market for Vaccines

(a) Fixing $\theta = 0.8$

(b) Fixing Correlation $= 0.5$

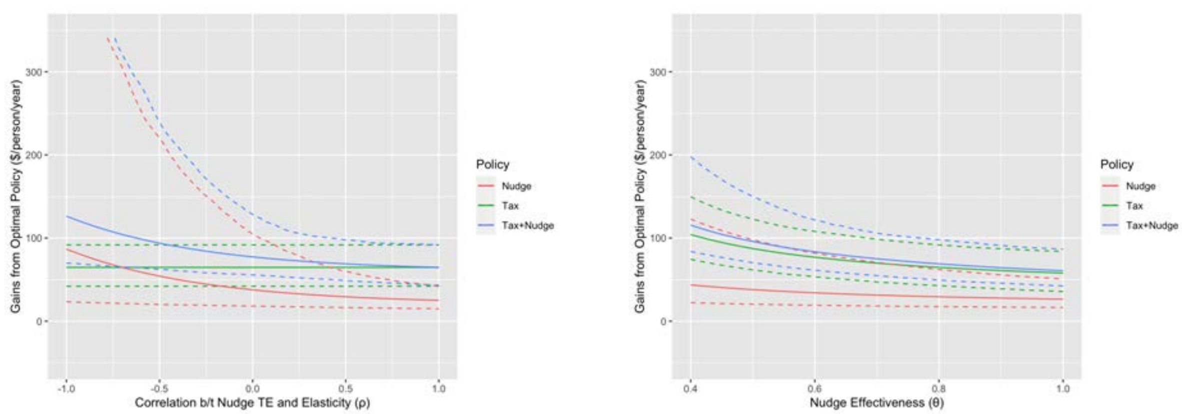

Notes: The figures illustrates welfare effects of nudges, optimal taxes, and a combination of the two policies in the market for influenza vaccines. Panel a) reports welfare effects for different correlations between nudge treatment effect and price elasticity, while assuming that the nudge is 80% effective in reducing the behavioral bias. Panel b) reports welfare effects for different values of nudge effectiveness, while assuming a correlation between nudge treatment effect and price elasticity of 50%. Dashed lines indicate confidence intervals.

Figure 6: Distribution of Nudge Effects and Price Elasticities for Electricity Consumption

(a) Nudge Effects

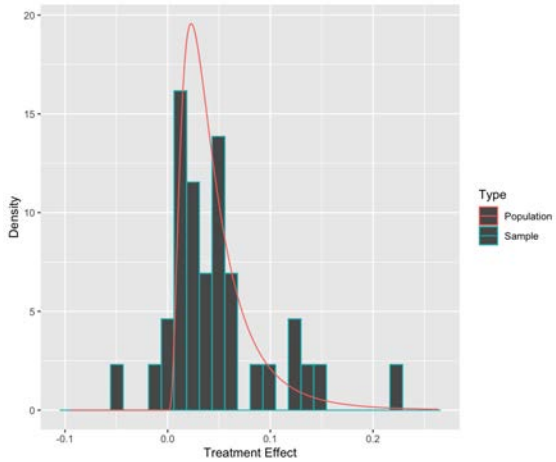

(b) Price Elasticities

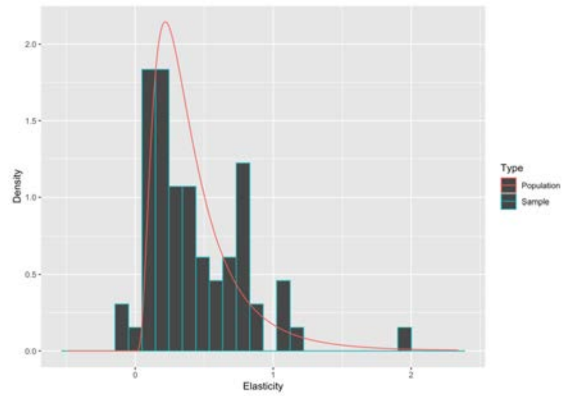

Notes: The figures illustrates the empirical distributions of nudge treatment effects (panel a) and of price elasticities (panel b) in the market for household electricity. Positive values indicate by how much the intervention *decreased* electricity consumption. The red line is the estimated log-normal distribution.

Figure 7: Welfare Effects in the Market for Household Electricity

(a) Fixing $\theta = 0.8$

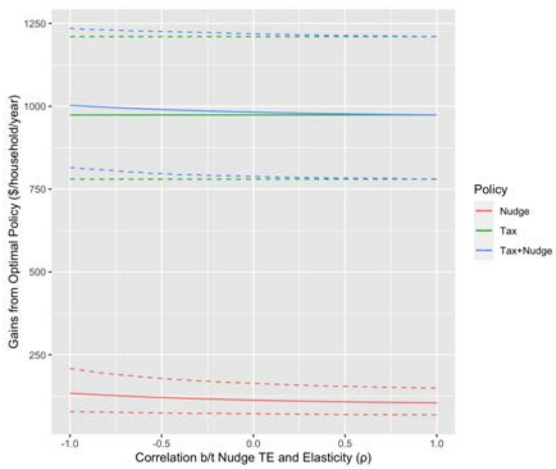

(b) Fixing Correlation $= 0.5$

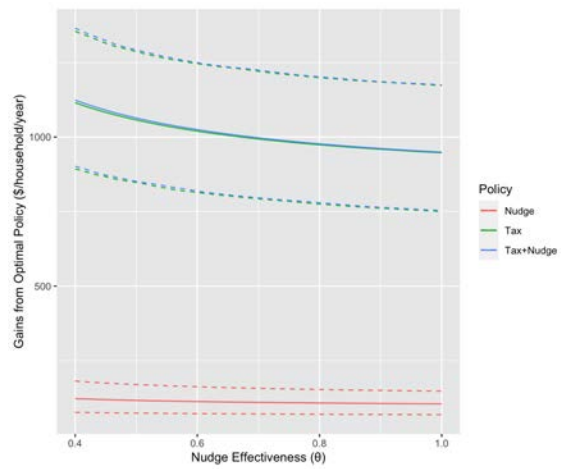

Notes: The figures illustrates welfare effects of nudges, optimal taxes, and a combination of the two policies in the market for household electricity. Panel a) reports welfare effects for different correlations between nudge treatment effect and price elasticity, while assuming that the nudge is 80% effective in reducing the behavioral bias. Panel b) reports welfare effects for different values of nudge effectiveness, while assuming a correlation between nudge treatment effect and price elasticity of 50%. Dashed lines indicate confidence intervals.

Figure 8: Welfare Effects when Nudges are Deceptive

(a) Cigarette Consumption

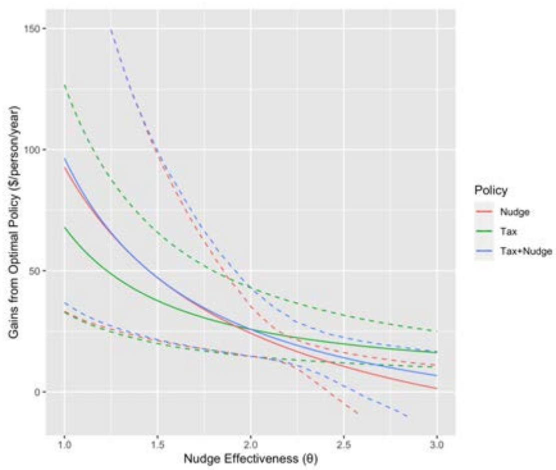

(b) Influenza Vaccines

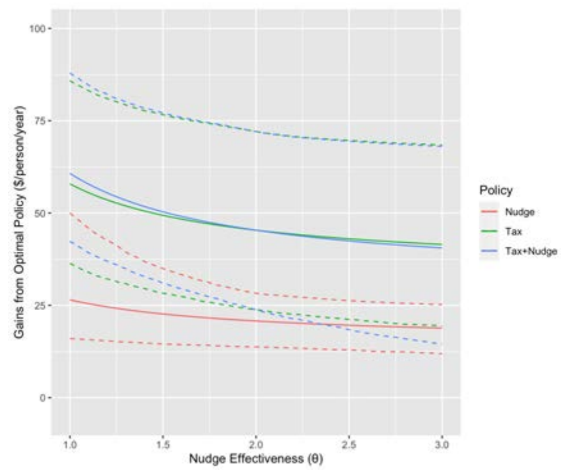

(c) Household Electricity

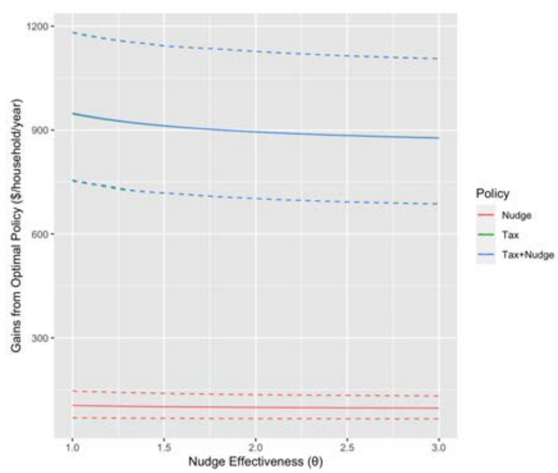

Notes: The figures illustrates welfare effects under the alternative model in which nudges can deceive consumers. A nudge effectiveness above one implies that the bias changes signs. The implication is that under nudging consumers deviate from their privately optimal consumption by smoking too little, being too likely to get vaccinated, and consuming not enough energy. All panels fix the correlation between nudge and price effects at 0.5. Dashed lines indicate confidence intervals.

Figure 9: Debiasing & Deception vs. Moaral Taxation

(a) Cigarettes

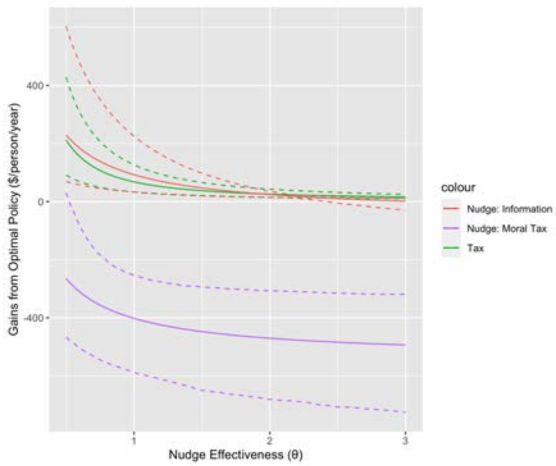

(b) Influenza Vaccines

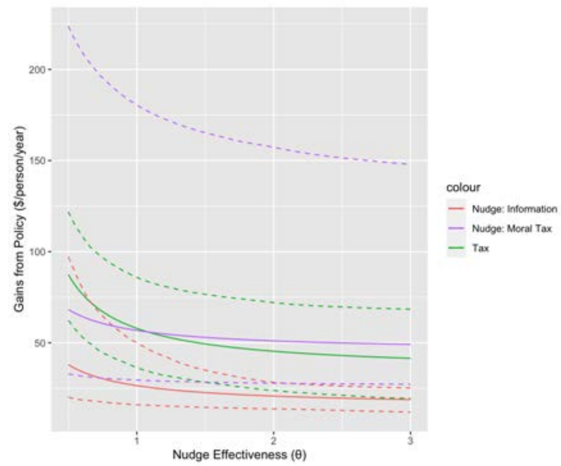

(c) Household Electricity

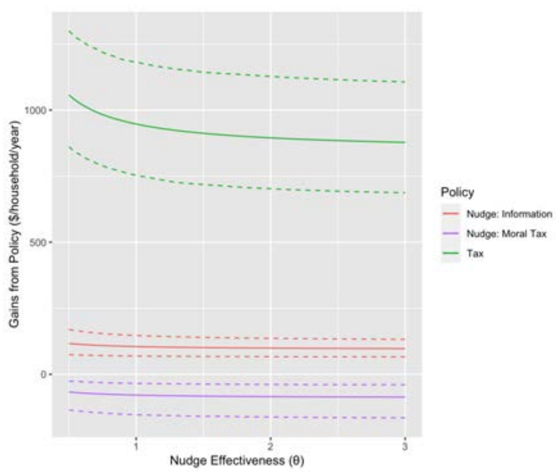

Notes: The figures compares welfare effects of moral taxes to financial taxes, as well as to alternative models of nudges. The purple curve shows the welfare effect of a moral tax. The red curve represents the welfare effect of an information nudge that is either debiasing ($\theta \leq 1$) or deceptive ($\theta > 1$). The green curve shows the welfare effect of a tax. Panel a) displays the cigarette market. Panel b) displays the influenza vaccine market. Panel c) displays the household electricity market. Dashed lines indicate confidence intervals.

Table 1: Welfare Effects, $\theta = 0.8$, $\rho = 0.5$

	Cigarettes (per consumer per year)	Influenza Vaccines (per person per year)	Electricity (per household per year)
Optimal Tax in isolation	\$97 [\$45,\$185]	\$65 [\$42,\$92]	\$974 [\$780, \$1,209]
Nudge in isolation	\$127 [\$42,\$320]	\$29 [\$16,\$60]	\$108 [\$70,\$152]
Nudge and optimal tax in combination	\$139 [\$51,\$338]	\$69 [\$49,\$98]	\$977 [\$785,\$1,210]

Notes: This table reports welfare effects of different policies in the market for cigarettes, influenza vaccines and household electricity. 95%-CIs are in parentheses. The first row shows welfare effects of implementing the optimal tax, while the second row reports welfare effects of using nudges. The final row gives the welfare effects of using both tools in combination. For the estimations, we use our baseline assumptions that the nudge is 80% effective in reducing the behavioral bias, $\theta = 0.8$, and that the correlation between nudge treatment effects and price elasticities is $\rho = 0.5$. See the Figures 3, 5, and 7 for a wide range of alternative assumptions.

Table 2: Optimal Taxes

	Cigarettes (per pack)	Influenza Vaccines (per vaccine)	Electricity (per kWh)
EPM of behavioral bias ($\theta = 1$)	\$2.25 (\$0.50)	-\$43 (\$85)	\$0.01 (\$0.003)
Optimal isolated tax ($\theta = 0.8$)	\$3.49 (\$0.62)	-\$206 (\$107)	\$0.21 (\$0.004)
Optimal tax with nudge ($\theta = 0.8$)	\$1.24 (\$0.12)	-\$164 (\$21)	\$0.19 (\$0.001)
Generalized targeting ratio of nudge ($\theta = 0.8$, $\rho = 0.5$)	1.9 (0.92)	0.46 (0.78)	0.069 (0.018)

Notes: This table reports the equivalent price metric in each market, as well as the size of the optimal tax with and without nudge. The last row indicates the generalized targeting ratio, as defined in Equation (23). Standard errors are in parentheses.

Table 3: Heterogeneity by Nudge Type

	Targeting Ratio ($\theta = 0.8$)	Average Treatment Effect	Std. Dev. of Treatment Effect	Number of Studies
		Cigarettes		
Social Norms Interventions	1.78 (0.78)	-7.4% (3.2%)	4.0% (2.0%)	4
Reminders & Feedback	1.44 (0.51)	-5.8% (2.2%)	2.0% (2.2%)	8
Planning Prompts & Goal Setting	1.86 (0.57)	-6.0% (1.7%)	0.0% (0.9%)	6
Other Information Provision	2.90 (1.01)	-9.5% (3.7%)	14.2% (5.9%)	21
Other	2.50 (0.74)	-8.9% (5.1%)	10.8% (3.6%)	4
Multi-Nudge	1.36 (0.39)	-4.7% (1.7%)	5.0% (2.1%)	13
		Influenza Vaccinations		
Reminders & Feedback	0.343 (0.118)	2.6% (8.8%)	23.5% (12.5%)	11
Other Information Provision	0.282 (0.094)	23.6% (6.1%	6.6% (7.0%)	8
Other	0.726 (0.228)	53.6% (21.3%)	59.0% (16.8%)	8
Multi-Nudge	0.660 (0.230)	51.4% (21.0%)	50.1% (14.8%)	6
		Household Electricity		
Reminders & Feedback	0.089 (0.029)	-5.2% (1.6%)	4.0% (1.2%)	21
Planning Prompts & Goal Setting	0.043 (0.024)	-2.1% (1.0%)	0.0% (0.5%)	3
Other Information Provision	0.116 (0.052)	-5.7% (3.2%)	6.5% (2.5%)	8
Multi-Nudge (including social norms)	0.080 (0.027)	-4.9% (1.3%)	3.2% (2.0%)	29

Notes: This table reports results from a heterogeneity analysis by different categories of nudge intervention. Each row subsets to a different category of nudge and reports the generalized targeting ratio, as defined in Equation (23). Nudges dominate taxes when this value is above 1, and are inferior to taxes when it is below 1. In addition, the table reports the average nudge treatment effect, the standard deviation of treatment effect across studies, and the total number of studies within a nudge category. Standard errors are in parentheses.

References

Allcott, H. (2011), 'Social norms and energy conservation', *Journal of Public Economics* **95**(9–10), 1082–1095.

Allcott, H., Cohen, D., Morrison, W. and Taubinsky, D. (2022), When do ”nudges” increase welfare?, NBER Working Paper 30740.

Allcott, H. and Kessler, J. B. (2019), ‘The welfare effects of nudges: A case study of energy use social comparisons’, *American Economic Journal: Applied Economics* **11**(1), 236–76.

Allcott, H. and Knittel, C. (2019), ‘Are consumers poorly informed about fuel economy? evidence from two experiments’, *American Economic Journal: Economic Policy* **11**(1), 1–37.

Allcott, H., Lockwood, B. B. and Taubinsky, D. (2019), ‘Regressive sin taxes, with an application to the optimal soda tax’, *The Quarterly Journal of Economics* **134**(3), 1557–1626.

Allcott, H., Mullainathan, S. and Taubinsky, D. (2014), ‘Energy policy with externalities and internalities’, *Journal of Public Economics* **112**, 72–88.

Allcott, H. and Rogers, T. (2014), ‘The short-run and long-run effects of behavioral interventions: Experimental evidence from energy conservation’, *The American Economic Review* **104**(10), 3003–3037.

Allcott, H. and Taubinsky, D. (2015), ‘Evaluating behaviorally motivated policy: Experimental evidence from the lightbulb market’, *The American Economic Review* **105**(8), 2501–2538.

Allcott, H. and Wozny, N. (2014), ‘Gasoline prices, fuel economy, and the energy paradox’, *Review of Economics and Statistics* **96**(5), 779–795.

Andor, M., Gerster, A., Peters, J. and Schmidt, C. M. (2018), ‘Social norms and energy conservation beyond the us’, *USAEE Working Paper* .

Antinyan, A. and Asatryan, Z. (2019), Nudging for tax compliance: A meta-analysis, ZEW-Centre for European Economic Research Discussion Paper 19-055.

Armitage, C. J. and Arden, M. A. (2008), ‘How useful are the stages of change for targeting interventions? randomized test of a brief intervention to reduce smoking.’, *Health Psychology* **27**(6), 789.

Barahona, N., Otero, C., Otero, S. and Kim, J. (2023), ‘Equilibrium effects of food labeling policies’, *Econometrica* .

Barseghyan, L., Molinari, F., O’Donoghue, T. and Teitelbaum, J. C. (2013), ‘The nature of risk preferences: Evidence from insurance choices’, *American Economic Review* **103**(6), 2499–2529.

Benartzi, S., Beshears, J., Milkman, K. L., Sunstein, C. R., Thaler, R. H., Shankar, M., Tucker-Ray, W., Congdon, W. J. and Galing, S. (2017), 'Should governments invest more in nudging?', *Psychological science* **28**(8), 1041–1055.

Bernheim, B. D. (2023), The challenges of behavioral welfare analysis, Working Paper.

Bernheim, B. D. and Rangel, A. (2009), 'Beyond revealed preference: Choice-theoretic foundations for behavioral welfare economics', *The Quarterly Journal of Economics* **124**(1), 51–104.

Borland, R., Balmford, J. and Swift, E. (2015), 'Effects of Encouraging Rapid Implementation and/or Structured Planning of Quit Attempts on Smoking Cessation Outcomes: a Randomized Controlled Trial', *Annals of Behavioral Medicine* **49**(5), 732–742.

Butera, L., Metcalfe, R., Morrison, W. and Taubinsky, D. (2022), 'Measuring the welfare effects of shame and pride', *American Economic Review* **112**(1), 122–68.

Chetty, R., Looney, A. and Kroft, K. (2009), 'Salience and Taxation: Theory and Evidence', *The American Economic Review* **99**(4), 1145–1177.

Conlin, M., O'Donoghue, T. and Vogelsang, T. J. (2007), 'Projection bias in catalog orders', *American Economic Review* **97**(4), 1217–1249.

Darby, S. et al. (2006), 'The effectiveness of feedback on energy consumption', *A Review for DEFRA of the Literature on Metering, Billing and direct Displays* **486**(2006).

DellaVigna, S. (2018), Structural behavioral economics, *in* 'Handbook of Behavioral Economics: Applications and Foundations 1', Vol. 1, Elsevier, pp. 613–723.

DellaVigna, S. and Linos, E. (2022), 'Rcts to scale: Comprehensive evidence from two nudge units', *Econometrica* **90**(1), 81–116.

DellaVigna, S., List, J. A. and Malmendier, U. (2012), 'Testing for altruism and social pressure in charitable giving', *The Quarterly Journal of Economics* **127**(1), 1–56.
URL: *http://qje.oxfordjournals.org/content/127/1/1*

DellaVigna, S., List, J. A., Malmendier, U. and Rao, G. (2016), 'Voting to tell others', *The Review of Economic Studies* **84**(1), 143–181.

DellaVigna, S., List, J. A., Malmendier, U. and Rao, G. (2022), 'Estimating social preferences and gift exchange at work', *American Economic Review* **112**(3), 1038–74.

Delmas, M. A., Fischlein, M. and Asensio, O. I. (2013), 'Information strategies and energy conservation behavior: A meta-analysis of experimental studies from 1975 to 2012', *Energy Policy* **61**, 729–739.

Diamond, P. A. (1973), 'Consumption externalities and imperfect corrective pricing', *The Bell Journal of Economics and Management Science* pp. 526–538.

Dubois, P., Griffith, R. and O'Connell, M. (2018), 'The effects of banning advertising in junk food markets', *The Review of Economic Studies* **85**(1), 396–436.

Farhi, E. and Gabaix, X. (2020), 'Optimal taxation with behavioral agents', *American Economic Review* **110**(1), 298–336.

Finkelstein, A. and Hendren, N. (2020), 'Welfare analysis meets causal inference', *Journal of Economic Perspectives* **34**(4), 146–167.

Frank, J. W., McMurray, L. and Henderson, M. (1985), 'Influenza vaccination in the elderly: 2. the economics of sending reminder letters', *Canadian Medical Association Journal* **132**(5), 516.

Gabaix, X. (2014), 'A sparsity-based model of bounded rationality', *The Quarterly Journal of Economics* **129**(4), 1661–1710.

Gallet, C. A. and List, J. A. (2003), 'Cigarette demand: a meta-analysis of elasticities', *Health Economics* **12**(10), 821–835.

Glaeser, E. L. (2006), 'Paternalism and psychology', *The University of Chicago Law Review* **73**(1), 133.

Goldin, J. and Reck, D. (2020), 'Revealed-preference analysis with framing effects', *Journal of Political Economy* **128**(7), 2759–2795.

Goldin, J. and Reck, D. (2022), 'Optimal defaults with normative ambiguity', *Review of Economics and Statistics* **104**(1), 17–33.

Gruber, J. (2001), 'Tobacco at the crossroads: the past and future of smoking regulation in the united states', *Journal of Economic Perspectives* **15**(2), 193–212.

Hänsel, M. C., Drupp, M. A., Johansson, D. J., Nesje, F., Azar, C., Freeman, M. C., Groom, B. and Sterner, T. (2020), 'Climate economics support for the un climate targets', *Nature Climate Change* **10**(8), 781–789.

Harberger, A. C. (1964), 'The measurement of waste', *The American Economic Review* **54**(3), 58–76.

Hedblom, D., Hickman, B. R. and List, J. A. (2019), Toward an understanding of corporate social responsibility: Theory and field experimental evidence, Technical report, National Bureau of Economic Research.

Hendren, N. and Sprung-Keyser, B. (2022), The case for using the mvpf in empirical welfare analysis, Technical report, National Bureau of Economic Research.

Hennrikus, D. J., Lando, H. A., McCarty, M. C., Klevan, D., Holtan, N., Huebsch, J. A., Jestus, S., Pentel, P. R., Pine, D., Sullivan, S., Swenson, K. and Vessey, J. (2005), 'The team project: the effectiveness of smoking cessation intervention with hospital patients', *Preventive Medicine* **40**(3), 249–258.

Houde, S. (2018), 'How consumers respond to product certification and the value of energy information', *The RAND Journal of Economics* **49**(2), 453–477.

Hummel, D. and Maedche, A. (2019), 'How effective is nudging? a quantitative review on the effect sizes and limits of empirical nudging studies', *Journal of Behavioral and Experimental Economics* **80**, 47–58.

Interagency Working Group (2021), Technical support document: Social cost of carbon, methane, and nitrous oxide interim estimates under executive order 13990, Technical report, Interagency Working Group on Social Cost of Greenhouse Gases, United States Government.

Jessoe, K. and Rapson, D. (2014), 'Knowledge is (less) power: Experimental evidence from residential energy use', *American Economic Review* **104**(4), 1417–38.

Labandeira, X., Labeaga, J. M. and López-Otero, X. (2017), 'A meta-analysis on the price elasticity of energy demand', *Energy policy* **102**, 549–568.

Laibson, D., Repetto, A. and Tobacman, J. (2007), Estimating discount functions with consumption choices over the lifecycle, Technical report.

Loewenstein, G. and O'Donoghue, T. (2006), 'We can do this the easy way or the hard way: Negative emotions, self-regulation, and the law', *The University of Chicago Law Review* **73**(1), 183–206.

Löschel, A., Rodemeier, M. and Werthschulte, M. (2022), Can self-set goals encourage resource conservation? field experimental evidence from a smartphone app, Working Paper.

Martín-Martín, A., Thelwall, M., Orduna-Malea, E. and Delgado López-Cózar, E. (2021), 'Google scholar, microsoft academic, scopus, dimensions, web of science, and opencitations' coci: a multidisciplinary comparison of coverage via citations', *Scientometrics* **126**(1), 871–906.

Milkman, K. L., Beshears, J., Choi, J. J., Laibson, D. and Madrian, B. C. (2011), 'Using implementation intentions prompts to enhance influenza vaccination rates', *Proceedings of the National Academy of Sciences* **108**(26), 10415–10420.

Mullainathan, S., Schwartzstein, J. and Congdon, W. J. (2012), 'A reduced-form approach to behavioral public finance', *Annual Review of Economics* **4**(1), 511–540.

Reck, D. and Seibold, A. (2023), The welfare economics of reference dependence, Technical report, National Bureau of Economic Research.

Rodemeier, M. (2020), Buy baits and consumer sophistication: Theory and field evidence from large-scale rebate promotions, CAWM Working Paper 124.

Rodemeier, M. (2023), Willingness to pay for carbon mitigation: Field evidence from the market for carbon offsets, Working paper.

Rodemeier, M. and Löschel, A. (2020), 'The welfare effects of persuasion and taxation: Theory and evidence from the field', *ZEW-Centre for European Economic Research Discussion Paper* (20-019).

Rodemeier, M. and Löschel, A. (2022), Information nudges, subsidies, and crowding out of attention: Field evidence from energy efficiency investments, Working Paper.

Sloan, F. A., Ostermann, J., Conover, C., Taylor Jr, D. H. and Picone, G. (2004), *The price of smoking*, MIT press.

Srinivasan, M., Huntman, J., Nelson, M. and Mathew, S. (2020), 'Use of Peer Comparison, Provider Education, and Electronic Medical Record Triggers to Increase Influenza Vaccination Rates in Hospitalized Children', *Hospital Pediatrics* **10**(1), 76–83.

Stefanski, L. A. and Carroll, R. J. (1990), 'Deconvolving kernel density estimators', *Statistics* **21**(2), 169–184.

Viscusi, W. K. (1995), 'Cigarette taxation and the social consequences of smoking', *Tax policy and the economy* **9**, 51–101.

White, C. (2021), 'Measuring social and externality benefits of influenza vaccination', *Journal of Human Resources* **56**(3), 749–785.

Zhu, X., Li, L., Zhou, K., Zhang, X. and Yang, S. (2018), 'A meta-analysis on the price elasticity and income elasticity of residential electricity demand', *Journal of Cleaner Production* **201**, 169–177.

References on Interventions in the Market for Cigarettes

Abernethy, A. M. and Teel, J. E. (1986), 'Advertising regulation's effect upon demand for cigarettes', *Journal of Advertising* **15**(4), 51–55.

Armitage, C. J. (2007), 'Efficacy of a brief worksite intervention to reduce smoking: The roles of behavioral and implementation intentions.', *Journal of Occupational Health Psychology* **12**(4), 376–390.

Atkinson, A. B. and Skegg, J. L. (1973), 'Anti-smoking publicity and the demand for tobacco in the UK', *The Manchester School* **41**(3), 265–282.

Baltagi, B. H. and Levin, D. (1986), 'Estimating dynamic demand for cigarettes using panel data: The effects of bootlegging, taxation and advertising reconsidered', *The Review of Economics and Statistics* **68**(1), 148–155.

Bardsley, P. and Olekalns, N. (1999), 'Cigarette and tobacco consumption: Have anti-smoking policies made a difference?', *Economic Record* **75**(3), 225–240.

Barnett, P. G., Keeler, T. E. and Hu, T.-w. (1995), 'Oligopoly structure and the incidence of cigarette excise taxes', *Journal of Public Economics* **57**(3), 457–470.

Baskerville, N. B., Struik, L. L., Guindon, G. E., Norman, C. D., Whittaker, R., Burns, C., Hammond, D., Dash, D. and Brown, K. S. (2018), 'Effect of a mobile phone intervention on quitting smoking in a young adult population of smokers: Randomized controlled trial', *JMIR mHealth and uHealth* **6**(10), e10893.

Bass, F. M. (1969), 'A simultaneous equation regression study of advertising and sales of cigarettes', *Journal of Marketing Research* **6**(3), 291–300.

Becker, G. S., Grossman, M. and Murphy, K. M. (1994), 'An empirical analysis of cigarette addiction', *The American Economic Review* **84**(3), 396–418.

Bishop, J. A. and Yoo, J. H. (1985), '"health scare," excise taxes and advertising ban in the cigarette demand and supply', *Southern Economic Journal* **52**(2), 402–411.

Blaine, T. W. and Reed, M. R. (1994), 'US cigarette smoking and health warnings: New evidence from post world war II data', *Journal of Agricultural and Applied Economics* **26**(2), 535–544.

Bolman, C., Eggers, S. M., van Osch, L., Te Poel, F., Candel, M. and de Vries, H. (2015), 'Is action planning helpful for smoking cessation? assessing the effects of action planning in a web-based computer-tailored intervention', *Substance Use & Misuse* **50**(10), 1249–1260.

Borland, R., Balmford, J. and Swift, E. (2015), 'Effects of encouraging rapid implementation and/or structured planning of quit attempts on smoking cessation outcomes: a randomized controlled trial', *Annals of Behavioral Medicine* **49**(5), 732–742.

Boyd, R. and Seldon, B. J. (1990), 'The fleeting effect of advertising: Empirical evidence from a case study', *Economics Letters* **34**(4), 375–379.

Bronson, D. (2001), 'Written, personalised feedback in addition to a standard intervention increased smoking cessation', *Evidence-Based Mental Health* **4**(4), 107–107.

Brown, J., Michie, S., Geraghty, A. W., Yardley, L., Gardner, B., Shahab, L., Stapleton, J. A. and West, R. (2014), 'Internet-based intervention for smoking cessation (StopAdvisor) in people with low and high socioeconomic status: a randomised controlled trial', *The Lancet Respiratory Medicine* **2**(12), 997–1006.

Campbell, R., Starkey, F., Holliday, J., Audrey, S., Bloor, M., Parry-Langdon, N., Hughes, R. and Moore, L. (2008), 'An informal school-based peer-led intervention for smoking prevention in adolescence (ASSIST): a cluster randomised trial', *The Lancet* **371**(9624), 1595–1602.

Chaloupka, F. (1991), 'Rational addictive behavior and cigarette smoking', *Journal of Political Economy* **99**(4), 722–742.

Chaloupka, F. (1992), 'Clean indoor air laws, addiction and cigarette smoking', *Applied Economics* **24**(2), 193–205.

Chaloupka, F. J. and Grossman, M. (1996-09), ‘Price, tobacco control policies and youth smoking’.

Chaloupka, F. J. and Wechsler, H. (1997), ‘Price, tobacco control policies and smoking among young adults’, *Journal of Health Economics* **16**(3), 359–373.

Chansatitporn, N., Charoenca, N., Sidhu, A., Lapvongwatana, P., Kungskulniti, N. and Sussman, S. (2016), ‘Three-month effects of project EX: A smoking intervention pilot program with thai adolescents’, *Addictive Behaviors* **61**, 20–24.

Chetwynd, J., Coope, P., Brodie, R. J. and Wells, E. (1988), ‘Impact of cigarette advertising on aggregate demand for cigarettes in new zealand’, *British Journal of Addiction* **83**(4), 409–414.

Cole-Lewis, H., Augustson, E., Sanders, A., Schwarz, M., Geng, Y., Coa, K. and Hunt, Y. (2017), ‘Analysing user-reported data for enhancement of SmokefreeTXT: a national text message smoking cessation intervention’, *Tobacco Control* **26**(6), 683–689.

Cox, H. and Smith, R. (1984), ‘Political approaches to smoking control: A comparative analysis’, *Applied Economics* **16**(4), 569–582.

Dominguez, L. V. and Page, A. L. (1971), ‘A note on a simultaneous-equation regression study of advertising and sales of cigarettes’, *JMR, Journal of Marketing Research (pre-1986)* **8**(3), 386.

Duffy, M. (1996*a*), ‘Econometric studies of advertising, advertising restrictions and cigarette demand: A survey’, *International Journal of Advertising* **15**(1), 1–23.

Duffy, M. (1996*b*), ‘An econometric study of advertising and cigarette demand in the united kingdom’, *International Journal of Advertising* **15**(3), 262–284.

Emmons, K. M., Puleo, E., Sprunck-Harrild, K., Ford, J., Ostroff, J. S., Hodgson, D., Greenberg, M., Diller, L., de Moor, J. and Tyc, V. (2013), ‘Partnership for health-2, a web-based versus print smoking cessation intervention for childhood and young adult cancer survivors: Randomized comparative effectiveness study’, *Journal of Medical Internet Research* **15**(11), e218.

Espada, J. P., Gonzálvez, M. T., Orgilés, M. and Sussman, S. (2017), ‘One-year effects of project EX: A smoking intervention pilot program with spanish adolescents’, *Journal of Health Psychology* **22**(8), 1067–1074.

Evans, W. N. and Farrelly, M. C. (1998), 'The compensating behavior of smokers: Taxes, tar, and nicotine', *The RAND Journal of Economics* **29**(3), 578–595.

Franke, G. R. (1994), 'U.s. cigarette demand, 1961–1990: Econometric issues, evidence, and implications', *Journal of Business Research* **30**(1), 33–41.

Fujii, E. T. (1980), 'The demand for cigarettes: further empirical evidence and its implications for public policy', *Applied Economics* **12**(4), 479–489.

Galbraith, J. W. and Kaiserman, M. (1997), 'Taxation, smuggling and demand for cigarettes in canada: Evidence from time-series data', *Journal of Health Economics* **16**(3), 287–301.

Gebauer, C., Kwo, C.-Y., Haynes, E. F. and Wewers, M. E. (1998), 'A nurse-managed smoking cessation intervention during pregnancy', *Journal of Obstetric, Gynecologic & Neonatal Nursing* **27**(1), 47–53.

Goel, R. K. and Morey, M. J. (1995), 'The interdependence of cigarette and liquor demand', *Southern Economic Journal* **62**(2), 451–459.

Graham, A. L., Cobb, N. K., Raymond, L., Sill, S. and Young, J. (2007), 'Effectiveness of an internet-based worksite smoking cessation intervention at 12 months', *Journal of Occupational & Environmental Medicine* **49**(8), 821–828.

Grogan, S., Flett, K., Clark-Carter, D., Conner, M., Davey, R., Richardson, D. and Rajaratnam, G. (2011), 'A randomized controlled trial of an appearance-related smoking intervention.', *Health Psychology* **30**(6), 805–809.

Groner, J. A., Ahijevych, K., Grossman, L. K. and Rich, L. N. (2000), 'The impact of a brief intervention on maternal smoking behavior', *Pediatrics* **105**, 267–271.

Haden, K. (1990), 'The demand for cigarettes in japan', *American Journal of Agricultural Economics* **72**(2), 446–450.

Hajek, P., West, R., Lee, A., Foulds, J., Owen, L., Eiser, J. R. and Main, N. (2001), 'Randomized controlled trial of a midwife-delivered brief smoking cessation intervention in pregnancy', *Addiction* **96**(3), 485–494.

Harris, J. E. and Chan, S. W. (1999), 'The continuum-of-addiction: cigarette smoking in relation to price among americans aged 15–29', *Health Economics* **8**(1), 81–86.

Hartmann, K., Thorpjr, J., Pahelshort, L. and Koch, M. (1996), ‘A randomized controlled trial of smoking cessation intervention in pregnancy in an academic clinic’, *Obstetrics & Gynecology* **87**(4), 621–626.

Hennrikus, D., Lando, H., Mccarty, M., Klevan, D., Holtan, N., Huebsch, J., Jestus, S., Pentel, P., Pine, D., Sullivan, S., Swenson, K. and Vessey, J. (2005), ‘The TEAM project: the effectiveness of smoking cessation intervention with hospital patients’, *Preventive Medicine* **40**(3), 249–258.

Herbec, A., Brown, J., Tombor, I., Michie, S. and West, R. (2014), ‘Pilot randomized controlled trial of an internet-based smoking cessation intervention for pregnant smokers (‘MumsQuit’)’, *Drug and Alcohol Dependence* **140**, 130–136.

Hetherington, J., Coutts, R. and Davison, K. (2012), ‘An evaluation of a novel biomarker feedback intervention on smoking cessation: A pilot study’, *Journal of Smoking Cessation* **7**(2), 80–88.

Hishida, A., Terazawa, T., Mamiya, T., Ito, H., Matsuo, K., Tajima, K. and Hamajima, N. (2010), ‘Efficacy of genotype notification to japanese smokers on smoking cessation—an intervention study at workplace’, *Cancer Epidemiology* **34**(1), 96–100.

Holak, S. L. and Reddy, S. K. (1986), ‘Effects of a television and radio advertising ban: A study of the cigarette industry’, *Journal of Marketing* **50**(4), 219–227.

Hsieh, C.-R., Hu, T.-w. and Lin, C.-F. J. (1999), ‘The demand for cigarettes in Taiwan: domestic versus imported cigarettes’, *Contemporary Economic Policy* **17**(2), 223–234.

Ippolito, R. A., Murphy, R. D. and Sant, D. (1979), *Staff report on consumer responses to cigarette health information*, Federal Trade Commission, Bureau of Economics.

Ito, H., Matsuo, K., Wakai, K., Saito, T., Kumimoto, H., Okuma, K., Tajima, K. and Hamajima, N. (2006), ‘An intervention study of smoking cessation with feedback on genetic cancer susceptibility in japan’, *Preventive Medicine* **42**(2), 102–108.

Jackson, J. D. and Saba, R. P. (1997), ‘Some limits on taxing sin: Cigarette taxation and health care finance’, *Southern Economic Journal* **63**(3), 761–775.

Jenkins, C. N., McPhee, S. J., Le, A., Pham, G. Q., Ha, N. T. and Stewart, S. (1997), ‘The effectiveness of a media-led intervention to reduce smoking among vietnamese-american men.', *American Journal of Public Health* **87**(6), 1031–1034.

Johnson, L. W. (1986), 'Advertising expenditure and aggregate demand for cigarettes in australia', *International Journal of Advertising* **5**(1), 45–58.

Jones, A. (1989), 'The UK demand for cigarettes 1954–1986, a double-hurdle approach', *Journal of Health Economics* **8**(1), 133–141.

Kao, K. and Tremblay, V. J. (1988), 'Cigarette "health scare," excise taxes, and advertising ban: Comment', *Southern Economic Journal* **54**(3), 770–776.

Keeler, T. E., Hu, T.-W., Barnett, P. G. and Manning, W. G. (1993), 'Taxation, regulation, and addiction: A demand function for cigarettes based on time-series evidence', *Journal of Health Economics* **12**(1), 1–18.

Kim, S. S., Lee, S. A., Mejia, J., Cooley, M. E. and Demarco, R. F. (2020), 'Pilot randomized controlled trial of a digital storytelling intervention for smoking cessation in women living with HIV', *Annals of Behavioral Medicine* **54**(6), 447–454.

Koutsoyannis, A. P. (1963), 'Demand functions for tobacco', *The Manchester School* **31**(1), 1–19.

Krist, L., Lotz, F., Bürger, C., Ströbele-Benschop, N., Roll, S., Rieckmann, N., Müller-Nordhorn, J., Willich, S. N. and Müller-Riemenschneider, F. (2016), 'Long-term effectiveness of a combined student-parent and a student-only smoking prevention intervention among 7th grade school children in berlin, germany: Evaluation of a combined parent-student smoking prevention intervention', *Addiction* **111**(12), 2219–2229.

Lanoie, P. and Leclair, P. (1998), 'Taxation or regulation: Looking for a good anti-smoking policy', *Economics Letters* **58**(1), 85–89.

Laugesen, M. and Meads, C. (1991), 'Tobacco advertising restrictions, price, income and tobacco consumption in OECD countries, 1960–1986', *British Journal of Addiction* **86**(10), 1343–1354.

Lenert, L., Muñoz, R. F., Perez, J. E. and Bansod, A. (2004), 'Automated e-mail messaging as a tool for improving quit rates in an internet smoking cessation intervention', *Journal of the American Medical Informatics Association* **11**(4), 235–240.

Leu, R. E. (1984), 'Anti-smoking publicity, taxation, and the demand for cigarettes', *Journal of Health Economics* **3**(2), 101–116.

Lewit, E. M. and Coate, D. (1982), 'The potential for using excise taxes to reduce smoking', *Journal of Health Economics* **1**(2), 121–145.

Lewit, E. M., Coate, D. and Grossman, M. (1981), 'The effects of government regulation on teenage smoking', *The Journal of Law & Economics* **24**(3), 545–569.

Li, W. H. C., Ho, K. Y., Wang, M. P., Cheung, D. Y. T., Lam, K. K. W., Xia, W., Cheung, K. Y., Wong, C. K. H., Chan, S. S. C. and Lam, T. H. (2020), 'Effectiveness of a brief self-determination theory–based smoking cessation intervention for smokers at emergency departments in hong kong: A randomized clinical trial', *JAMA Internal Medicine* **180**(2), 206–214.

Liao, Y., Wu, Q., Kelly, B. C., Zhang, F., Tang, Y.-Y., Wang, Q., Ren, H., Hao, Y., Yang, M., Cohen, J. and Tang, J. (2018), 'Effectiveness of a text-messaging-based smoking cessation intervention ("happy quit") for smoking cessation in china: A randomized controlled trial', *PLOS Medicine* **15**(12), e1002713.

Lin, P. R., Zhao, Z. W., Cheng, K.-K. and Lam, T.-H. (2013), 'The effect of physician's 30 s smoking cessation intervention for male medical outpatients: a pilot randomized controlled trial', *Journal of Public Health* **35**(3), 375–383.

Maier, F. H. (1955), 'Consumer demand for cigarettes estimated from state data', *Journal of Farm Economics* **37**(4), 690–704.

McAuliffe, R. (1988), 'The FTC and the effectiveness of cigarette advertising regulations', *Journal of Public Policy & Marketing* **7**(1), 49–64.

Mcbride, C., Bepler, G., Lipkus, I., Lyna, P., Samsa, G., Albright, J., Datta, S. and Rimer, B. (2002), 'Incorporating genetic susceptibility feedback into a smoking cessation program for african-american smokers with low income', *Cancer epidemiology, biomarkers & prevention* **11**(6), 521–8.

McClure, J. B., Peterson, D., Derry, H., Riggs, K., Saint-Johnson, J., Nair, V., An, L. and Shortreed, S. M. (2014), 'Exploring the "active ingredients" of an online smoking intervention: A randomized factorial trial', *Nicotine & Tobacco Research* **16**(8), 1129–1139.

McGuinness, T. and Cowling, K. (1975), 'Advertising and the aggregate demand for cigarettes', *European Economic Review* **6**(3), 311–328.

McLeod, P. B. (1986), 'Advertising bans, tobacco and cigarette consumption', *Economics Letters* **20**(4), 391–396.

Mipatrini, D., Mannocci, A., Pizzi, C. and La Torre, G. (2016), 'School-based anti-smoking intervention for physiotherapy students: a three-year non-randomized trial', *Journal of Preventive Medicine and Hygiene* **57**(2), E91–94.

Mohammed, M., Eggers, S. M., Alotaiby, F. F., de Vries, N. and de Vries, H. (2016), 'Effects of a randomized controlled trial to assess the six-months effects of a school based smoking prevention program in saudi arabia', *Preventive Medicine* **90**, 100–106.

Muckelbauer, R., Englert, H., Rieckmann, N., Chen, C.-M., Wegscheider, K., Völler, H., Katus, H. A., Willich, S. N. and Müller-Nordhorn, J. (2015), 'Long-term effect of a low-intensity smoking intervention embedded in an adherence program for patients with hypercholesterolemia: Randomized controlled trial', *Preventive Medicine* **77**, 155–161.

Müssener, U., Bendtsen, M., Karlsson, N., White, I. R., McCambridge, J. and Bendtsen, P. (2016), 'Effectiveness of short message service text-based smoking cessation intervention among university students: A randomized clinical trial', *JAMA Internal Medicine* **176**(3), 321.

Noonan, D., Silva, S., Njuru, J., Bishop, T., Fish, L. J., Simmons, L. A., Choi, S. H. and Pollak, K. I. (2018), 'Feasibility of a text-based smoking cessation intervention in rural older adults', *Health Education Research* **33**(1), 81–88.

Ojedokun, J. (2013), 'Lung age bio-feedback using a portable lung age meter with brief advice during routine consultations promote smoking cessation ? know2quit multicenter randomized control trial', *Journal of General Practice* **01**(3).

on behalf of the PAADRN Investigators, Wolinsky, F. D., Lou, Y., Edmonds, S. W., Saag, K. G., Roblin, D. W., Wright, N. C., Jones, M. P. and Cram, P. (2017), 'The effects of a patient activation intervention on smoking and excessive drinking cessations: results from the PAADRN randomized controlled trial', *Osteoporosis International* **28**(10), 3055–3060.

Patten, C. A., Koller, K. R., Flanagan, C. A., Hiratsuka, V. Y., Hughes, C. A., Wolfe, A. W., Decker, P. A., Fruth, K., Brockman, T. A., Korpela, M., Gamez, D., Bronars, C., Murphy, N. J., Hatsukami, D., Benowitz, N. L. and Thomas, T. K. (2019), 'Biomarker feedback intervention for smoking cessation among alaska native pregnant women: Randomized pilot study', *Patient Education and Counseling* **102**(3), 528–535.

Pekurinen, M. (1989), 'The demand for tobacco products in finland', *British Journal of Addiction* **84**(10), 1183–1192.

Peto, J. (1974), 'Price and consumption of cigarettes: a case for intervention?', *British journal of preventive & social medicine* **28**(4), 241–245.

Porter, R. H. (1986), 'The impact of government policy on the US cigarette industry', *Empirical approaches to consumer protection economics* pp. 447–481.

Prest, A. R. (1949), 'Some experiments in demand analysis', *The Review of Economics and Statistics* **31**(1), 33–49.

Prochaska, J. O., Velicer, W. F., Fava, J. L., Rossi, J. S. and Tsoh, J. Y. (2001), 'Evaluating a population-based recruitment approach and a stage-based expert system intervention for smoking cessation', *Addictive Behaviors* **26**(4), 583–602.

Quist-Paulsen, P. and Gallefoss, F. (2003), 'Randomised controlled trial of smoking cessation intervention after admission for coronary heart disease', *BMJ* **327**(7426), 1254–1257.

Rajaee, S., Holder, T., Indes, J. E., Muhs, B., Sarac, T., Sumpio, B., Toll, B. A. and Ochoa Chaar, C. I. (2019), 'A pilot study of a standardized smoking cessation intervention for patients with vascular disease', *Annals of Vascular Surgery* **61**, 91–99.e3.

Reekie, W. D. (1994), 'Consumers' surplus and the demand for cigarettes', *Managerial and Decision Economics* **15**(3), 223–234.

Russell, M. (1973), 'Changes in cigarette price and consumption by men in britain, 1946-71: a preliminary analysis.', *British journal of preventive & social medicine* **27**(1), 1.

Sackrin, S. (1962), 'Factors affecting the demand for cigarettes', *Agricultural Economics Research* **14**(1489), 81–88.

Saffer, H. and Chaloupka, F. (1999-02), 'Tobacco advertising: Economic theory and international evidence'.

Schnabel, M. (1972), 'An oligopoly model of the cigarette industry', *Southern Economic Journal* **38**(3), 325–335.

Schneider, L., Klein, B. and Murphy, K. M. (1981), 'Governmental regulation of cigarette health information', *The Journal of Law and Economics* **24**(3), 575–612.

Schoenberg, E. H. (1933), 'The demand curve for cigarettes', *The Journal of Business of the University of Chicago* **6**(1), 15–35.

Seldon, B. J. and Boyd, R. (1991), 'The stability of cigarette demand', *Applied Economics* **23**(2), 319–326.

Seldon, B. J. and Doroodian, K. (1989), 'A simultaneous model of cigarette advertising: Effects on demand and industry response to public policy', *The Review of Economics and Statistics* **71**(4), 673–677.

Shahab, L., West, R. and McNeill, A. (2011), 'A randomized, controlled trial of adding expired carbon monoxide feedback to brief stop smoking advice: Evaluation of cognitive and behavioral effects.', *Health Psychology* **30**(1), 49–57.

Shi, H.-J., Jiang, X.-X., Yu, C.-Y. and Zhang, Y. (2013), 'Use of mobile phone text messaging to deliver an individualized smoking behaviour intervention in chinese adolescents', *Journal of Telemedicine and Telecare* **19**(5), 282–287.

Simmons, V. N., Heckman, B. W., Fink, A. C., Small, B. J. and Brandon, T. H. (2013), 'Efficacy of an experiential, dissonance-based smoking intervention for college students delivered via the internet.', *Journal of Consulting and Clinical Psychology* **81**(5), 810–820.

Simonich, W. L. (1991), *Government antismoking policies*, P. Lang.

Stanczyk, N. E., Smit, E. S., Schulz, D. N., de Vries, H., Bolman, C., Muris, J. W. M. and Evers, S. M. A. A. (2014), 'An economic evaluation of a video- and text-based computer-tailored intervention for smoking cessation: A cost-effectiveness and cost-utility analysis of a randomized controlled trial', *PLoS ONE* **9**(10), e110117.

Stavrinos, V. G. (1987), 'The effects of an anti-smoking campaign on cigarette consumption: empirical evidence from greece', *Applied Economics* **19**(3), 323–329.

Stevens, V. J., Glasgow, R. E., Hollis, J. F., Lichtenstein, E. and Vogt, T. M. (1993), 'A smoking-cessation intervention for hospital patients', *Medical Care* **31**(1), 65–72.

Stone, R. (1945), 'The analysis of market demand', *Journal of the Royal Statistical Society* **108**(3), 286–391.

Sumner, M. T. (1971), 'The demand for tobacco in the U.K.', *The Manchester School* **39**(1), 23–36.

Sung, H.-Y., Hu, T.-W. and Keeler, T. E. (1994), 'Cigarette taxation and demand: an empirical model', *Contemporary Economic Policy* **12**(3), 91–100.

Sutton, S. and Hallett, R. (1988), 'Smoking intervention in the workplace using videotapes and nicotine chewing gum', *Preventive Medicine* **17**(1), 48–59.

Tansel, A. (1993), 'Cigarette demand, health scares and education in turkey', *Applied Economics* **25**(4), 521–529.

Tegene, A. (1991), 'Kalman filter and the demand for cigarettes', *Applied Economics* **23**(7), 1175–1182.

Townsend, J. L. (1987), 'Cigarette tax, economic welfare and social class patterns of smoking', *Applied Economics* **19**(3), 355–365.

Tremblay, C. H. and Tremblay, V. J. (1995), 'The impact of cigarette advertising on consumer surplus, profit, and social welfare', *Contemporary Economic Policy* **13**(1), 113–124.

Unrod, M., Smith, M., Spring, B., DePue, J., Redd, W. and Winkel, G. (2007), 'Randomized controlled trial of a computer-based, tailored intervention to increase smoking cessation counseling by primary care physicians', *Journal of General Internal Medicine* **22**(4), 478–484.

Valdés, B. (1993), 'Cigarette consumption in spain: empirical evidence and implications for public health policy', *Applied Economics* **25**(2), 149–156.

van der Aalst, C. M., de Koning, H. J., van den Bergh, K. A., Willemsen, M. C. and van Klaveren, R. J. (2012), 'The effectiveness of a computer-tailored smoking cessation intervention for participants in lung cancer screening: A randomised controlled trial', *Lung Cancer* **76**(2), 204–210.

Vernon, J. M., Rives, N. W. and Naylor, T. H. (1969), 'An econometric model of the tobacco industry', *The Review of Economics and Statistics* **51**(2), 149–158.

Wangberg, S. C., Nilsen, O., Antypas, K. and Gram, I. T. (2011), 'Effect of tailoring in an internet-based intervention for smoking cessation: Randomized controlled trial', *Journal of Medical Internet Research* **13**(4), e121.

Warner, K. E. (1981), 'Cigarette smoking in the 1970's: The impact of the antismoking campaign on consumption', *Science* **211**(4483), 729–731.

Wasserman, J., Manning, W. G., Newhouse, J. P. and Winkler, J. D. (1991), 'The effects of excise taxes and regulations on cigarette smoking', *Journal of Health Economics* **10**(1), 43–64.

Whittaker, R., Dorey, E., Bramley, D., Bullen, C., Denny, S., Elley, C. R., Maddison, R., McRobbie, H., Parag, V., Rodgers, A. and Salmon, P. (2011), 'A theory-based video messaging mobile phone intervention for smoking cessation: Randomized controlled trial', *Journal of Medical Internet Research* **13**(1), e10.

Winkleby, M. A., Feighery, E., Dunn, M., Kole, S., Ahn, D. and Killen, J. D. (2004), 'Effects of an advocacy intervention to reduce smoking among teenagers', *Archives of Pediatrics & Adolescent Medicine* **158**(3), 269–275.

Wisborg, K., Brink Henriksen, T. and Jørgen Secher, N. (1998), 'A prospective intervention study of stopping smoking in pregnancy in a routine antenatal care setting', *BJOG: An International Journal of Obstetrics & Gynaecology* **105**(11), 1171–1176.

Witt, S. F. and Pass, C. L. (1981), 'The effects of health warnings and advertising on the demand for cigarettes', *Scottish Journal of Political Economy* **28**(1), 86–91.

Yilmaz, G., Karacan, C., Yöney, A. and Yilmaz, T. (2006), 'Brief intervention on maternal smoking: a randomized controlled trial', *Child: Care, Health and Development* **32**(1), 73–79.

Yingst, J. M., Veldheer, S., Hrabovsky, S., Hammett, E., Nicholson, J., Berg, A. and Foulds, J. (2018), 'Pilot randomized trial of an automated smoking cessation intervention via mobile phone text messages as an adjunct to varenicline in primary care', *Journal of Health Communication* **23**(4), 370–378.

Young, T. (1983), 'The demand for cigarettes: alternative specifications of fujii's model', *Applied Economics* **15**(2), 203–211.

Yu, S., Galimov, A., Sussman, S., Jeong, G. C. and Shin, S. R. (2019), ‘Three-month effects of project EX: A smoking intervention pilot program with korean adolescents’, *Addictive Behaviors Reports* **9**, 100152.

References on Interventions in the Market for Influenza Vaccination

Abramson, Z. H., Avni, O., Levi, O. and Miskin, I. N. (2011), ‘Is the influenza vaccination rate of elderly patients affected by raising the vaccination rate of the staff at their primary health care clinics?’, *The Israel Medical Association journal: IMAJ* **13**(6), 325–328.

Arthur, A. J., Matthews, R. J., Jagger, C., Clarke, M., Hipkin, A. and Bennison, D. P. (2002), ‘Improving uptake of influenza vaccination among older people: a randomised controlled trial’, *British Journal of General Practice* **52**(482), 717–8, 720–2.

Bartolo, S., Deliege, E., Mancel, O., Dufour, P., Vanderstichele, S., Roumilhac, M., Hammou, Y., Carpentier, S., Dessein, R., Subtil, D. and Faure, K. (2019), ‘Determinants of influenza vaccination uptake in pregnancy: a large single-centre cohort study’, *BMC Pregnancy and Childbirth* **19**(1), 510.

Barton, M. B. and Schoenbaum, S. C. (1990), ‘Improving influenza vaccination performance in an HMO setting: the use of computer-generated reminders and peer comparison feedback.’, *American Journal of Public Health* **80**(5), 534–536.

Buchan, S. A., Rosella, L. C., Finkelstein, M., Juurlink, D., Isenor, J., Marra, F., Patel, A., Russell, M. L., Quach, S., Waite, N. and Kwong, J. C. (2017), ‘Impact of pharmacist administration of influenza vaccines on uptake in canada’, *Canadian Medical Association Journal* **189**(4), E146–E152.

Buchner, D. M., Larson, E. B. and White, R. F. (1987), ‘Influenza vaccination in community elderly: A controlled trial of postcard reminders’, *Journal of the American Geriatrics Society* **35**(8), 755–760.

Chapman, G. B., Li, M., Colby, H. and Yoon, H. (2010), ‘Opting in vs opting out of influenza vaccination’, *JAMA* **304**(1), 43.

Conner, M., Sandberg, T., Nekitsing, C., Hutter, R., Wood, C., Jackson, C., Godin, G. and Sheeran, P. (2017), ‘Varying cognitive targets and response rates to enhance the question-behaviour effect: An 8-arm randomized controlled trial on influenza vaccination uptake’, *Social Science & Medicine* **180**, 135–142.

Dey, P., Halder, S., Collins, S., Benons, L. and Woodman, C. (2001), ‘Promoting uptake of influenza vaccination among health care workers: a randomized controlled trial’, *Journal of Public Health* **23**(4), 346–348.

Dombkowski, K. J., Cowan, A. E., Reeves, S. L., Foley, M. R. and Dempsey, A. F. (2017), ‘The impacts of email reminder/recall on adolescent influenza vaccination’, *Vaccine* **35**(23), 3089–3095.

Dombkowski, K. J., Harrington, L. B., Dong, S. and Clark, S. J. (2012), ‘Seasonal influenza vaccination reminders for children with high-risk conditions: A registry-based randomized trial’, *American Journal of Preventive Medicine* **42**(1), 71–75.

Esposito, S., Pelucchi, C., Tel, F., Chiarelli, G., Sabatini, C., Semino, M., Marseglia, G. L., De Mattia, D. and Principi, N. (2009), ‘Factors conditioning effectiveness of a reminder/recall system to improve influenza vaccination in asthmatic children’, *Vaccine* **27**(5), 633–635.

Frank, J. W., McMurray, L. and Henderson, M. (1985), ‘Influenza vaccination in the elderly: 2. the economics of sending reminder letters’, *Canadian Medical Association Journal* **132**(5), 516–518, 521.

Goebel, L. J., Neitch, S. M. and Mufson, M. A. (2005), ‘Standing orders in an ambulatory setting increases influenza vaccine usage in older people: Standing orders influenza vaccine’, *Journal of the American Geriatrics Society* **53**(6), 1008–1010.

Goodman, K., Mossad, S. B., Taksler, G. B., Emery, J., Schramm, S. and Rothberg, M. B. (2015), ‘Impact of video education on influenza vaccination in pregnancy’, *The Journal of Reproductive Medicine* **60**(11), 471–479.

Grivas, P. D., Devata, S., Khoriaty, R., Boonstra, P. S., Ruch, J., McDonnell, K., Hernandez-Aya, L., Wilfong, J., Smerage, J., Ison, M. G., Eisenberg, J. N. S., Silveira, M., Cooney, K. A. and Worden, F. P. (2017), ‘Low-cost intervention to increase influenza vaccination rate at a comprehensive cancer center’, *Journal of Cancer Education* **32**(4), 871–877.

Hanchak, N., Murray, J., Harmon-Weiss, S. and Schlackman, N. (1996), ‘The effectiveness of an influenza vaccination program in an HMO setting’, *The American Journal of Managed Care* **2**, 661–6.

Hou, Z., Jie Chang, Yue, D., Fang, H., Meng, Q. and Zhang, Y. (2014), ‘Determinants of willingness to pay for self-paid vaccines in china’, *Vaccine* **32**(35), 4471–4477.

Ibuka, Y. and Bessho, S.-i. (2015), ‘Subsidies for influenza vaccination, vaccination rates, and health outcomes among the elderly in japan’, *Japan and the World Economy* **36**, 56–66.

Ives, D., Lave, J., Traven, N. and Kuller, L. (1994), ‘Impact of medicare reimbursement on influenza vaccination rates in the elderly’, *Preventive Medicine* **23**(2), 134–141.

Jung, Y., Kwon, M. and Song, J. (2017), ‘Stepwise intervention including 1-on-1 counseling is highly effective in increasing influenza vaccination among health care workers’, *American Journal of Infection Control* **45**(6), 635–641.

Kempe, A., Saville, A. W., Albertin, C., Helmkamp, L., Zhou, X., Vangela, S., Dickinson, L. M., Tseng, C.-H., Campbell, J. D., Whittington, M., Gurfinkel, D., Roth, H., Hoefer, D. and Szilagyi, P. (2020), ‘Centralized reminder/recall to increase influenza vaccination rates: A two-state pragmatic randomized trial’, *Academic Pediatrics* **20**(3), 374–383.

Klatt, T. E. and Hopp, E. (2012), ‘Effect of a best-practice alert on the rate of influenza vaccination of pregnant women’, *Obstetrics & Gynecology* **119**(2), 301–305.

Larson, E. B., Bergman, J., Heidrich, F., Alvin, B. L. and Schneeweiss, R. (1982), ‘Do postcard reminders improve influenza vaccination compliance?: A prospective trial of different postcard ”cues”’, *Medical Care* **20**(6), 639–648.

Lehmann, B. (2015), ‘Small decision, big impact: promoting influenza vaccination uptake among health care workers’.

Leitmeyer, K., Buchholz, U., Kramer, M., Schenkel, K., Stahlhut, H., Köllstadt, M., Haas, W. and Meyer, C. (2006), ‘Influenza vaccination in german health care workers: Effects and findings after two rounds of a nationwide awareness campaign’, *Vaccine* **24**(47), 7003–7008.

Leung, K. C., Mui, C., Chiu, W. Y., Ng, Y. Y., Chen, M. H. Y., Ho, P. H., Kwok, C. P., Lam, S. S. M., Wong, C. Y., Wong, K. Y. and Pang, H. H. (2017), ‘Impact of patient education on influenza vaccine uptake among community-dwelling elderly: a randomized controlled trial’, *Health Education Research* **32**(5), 455–464.

MacIntyre, C., Kainer, M. and Brown, G. (2003), 'A randomised, clinical trial comparing the effectiveness of hospital and community-based reminder systems for increasing uptake of influenza and pneumococcal vaccine in hospitalised patients aged 65 years and over', *Gerontology* **49**(1), 33–40.

Margolis, K. L. (1988), 'Organizational strategies to improve influenza vaccine delivery: A standing order in a general medicine clinic', *Archives of Internal Medicine* **148**(10), 2205.

McCreary, L. (2013), 'Increasing the rate of influenza vaccination in children with asthma using a clinic staff and provider educational intervention', *Journal of Asthma & Allergy Educators* **4**(6), 277–281.

McDerby, N. C., Kosari, S., Bail, K. S., Shield, A. J., MacLeod, T., Peterson, G. M. and Naunton, M. (2019), 'Pharmacist-led influenza vaccination services in residential aged care homes: A pilot study', *Australasian Journal on Ageing* **38**(2), 132–135.

McDowell, I., Newell, C. and Rosser, W. (1986), 'Comparison of three methods of recalling patients for influenza vaccination', *CMAJ: Canadian Medical Association journal* **135**(9), 991–997.

Milkman, K. L., Beshears, J., Choi, J. J., Laibson, D. and Madrian, B. C. (2011), 'Using implementation intentions prompts to enhance influenza vaccination rates', *Proceedings of the National Academy of Sciences* **108**(26), 10415–10420.

Mouzoon, M. E., Munoz, F. M., Greisinger, A. J., Brehm, B. J., Wehmanen, O. A., Smith, F. A., Markee, J. A. and Glezen, W. P. (2010), 'Improving influenza immunization in pregnant women and healthcare workers', *The American Journal of Managed Care* **16**(3), 209–216.

Nexøe, J., Kragstrup, J. and Rønne, T. (1997), 'Impact of postal invitations and user fee on influenza vaccination rates among the elderly: A randomized controlled trial in general practice', *Scandinavian Journal of Primary Health Care* **15**(2), 109–112.

Nowalk, M. P., Lin, C. J., Toback, S. L., Rousculp, M. D., Eby, C., Raymund, M. and Zimmerman, R. K. (2010), 'Improving influenza vaccination rates in the workplace: A randomized trial', *American Journal of Preventive Medicine* **38**(3), 237–246.

Ofstead, C. L., Amelang, M. R., Wetzler, H. P. and Tan, L. (2017), 'Moving the needle on nursing staff influenza vaccination in long-term care: Results of an evidence-based intervention', *Vaccine* **35**(18), 2390–2395.

Ogburn, T., Espey, E. L., Contreras, V. and Arroyo, P. (2007), 'Impact of clinic interventions on the rate of influenza vaccination in pregnant women', *The Journal of Reproductive Medicine* **52**(9), 753–756.

Oguz, M. M. (2019), 'Improving influenza vaccination uptake among healthcare workers by on-site influenza vaccination campaign in a tertiary children hospital', *Human Vaccines & Immunotherapeutics* **15**(5), 1060–1065.

Ohkusa, Y. (2005), 'Policy evaluation for the subsidy for influenza vaccination in elderly', *Vaccines and Immunisation. Based on the Fourth World Congress on Vaccines and Immunisation* **23**(17), 2256–2260.

Olanipekun, T., Effoe, V. S., Olanipekun, O., Igbinomwanhia, E., Kola-Kehinde, O., Fotzeu, C., Bakinde, N. and Harris, R. (2020), 'Factors influencing the uptake of influenza vaccination in african american patients with heart failure: Findings from a large urban public hospital', *Heart & Lung* **49**(3), 233–237.

ORIG association, Rothan-Tondeur, M., Filali-Zegzouti, Y., Belmin, J., Lejeune, B., Golmard, J.-L., de Wazières, B., Carrat, F., Piette, F., Mouala, C. and Gavazzi, G. (2010), 'Assessment of healthcare worker influenza vaccination program in french geriatric wards: a cluster-randomized controlled trial', *Aging Clinical and Experimental Research* **22**(5), 450–455.

Patwardhan, A., Kelleher, K., Cunningham, D., Menke, J. and Spencer, C. (2011), 'The use of a mandatory best practice reminder in the electronic record improves influenza vaccination rate in a pediatric rheumatology clinic', *Clinical Governance: An International Journal* **16**(4), 308–319.

Pierson, R. C., Malone, A. M. and Haas, D. M. (2015), 'Increasing influenza vaccination rates in a busy urban clinic', *Journal of Nature and Science* **1**(3), e57.

Poehling, K. A., Fairbrother, G., Zhu, Y., Donauer, S., Ambrose, S., Edwards, K. M., Staat, M. A., Prill, M. M., Finelli, L., Allred, N. J., Bardenheier, B., Szilagyi, P. G. and for the New Vaccine Surveillance Network (2010), 'Practice and child characteristics associated with influenza vaccine uptake in young children', *Pediatrics* **126**(4), 665–673.

Ribner, B. S., Hall, C., Steinberg, J. P., Bornstein, W. A., Chakkalakal, R., Emamifar, A., Eichel, I., Lee, P. C., Castellano, P. Z. and Grossman, G. D. (2008), 'Use of a mandatory declination form in a program for influenza vaccination of healthcare workers', *Infection Control & Hospital Epidemiology* **29**(4), 302–308.

Ruan, M. (2015), 'Factors affecting vaccination demand in the united states', *Summer Program for Undergraduate Research (SPUR)* .

Sartor, C., Tissot-Dupont, H., Zandotti, C., Martin, F., Roques, P. and Drancourt, M. (2004), 'Use of a mobile cart influenza program for vaccination of hospital employees', *Infection Control & Hospital Epidemiology* **25**(11), 918–922.

Satterthwaite, P. (1997), 'A randomised intervention study to examine the effect on immunisation coverage of making influenza vaccine available at no cost', *The New Zealand medical journal* **110**(1038), 58–60.

Schmidtke, K. A., Nightingale, P. G., Reeves, K., Gallier, S., Vlaev, I., Watson, S. I. and Lilford, R. J. (2020), 'Randomised controlled trial of a theory-based intervention to prompt front-line staff to take up the seasonal influenza vaccine', *BMJ Quality & Safety* **29**(3), 189–197.

Shores, D., Wilson, L. and Oliva-Hemker, M. (2019), 'Utilizing information technology to improve influenza vaccination in pediatric patients with inflammatory bowel disease', *Gastroenterology Nursing* **42**(4), 370–374.

Siriwardena, A. N., Rashid, A., Johnson, M., Hazelwood, L. and Wilburn, T. (2003), 'Improving influenza and pneumococcal vaccination uptake in high risk groups in lincolnshire: a quality improvement report from a large rural county', *Quality in Primary Care* **11**(1), 19–28.

Song, J. Y., Park, C. W., Jeong, H. W., Cheong, H. J., Kim, W. J. and Kim, S. R. (2006), 'Effect of a hospital campaign for influenza vaccination of healthcare workers.', *Infection Control and Hospital Epidemiology* **27**(6), 612–617.

Srinivasan, M., Huntman, J., Nelson, M. and Mathew, S. (2020), 'Use of peer comparison, provider education, and electronic medical record triggers to increase influenza vaccination rates in hospitalized children', *Hospital Pediatrics* **10**(1), 76–83.

Stenqvist, K., Hellvin, M.-A. A., Hellke, P., Höglund, D. and von Sydow, H. (2006), 'Influenza work on the regional level in sweden: An integrated program for vaccination of risk groups, surveillance and pandemic planning which focuses on the role of the health care worker', *Vaccine* **24**(44), 6712–6716.

Szilagyi, P. G., Albertin, C. S., Saville, A. W., Valderrama, R., Breck, A., Helmkamp, L., Zhou, X., Vangala, S., Dickinson, L. M., Tseng, C.-H., Campbell, J. D., Whittington, M. D., Roth, H., Rand, C. M., Humiston, S. G., Hoefer, D. and Kempe, A. (2020), 'Effect of state immunization information system based reminder/recall for influenza vaccinations: A randomized trial of autodialer, text, and mailed messages', *The Journal of Pediatrics* **221**, 123–131.e4.

Szilagyi, P. G., Rodewald, L. E., Savageau, J., Yoos, L. and Doane, C. (1992), 'Improving influenza vaccination rates in children with asthma: A test of a computerized reminder system and an analysis of factors predicting vaccination compliance', *Pediatrics* **90**(6), 871–875.

Szilagyi, P. G., Schaffer, S., Rand, C. M., Goldstein, N. P. N., Younge, M., Mendoza, M., Albertin, C. S., Concannon, C., Graupman, E., Hightower, A. D., Yoo, B.-K. and Humiston, S. G. (2019), 'Text message reminders for child influenza vaccination in the setting of school-located influenza vaccination: A randomized clinical trial', *Clinical Pediatrics* **58**(4), 428–436.

Trick, W. E., Das, K., Gerard, M. N., Charles-Damte, M., Murphy, G., Benson, I. and Morita, J. Y. (2009), 'Clinical trial of standing-orders strategies to increase the inpatient influenza vaccination rate', *Infection Control & Hospital Epidemiology* **30**(1), 86–88.

Vayisoglu, S. K. and Zincir, H. (2019), 'The health action process approach-based program's effects on influenza vaccination behavior', *The Journal for Nurse Practitioners* **15**(7), 517–524.

Ward, C. J. (2014), 'Influenza vaccination campaigns: Is an ounce of prevention worth a pound of cure?', *American Economic Journal: Applied Economics* **6**(1), 38–72.

References on Interventions in the Market for Household Energy

Agarwal, S., Rengarajan, S., Sing, T. F. and Yang, Y. (2017), 'Nudges from school children and electricity conservation: Evidence from the "project carbon zero" campaign in singapore', *Energy Economics* **61**, 29–41.

Al-Faris, A. R. F. (2002), 'The demand for electricity in the GCC countries', *Energy Policy* **30**(2), 117–124.

Al-Salman, M. H. (2007), ‘Household demand for energy in kuwait’, *Journal of King Saud University* **19**(1), 51–60.

Alahmad, M. A., Wheeler, P. G., Schwer, A., Eiden, J. and Brumbaugh, A. (2012), ‘A comparative study of three feedback devices for residential real-time energy monitoring’, *IEEE Transactions on Industrial Electronics* **59**(4), 2002–2013.

Alberini, A., Gans, W. and Velez-Lopez, D. (2011), ‘Residential consumption of gas and electricity in the u.s.: The role of prices and income’, *Energy Economics* **33**(5), 870–881.

Alberini, A. and Towe, C. (2015), ‘Information v. energy efficiency incentives: Evidence from residential electricity consumption in maryland’, *Energy Economics* **52**, S30–S40.

Allcott, H. (2011), ‘Social norms and energy conservation’, *Journal of Public Economics* **95**(9), 1082–1095.

Arvola, A., Uutela, A. and Anttila, U. (1993), ‘Billing feedback as means to encourage household electricity conservation: A field experiment in helsinki’, *Proceedings of the 1993 summer study of the European Council for an energy efficient economy* **7585**.

Asadoorian, M. O., Eckaus, R. S. and Schlosser, C. A. (2008), ‘Modeling climate feedbacks to electricity demand: The case of china’, *Energy Economics* **30**(4), 1577–1602.

Asensio, O. I. and Delmas, M. A. (2016), ‘The dynamics of behavior change: Evidence from energy conservation’, *Journal of Economic Behavior & Organization* **126**, 196–212.

Atakhanova, Z. and Howie, P. (2007), ‘Electricity demand in kazakhstan’, *Energy Policy* **35**(7), 3729–3743.

Aydin, E., Brounen, D. and Kok, N. (2018), ‘Information provision and energy consumption: Evidence from a field experiment’, *Energy Economics* **71**, 403–410.

Azevedo, I. M. L., Morgan, M. G. and Lave, L. (2011), ‘Residential and regional electricity consumption in the u.s. and EU: How much will higher prices reduce CO2 emissions?’, *The Electricity Journal* **24**(1), 21–29.

Badri, M. A. (1992), ‘Analysis of demand for electricity in the united states’, *Energy* **17**(7), 725–733.

Becker, L. J. (1978), ‘Joint effect of feedback and goal setting on performance: A field study of residential energy conservation.’, *Journal of Applied Psychology* **63**(4), 428–433.

Bentzen, J. and Engsted, T. (1993), 'Short- and long-run elasticities in energy demand: A cointegration approach', *Energy Economics* **15**(1), 9–16.

Bernstein, M. A. and Griffin, J. (2006), 'Regional differences in the price-elasticity of demand for energy'.

Bernstein, R. and Madlener, R. (2011), 'Responsiveness of residential electricity demand in OECD countries: A panel cointegation and causality analysis'.

Beznoska, M. (2014), 'Estimating a consumer demand system of energy, mobility and leisure: A microdata approach for germany'.

Bianco, V., Manca, O. and Nardini, S. (2009), 'Electricity consumption forecasting in italy using linear regression models', *Energy* **34**(9), 1413–1421.

Blundell, R. and Robin, J. M. (1999), 'Estimation in large and disaggregated demand systems: an estimator for conditionally linear systems', *Journal of Applied Econometrics* **14**(3), 209–232.

Borenstein, S. (2009), 'To what electricity price do consumers respond? residential demand elasticity under increasing-block pricing', *Preliminary Draft April* **30**, 95.

Brandon, G. and Lewis, A. (1999), 'Reducing household energy consumption: A qualitative and quantitative field study', *Journal of Environmental Psychology* **19**(1), 75–85.

Brenton, P. (1997), 'Estimates of the demand for energy using cross-country consumption data', *Applied Economics* **29**(7), 851–859.

Bulunga, A. A. L. and Thondhlana, G. (2018), 'Action for increasing energy-saving behaviour in student residences at rhodes university, south africa', *International Journal of Sustainability in Higher Education* **19**(4), 773–789.

Carroll, J., Lyons, S. and Denny, E. (2014), 'Reducing household electricity demand through smart metering: The role of improved information about energy saving', *Energy Economics* **45**, 234–243.

Casler, S. D. (1997), 'Applied production theory: explicit, flexible, and general functional forms', *Applied Economics* **29**(11), 1483–1492.

Chen, V. L., Delmas, M. A., Locke, S. L. and Singh, A. (2017), 'Information strategies for energy conservation: A field experiment in india', *Energy Economics* **68**, 215–227.

Christensen, T. H., Friis, F., Bettin, S., Throndsen, W., Ornetzeder, M., Skjølsvold, T. M. and Ryghaug, M. (2020-02), 'The role of competences, engagement, and devices in configuring the impact of prices in energy demand response: Findings from three smart energy pilots with households', *Energy Policy* **137**, 111142.

Costa, D. L. and Kahn, M. E. (2013), 'Energy conservation "nudges" and environmentalist ideology: Evidence from a randomized residential electricity field experiment', *Journal of the European Economic Association* **11**(3), 680–702.

Crago, C. L., Spraggon, J. M. and Hunter, E. (2020), 'Motivating non-ratepaying households with feedback and social nudges: A cautionary tale', *Energy Policy* **145**, 111764.

Dahan, A. A. (1996), 'Energy consumption in yemen: Economics and policy (1970-1990)'.

Dergiades, T. and Tsoulfidis, L. (2008), 'Estimating residential demand for electricity in the united states, 1965–2006', *Energy Economics* **30**(5), 2722–2730.

Dolan, P. and Metcalfe, R. D. (2015), 'Neighbors, knowledge, and nuggets: Two natural field experiments on the role of incentives on energy conservation', *SSRN Electronic Journal* .

Durant, I. (1991), 'Residential demand for electricity in barbados, 1966-88', *Central Bank of Barbados Working Papers* **1990**, 287–299.

Erlene Parece, T., Younos, T., Grossman, L. S. and Geller, E. S. (2013), 'A study of environmentally relevant behavior in university residence halls', *International Journal of Sustainability in Higher Education* **14**(4), 466–481.

Fell, H., Li, S. and Paul, A. (2014), 'A new look at residential electricity demand using household expenditure data', *International Journal of Industrial Organization* **33**, 37–47.

Fijnheer, J. D. L., van Oostendorp, H. and Veltkamp, R. C. (2019), Enhancing energy conservation by a household energy game, *in* M. Gentile, M. Allegra and H. Söbke, eds, 'Games and Learning Alliance', Vol. 11385, Springer International Publishing, pp. 257–266.

Filippini, M. (1995), 'Swiss residential demand for electricity by time-of-use', *Resource and Energy Economics* **17**(3), 281–290.

Filippini, M. (1999), 'Swiss residential demand for electricity', *Applied Economics Letters* **6**(8), 533–538.

Filippini, M. (2011), 'Short- and long-run time-of-use price elasticities in swiss residential electricity demand', *Sustainability of biofuels* **39**(10), 5811–5817.

Filippini, M. and Pachauri, S. (2004), 'Elasticities of electricity demand in urban indian households', *Energy Policy* **32**(3), 429–436.

Fouquet, R. (1995), 'The impact of VAT introduction on UK residential energy demand: An investigation using the cointegration approach', *Energy Economics* **17**(3), 237–247.

Fullerton, T. M., Juarez, D. A. and Walke, A. G. (2012), 'Residential electricity consumption in seattle', *Energy Economics* **34**(5), 1693–1699.

Galli, R. (1998), 'The relationship between energy intensity and income levels: Forecasting long term energy demand in asian emerging countries', *The Energy Journal* **19**(4), 85–105.

Ghesla, C., Grieder, M., Schmitz, J. and Stadelmann, M. (2020), 'Pro-environmental incentives and loss aversion: A field experiment on electricity saving behavior', *Energy Policy* **137**, 111131.

Graff Zivin, J. and Novan, K. (2016), 'Upgrading efficiency and behavior: Electricity savings from Residential weatherization programs', *The Energy Journal* **37**(4).

Gölz, S. (2017), 'Does feedback usage lead to electricity savings? analysis of goals for usage, feedback seeking, and consumption behavior', *Energy Efficiency* **10**(6), 1453–1473.

Haakana, M., Sillanpaeae, L. and Talsi, M. (1997), 'The effect of feedback and focused advice on household energy consumption'.

Harding, M. and Hsiaw, A. (2014), 'Goal setting and energy conservation', *Journal of Economic Behavior & Organization* **107**, 209–227.

Harries, T., Rettie, R., Studley, M., Burchell, K. and Chambers, S. (2013), 'Is social norms marketing effective?: A case study in domestic electricity consumption', *European Journal of Marketing* **47**(9), 1458–1475.

Hayes, S. C. and Cone, J. D. (1981), 'Reduction of residential consumption of electricity through simple monthly feedback', *Journal of Applied Behavior Analysis* **14**(1), 81–88.

He, H. and Kua, H. (2013), 'Lessons for integrated household energy conservation policy from singapore's southwest eco-living program', *Energy Policy* **55**, 105–116.

Henry, M. L., Ferraro, P. J. and Kontoleon, A. (2019), 'The behavioural effect of electronic home energy reports: Evidence from a randomised field trial in the united states', *Energy Policy* **132**, 1256–1261.

Houde, S., Todd, A., Sudarshan, A. and Carrie Armel, K. (2013), 'Real-time feedback and electricity consumption: A field experiment assessing the potential for savings and persistence', *The Energy Journal* **34**(1).

Hunter, E. (2016), 'Motivating household energy conservation using feedback and social nudges: A field experiment', *Masters Theses* .

Ito, K., Ida, T. and Tanaka, M. (2018), 'Moral suasion and economic incentives: Field experimental evidence from energy demand', *American Economic Journal: Economic Policy* **10**(1), 240–267.

Kendel, A., Lazaric, N. and Maréchal, K. (2017), 'What do people 'learn by looking' at direct feedback on their energy consumption? results of a field study in southern france', *Energy Policy* **108**, 593–605.

List, J. A., Metcalfe, R. D., Price, M. K. and Rundhammer, F. (2017), 'Harnessing policy complementarities to conserve energy: Evidence from a natural field experiment'.

Loock, C.-M., Staake, T. and Thiesse, F. (2013), 'Motivating energy-efficient behavior with green IS: An investigation of goal setting and the role of defaults', *MIS Quarterly* **37**(4), 1313–1332.

Lynham, J., Nitta, K., Saijo, T. and Tarui, N. (2016), 'Why does real-time information reduce energy consumption?', *Energy Economics* **54**, 173–181.

Matsukawa, I. (2004), 'The effects of information on residential demand for electricity', *The Energy Journal* **25**(1).

Matsukawa, I. (2017), Information feedback from in-home displays and salience effects: Evidence from residential electricity consumption, *in* 'Heading Towards Sustainable Energy Systems: Evolution or Revolution?, 15th IAEE European Conference, Sept 3-6, 2017', International Association for Energy Economics.

Midden, C. J., Meter, J. F., Weenig, M. H. and Zieverink, H. J. (1983), 'Using feedback, reinforcement and information to reduce energy consumption in households: A field-experiment', *Journal of Economic Psychology* **3**(1), 65–86.

Nabeel, M., Ali, B. and Hamdan, A. (2021), The effect of real-time feedback on consumer's behavior in the energy management sector: Empirical study, *in* A. E. Hassanien, R. Bhatnagar and A. Darwish, eds, 'Advanced Machine Learning Technologies and Applications', Vol. 1141, Springer Singapore, pp. 649–660.

Paul, A. C., Myers, E. C. and Palmer, K. L. (2009), 'A partial adjustment model of US electricity demand by region, season, and sector'.

Pellerano, J. A., Price, M. K., Puller, S. L. and Sánchez, G. E. (2017), 'Do extrinsic incentives undermine social norms? evidence from a field experiment in energy conservation', *Environmental and Resource Economics* **67**(3), 413–428.

Peschiera, G., Taylor, J. E. and Siegel, J. A. (2010), 'Response–relapse patterns of building occupant electricity consumption following exposure to personal, contextualized and occupant peer network utilization data', *Energy and Buildings* **42**(8), 1329–1336.

Ro, M., Brauer, M., Kuntz, K., Shukla, R. and Bensch, I. (2017), 'Making cool choices for sustainability: Testing the effectiveness of a game-based approach to promoting pro-environmental behaviors', *Journal of Environmental Psychology* **53**, 20–30.

Romero-Jordán, D., del Río, P. and Peñasco, C. (2014), 'Household electricity demand in spanish regions. public policy implications', *Public Policy Implications (June 18, 2014). IEB Working Paper* (2014/24).

Rubens, L., Le Conte, J., Assegond, C., Fairier, E., Salvazet, R., Bonnefoy, B. and Baud, A.-C. (2017), 'How do french social housing tenants interpret normative descriptive feedback connected with energy? / *¿Cómo interpretan los inquilinos de vivienda social en Francia los mensajes normativos descriptivos respecto al consumo de energía?*', *Psyecology* **8**(3), 323–353.

Schleich, J., Faure, C. and Klobasa, M. (2017), 'Persistence of the effects of providing feedback alongside smart metering devices on household electricity demand', *Energy Policy* **107**, 225–233.

Schultz, P. W., Estrada, M., Schmitt, J., Sokoloski, R. and Silva-Send, N. (2015), 'Using in-home displays to provide smart meter feedback about household electricity consumption: A randomized control trial comparing kilowatts, cost, and social norms', *Energy* **90**, 351–358.

Schultz, P. W., Nolan, J. M., Cialdini, R. B., Goldstein, N. J. and Griskevicius, V. (2007), 'The constructive, destructive, and reconstructive power of social norms', *Psychological Science* **18**(5), 429–434.

Seligman, C. and Darley, J. M. (1977), 'Feedback as a means of decreasing residential energy consumption.', *Journal of Applied Psychology* **62**(4), 363–368.

Shen, M., Young, R. and Cui, Q. (2016), 'The normative feedback approach for energy conservation behavior in the military community', *Energy Policy* **98**, 19–32.

Smith, G. (2013), 'Social norms, social distance, social approval and household electricity consumption: A field experiment in cape town'.

Stinson, J., Willis, A., Williamson, J. B., Currie, J. and Smith, R. (2015), 'Visualising energy use for smart homes and informed users', *Energy Procedia* **78**, 579–584.

Tedenvall, M. (2015), 'Are smart meters really smart?', *IIIEE Master thesis* .

van Dam, S. S., Bakker, C. A. and van Hal, J. D. M. (2010), 'Home energy monitors: impact over the medium-term', *Building Research & Information* **38**(5), 458–469.

Wemyss, D., Cellina, F., Lobsiger-Kägi, E., de Luca, V. and Castri, R. (2019), 'Does it last? long-term impacts of an app-based behavior change intervention on household electricity savings in switzerland', *Energy Research & Social Science* **47**, 16–27.

Winett, R. A., Leckliter, I. N., Chinn, D. E., Stahl, B. and Love, S. Q. (1985), 'Effects of television modeling on residential energy conservation', *Journal of Applied Behavior Analysis* **18**(1), 33–44.

Winett, R. A., Neale, M. S. and Grier, H. C. (1979), 'Effects of self-monitoring and feedback on residential electricity consumption', *Journal of Applied Behavior Analysis* **12**(2), 173–184.

Wisecup, A. K., Grady, D., Roth, R. A. and Stephens, J. (2017), 'A comparative study of the efficacy of intervention strategies on student electricity use in campus residence halls', *International Journal of Sustainability in Higher Education* **18**(4), 503–519.

Yim, D. (2011), Tale of two green communities: Energy informatics and social competition on energy conservation behavior, *in* 'Americas Conference on Information Systems'.

A General Welfare Formla

In this section, we derive the general welfare formula given in equation 21 for arbitrary values of nudge effectiveness, θ. Plugging in $b_1 = (1-\theta)b_0$ into Equation (4), welfare given a nudge and a tax of t is

$$W(t,1) = W^* + \frac{1}{2}\mathbb{E}[D']\mathbb{E}_W\left[(t-((1-\theta)b_0+\xi))^2\right]. \tag{28}$$

Setting the derivative of this expressions with respect to t equal to 0 immediately yields that optimal tax with a nudge is given by $t_1^* = \mathbb{E}_W[(1-\theta)b_0+\xi]$. Plugging this expressions into Equation (28) and rearranging terms yields the expressions in Equation 21.

B Derivation of Welfare Formulae Under Log-Normal Model

In this section, we derive the welfare formulae of equation (26). The following facts about log-normal random variables are used throughout. If $X \sim \log\mathcal{N}(\mu_X, \sigma_X^2)$ and $Y \sim \log\mathcal{N}(\mu_Y, \sigma_Y^2)$, then $\mathbb{E}[X] = \exp\left(\mu + \sigma^2/2\right)$ while $\mathrm{Var}(X) = [\exp(\sigma^2) - 1]\exp(2\mu + \sigma^2)$. Additionally, if $\log X, \log Y$ are correlated with correlation coefficient ρ, then $XY \sim \log\mathcal{N}(\mu_X + \mu_Y, \sigma_X^2 + \sigma_Y^2 + 2\rho\sigma_X\sigma_Y)$.

We are interested ultimately in computing the improvement, relative to no policy intervention of nudges, optimal taxes, or optimal taxes in conjunction with nudges. However, for technical reasons, it will be convenient to begin by computing the deviation of welfare under these policies from the (infeasible) first-best benchmark. Specifically, in the following subsections, we will respectively calculate $W^* - W(0,0)$, $W^* - W(t_0^*, 0)$, $W^* - W(0,1)$, and $W^*(t_1^*, 1)$. The formulae given in Equation (26) then follow from subtracting these various formulae from one another.

B.1 Welfare under Laissez-Faire

If no policy is implemented, the welfare deviation from the first-best is given by

$$-\frac{1}{2}\left(\mathbb{E}[\xi^2]\mathbb{E}[S] + 2\mathbb{E}[BS]\mathbb{E}[\xi] + \mathbb{E}[B^2S]\right). \tag{29}$$

Since monomials of log-normal random variables are again log-normal, we have that B^2S is lognormal with $\mathbb{E}[B^2S] = \exp\left(2\mu_B + \mu_S + [4\sigma_B^2 + 4\sigma_{BS} + \sigma_S^2]/2\right)$. Meanwhile, $\mathbb{E}[\xi^2] = \exp\left(2\mu_\xi + 2\sigma_\xi^2\right)$ and $\mathbb{E}[S] = \exp\left(\mu_S + \sigma_S^2/2\right)$. Putting this all together, we have that

$$\begin{aligned} W(0,0) = \frac{1}{2}\{&\exp\left(2\mu_\xi + \mu_S + [4\sigma_\xi^2 + \sigma_S^2]/2\right) \\ &+ 2\exp(\mu_B + \mu_S + \mu_\xi + [\sigma_B^2 + 2\sigma_{BS} + \sigma_S^2 + \sigma_\xi^2]/2) \\ &+ \exp\left(2\mu_B + \mu_S + [4\sigma_B^2 + 4\sigma_{BS} + \sigma_S^2]/2\right)\} \end{aligned} \tag{30}$$

B.2 Welfare under Optimal Tax in Isolation

Recall that the welfare formula in this case is given by $-\frac{1}{2}\mathbb{E}[S]\mathrm{Var}_S(b+\xi)$. Under the assumed independence between ξ and S, $\mathrm{Var}_S(\xi) = \mathrm{Var}(\xi) = [\exp(\sigma_\xi^2) - 1]\exp(2\mu_\xi + \sigma_\xi^2)$. On the other hand, by definition, we have that

$$\begin{aligned} -\mathbb{E}[S]\mathrm{Var}_S[B] &= \mathbb{E}[SB^2] - \frac{1}{\mathbb{E}[S]}\mathbb{E}[SB]^2 \\ &= \exp\left(\mu_S + 2\mu_B + [\sigma_S^2 + 4\sigma_{BS} + 4\sigma_B^2]/2\right) \\ &\quad - \exp\left(\mu_S + 2\mu_B + [\sigma_S^2 + 4\sigma_{BS} + 2\sigma_B^2]/2\right) \geq 0 \end{aligned} \tag{31}$$

with strict inequality whenever $\sigma_B^2 > 0$. Putting this all together, we have that

$$\begin{aligned} W_t(0,0) = W^* - \frac{1}{2}\{&[\exp(\sigma_\xi^2) - 1]\exp(2\mu_\xi + \mu_S + \sigma_\xi^2 + \sigma_S^2/2) \\ &+ \exp(\mu_S + 2\mu_B + [\sigma_S^2 + 4\sigma_{BS} + 4\sigma_B^2]/2) \\ &- \exp(\mu_S + 2\mu_B + [\sigma_S^2 + 4\sigma_{BS} + 2\sigma_B^2]/2)\} \end{aligned} \tag{32}$$

B.3 Welfare under Information Nudging in Isolation

Recall that, relative to the first-best, the welfare difference of the partial nudge is given by $-\frac{1}{2}\mathbb{E}[S]\mathbb{E}_S[((1-\theta)B+\xi)^2]$. This formula is equal to

$$\frac{1}{2}\left[(1-\theta)^2\mathbb{E}[SB^2] + 2(1-\theta)\mathbb{E}[SB]\mathbb{E}[\xi] + \mathbb{E}[S]\mathbb{E}[\xi^2]\right] \tag{33}$$

As before, each of the expectations above are of log-normal random variables, so we have

$$
\begin{aligned}
W_n(0,0) = W^* - \frac{1}{2}\{&(1-\theta)^2 \exp(\mu_S + 2\mu_B + [\sigma_S^2 + 4\sigma_{BS} + 4\sigma_B^2]/2) \\
&+2(1-\theta)\exp(\mu_S + \mu_B + \mu_\xi + [\sigma_S^2 + 2\sigma_{BS} + \sigma_B^2 + \sigma_\xi^2]/2) \\
&+\exp(\mu_S + 2\mu_\xi + [\sigma_S^2 + 4\sigma_\xi^2]/2)\}
\end{aligned} \tag{34}
$$

B.4 Welfare under Moral Taxation Nudging in Isolation

Equation (11) and (14) differ only in terms of the coefficient in front of b_0^2 and the extra term, $\mathbb{E}[\theta D b_0]$. Assuming that D is uncorrelated to b_0, this term is given by $\theta\mathbb{E}[D]\exp(\mu_B + \sigma_B^2/2)$, given our log-normal paramterization. The resulting expression for $W_n(0,0)$ is therefore given by

$$
\begin{aligned}
W_n(0,0) = W^* - \frac{1}{2}\{&(1-2\theta)\exp(\mu_S + 2\mu_B + [\sigma_S^2 + 4\sigma_{BS} + 4\sigma_B^2]/2) \\
&+2(1-\theta)\exp(\mu_S + \mu_B + \mu_\xi + [\sigma_S^2 + 2\sigma_{BS} + \sigma_B^2 + \sigma_\xi^2]/2) \\
&+\exp(\mu_S + 2\mu_\xi + [\sigma_S^2 + 4\sigma_\xi^2]/2)\} \\
&+\theta\mathbb{E}[D]\exp(\mu_B + \sigma_B^2/2)
\end{aligned} \tag{35}
$$

B.5 Welfare under Policy Mix

The welfare formula in this case is given by $W_{nt}(0,0) = W^* - \frac{1}{2}\mathbb{E}[S]\mathrm{Var}((1-\theta)B + \xi)$. Rearranging the calculations already done for the isolated tax case, we have that

$$
\begin{aligned}
W_{tn}(0,0) = W^* - \frac{1}{2}\{&[\exp(\sigma_\xi^2) - 1]\exp(2\mu_\xi + \mu_S + \sigma_\xi^2 + \sigma_S^2/2) \\
&+(1-\theta)^2[\exp(\mu_S + 2\mu_B + [\sigma_S^2 + 4\sigma_{BS} + 4\sigma_B^2]/2) \\
&-\exp(\mu_S + 2\mu_B + [\sigma_S^2 + 4\sigma_{BS} + 2\sigma_B^2]/2)]\}
\end{aligned} \tag{36}
$$

C Point Estimates and Standard Errors for Each Study

Figure 10: Nudge Treatment Effects on Smoking Cessation Probability

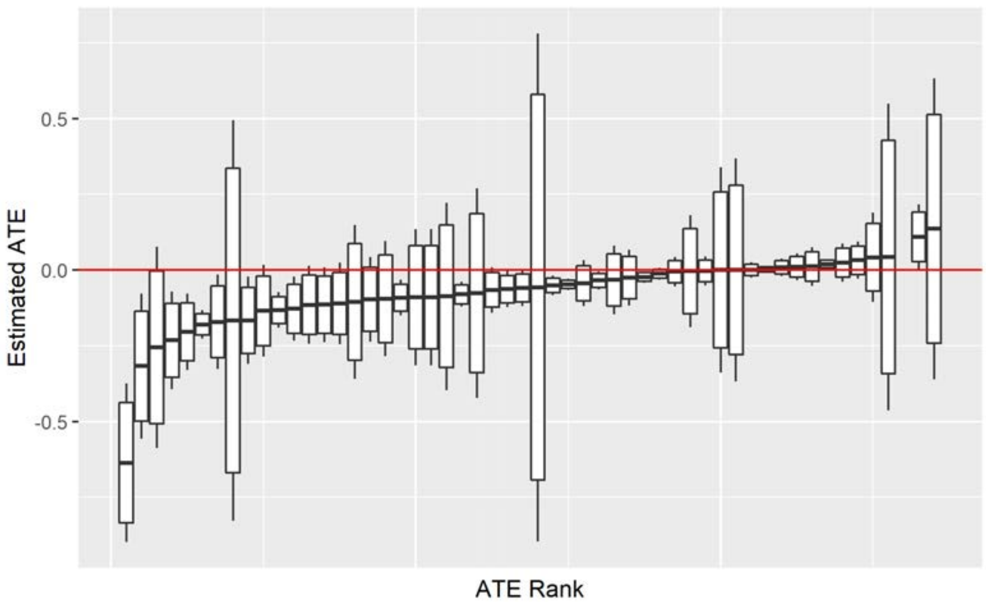

Notes: This figure plots average treatment effects of nudges on the cigarette cessation probability, together with 95%- and 99%-confidence intervals. Positive values indicate by many percent the nudge increased the probability to quit smoking. Point estimates are ranked from lowest to highest.

Figure 11: Cigarette Price Elasticities

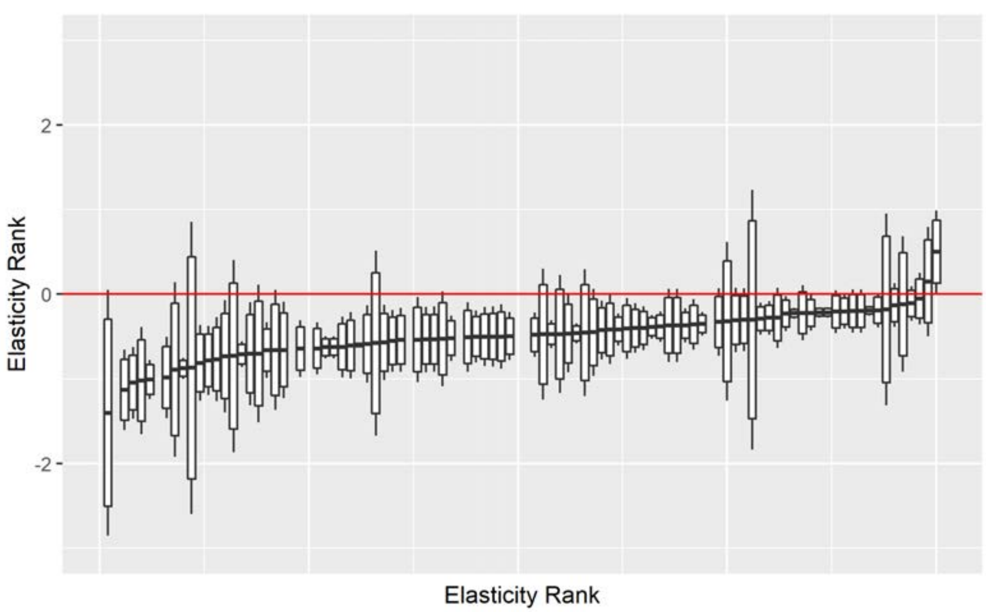

Notes: This figure plots cigarette price elasticities, together with 95%- and 99%-confidence intervals. Negative values indicate by how many percent cigarette demand decreases when the cigarette price increases by 1%. Point estimates are ranked from lowest to highest.

Figure 12: Nudge Treatment Effects on Influenza Vaccination Probability

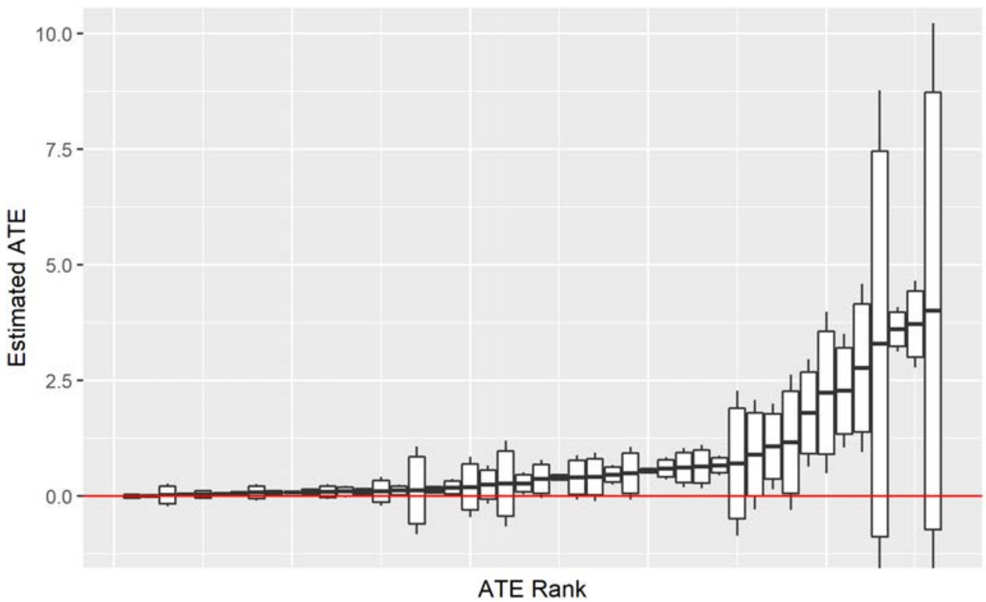

Notes: This figure plots average treatment effects of nudges on the influenza vaccination probability, together with 95%- and 99%-confidence intervals. Positive values indicate by how many percent the nudge increased the probability to get vaccinated. Point estimates are ranked from lowest to highest.

Figure 13: Influenza Vaccine Price Elasticities

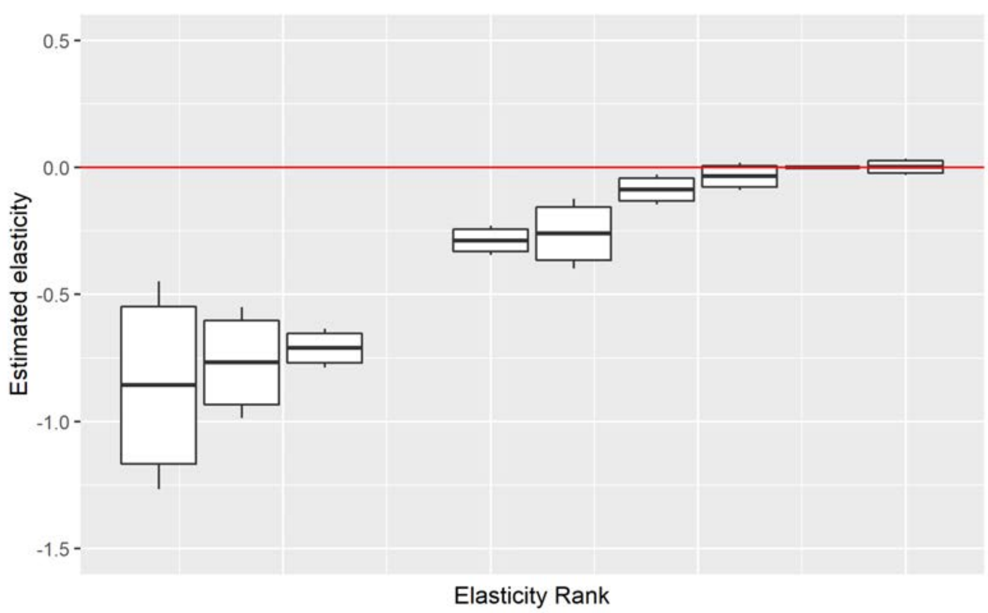

Notes: This figure plots price elasticities of influenza vaccines, together with 95%- and 99%-confidence intervals. Negative values indicate by how many percent demand for vaccines decreases when the vaccine price increases by 1%. Point estimates are ranked from lowest to highest.

Figure 14: Nudge Treatment Effects on Electricity Consumption

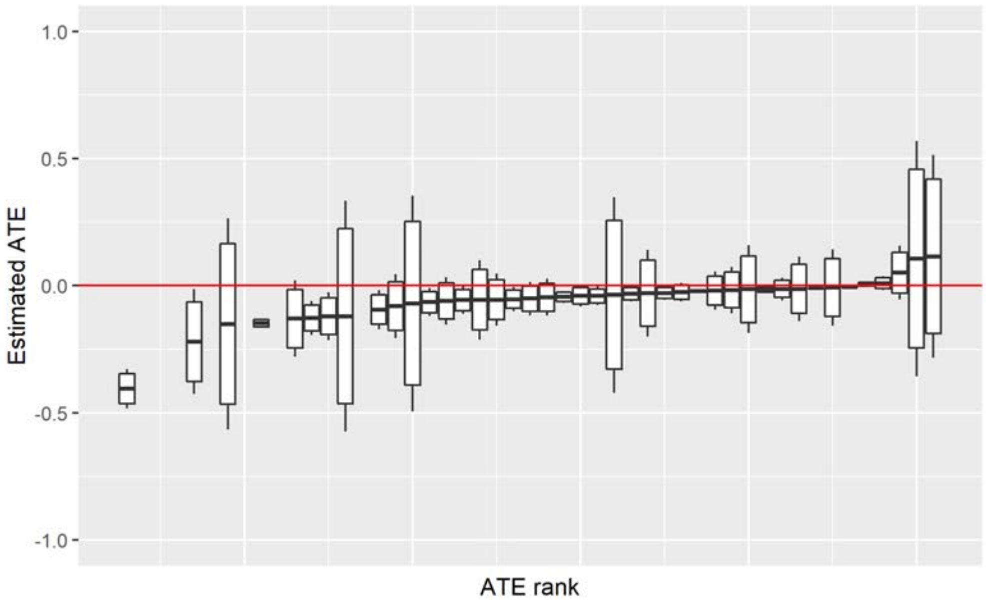

Notes: This figure plots average treatment effects of nudges on household electricity consumption, together with 95%- and 99%-confidence intervals. Negative values indicate by how many percent the nudge decreased the probability to quit smoking. Point estimates are ranked from lowest to highest.

Figure 15: Electricity Price Elasticities

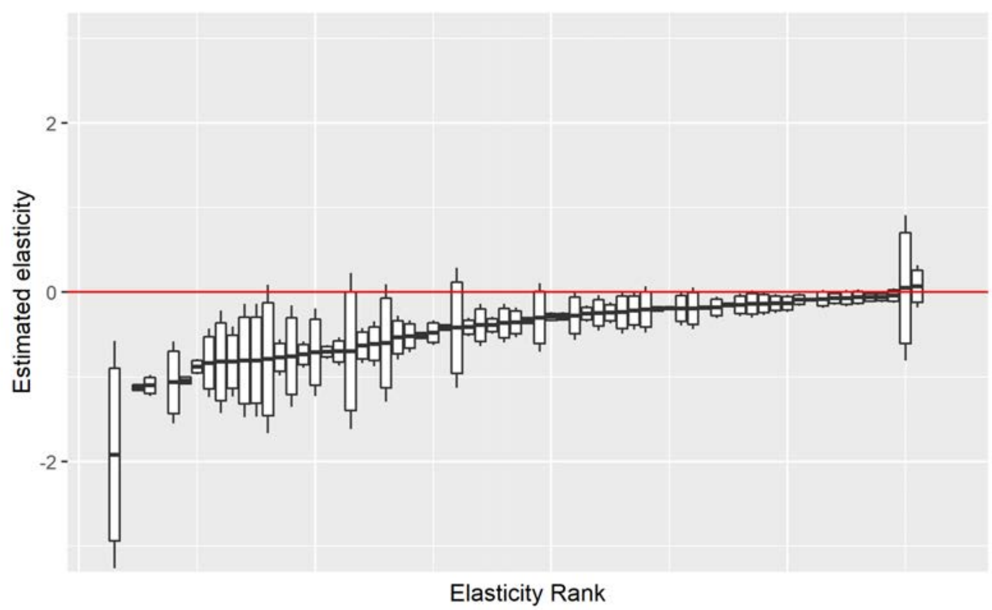

Notes: This figure plots electricity price elasticities, together with 95%- and 99%-confidence intervals. Negative values indicate by how many percent households' electricity demand decreases when the electricity price increases by 1%. Point estimates are ranked from lowest to highest.

D Correction for Publication Bias

In this section, we show that our key results are robust to accounting for publication bias. Sample size limitations preclude us from estimating a rich model of sample selection, so we instead opt to take a simple approach in a similar spirit to one of the approaches of DellaVigna and Linos (2022), which we henceforth refer to as DL. Specifically, we model publication bias as taking a particular form where publishing probability depends only on whether the main result in a paper is significant or not. Formally, for each study, s, we assume

$$\Pr[\text{Publish } s] = \begin{cases} p_0 & \text{study } s \text{ is not statistically significant} \\ p_1 & \text{study } s \text{ is statistically significant} \end{cases} \tag{37}$$

Given this simple model of publication bias, we can correct for publication bias simply by re-weighting our sample so that statistically insignificant studies receive p_1/p_0 times more weight than statistically significant studies.

While p_1 and p_0 are not separately identified given our data, their ratio is identified adopting a "regression discontinuity" style strategy as is taken in DL. Formally, we assume that un-selected z-scores across studies are distributed according to a smooth distribution $f(z)$. Then Equation (37) implies that the observed density of z scores, which we denote by g satisfies

$$g(z) \propto f(z) \left[\mathbb{1}\{|z| < z_{0.975}\}p_0 + \mathbb{1}\{|z| \geq z_{0.975}\}p_1\right]$$

where $z_{0.975} \approx 1.96$ is the threshold for statistical significance for a two-sided test at the 5% significance level. Then taking limits from the left and right as z approaches 1.96, we have

$$\frac{\lim_{z \to z_{0.975}^{+}} g(z)}{\lim_{z \to z_{0.975}^{-}} g(z)} = \frac{p_1}{p_0} \tag{38}$$

We approximate Equation (38) empirically by taking the ratio between N_1, the number of studies which are statistically significant and with z score within 0.25 of $z_{0.975}$ and N_0, the number of studies which are statistically insignificant and with z score within 0.25 of $z_{0.975}$.[18] Within our sample of studies, we find that $N_1 = 19$ and $N_0 = 10$, so $N_1/N_0 \approx 2$. This is somewhat smaller than the point estimate of approximately $5 = 10/2$ found by the comparable computation in DL, but we appear to have a larger mass of studies

[18] The results are not sensitive to changing the cutoff.

which are on the boundary of statistical significance, which makes our parameter somewhat more precisely estimated. Our implied publication bias parameter, $\gamma \equiv N_0/N_1$ is also well within the 95% confidence intervals implied by DL's structural estimates.

We thus proceed by taking $p_1/p_0 = 2$ and re-compute the main results of the paper by up-weighting statistically insignificant point estimates by a factor of 2. The resulting point estimates for average nudge treatment effects are now slightly lower. In the market for cigarettes, we now find a that nudges only decrease smoking probability by 5.7% whereas without correcting for publication bias, we find an effect of 7.5%. In the market for vaccines, we find an 31% increase in vaccinations, compared to the 35% effect found without correcting for publication bias. Finally, in the electricity market, we continue to find a 4.8% average treatment effect. Note, however that our welfare framework highlights the importance of *variance of bias*, not the *average treatment effect of nudge*, as the primary determinant of the efficiency of nudges. We thus turn to our welfare results next.

The main welfare effects are displayed in Figures 16 through 18. Meanwhile, Tables 4 and 5, are publication-bias corrected analogues of Tables 1 and 2. In the market for cigarettes and household electricity, we find somewhat smaller welfare benefits of nudges, reflecting lower estimated heterogeneity in behavioral bias after correcting for publication bias. In the market for influenza, we in fact find larger welfare effects of nudges. Nonetheless, the main qualitative insights of our paper remain intact: the markets for influenza vaccines and household electricity continue to yield point estimates that favor taxes while the markets for cigarettes continues to yield point estimates favoring nudges.

Table 4: Welfare Effects, $\theta = 0.8$, $\rho = 0.5$, Corrected for Publication Bias

	Cigarettes (per consumer per year)	Influenza Vaccines (per person per year)	Electricity (per household per year)
Optimal Tax in isolation	\$65 [\$29,\$127]	\$52 [\$43,\$94]	\$850 [\$628, \$1,098]
Nudge in isolation	\$67 [\$27,\$187]	\$28 [\$17,\$61]	\$90 [\$57,\$129]
Nudge and optimal tax in combination	\$77 [\$34,\$200]	\$61 [\$50,\$99]	\$839 [\$631,\$1,100]

Notes: This table reports welfare effects of different policies in the market for cigarettes, influenza vaccines and household electricity, correcting for publication bias. The first row shows welfare effects of implementing the optimal tax, while the second row reports welfare effects of using nudges. The final row gives the welfare effects of using both tools in combination. For the estimations, we use our baseline assumptions that the nudge is 80% effective in reducing the behavioral bias, $\theta = 0.8$, and that the correlation between nudge treatment effects and price elasticities is $\rho = 0.5$. See the Figures 3, 5, and 7 for a wide range of alternative assumptions.

Figure 16: Welfare Effects in The Cigarette Market, Corrected for Publication Bias

(a) Fixing $\theta = 0.8$ (b) Fixing Correlation $= 0.5$

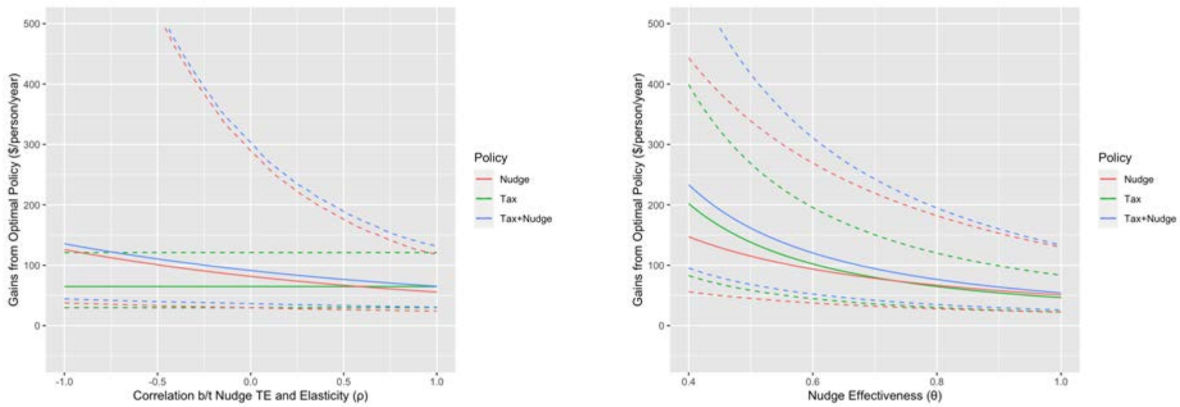

Notes: The figures illustrates welfare effects of nudges, optimal taxes, and a combination of the two policies in the market for cigarettes, correcting for publication bias. Panel a) reports welfare effects for different correlations between nudge treatment effect and price elasticity, while assuming that the nudge is 80% effective in reducing the behavioral bias. Panel b) reports welfare effects for different values of nudge effectiveness, while assuming a correlation between nudge treatment effect and price elasticity of 50%. Dashed lines indicate confidence intervals.

Figure 17: Welfare Effects in the Market for Vaccination Take-Up, Corrected for Publication Bias

(a) Fixing $\theta = 0.8$ (b) Fixing Correlation $= 0.5$

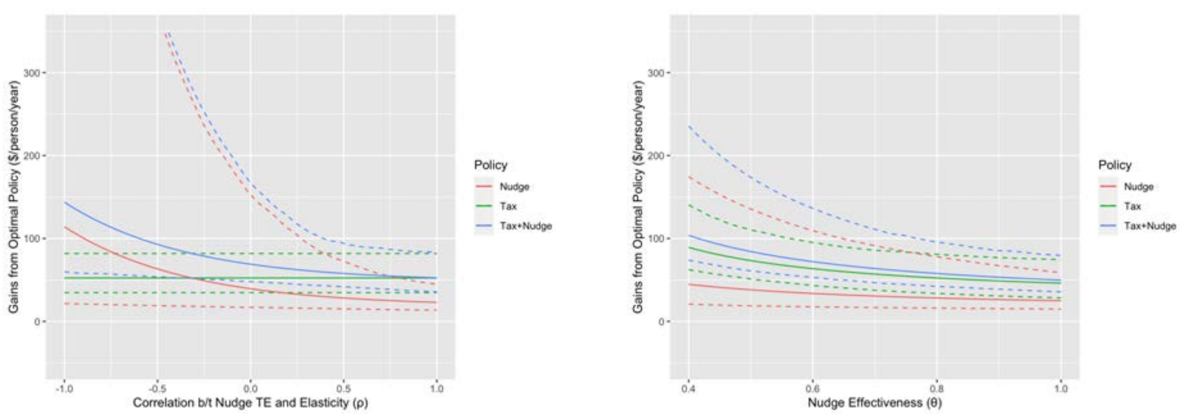

Notes: The figures illustrates welfare effects of nudges, optimal taxes, and a combination of the two policies in the market for influenza vaccines, correcting for publication bias. Panel a) reports welfare effects for different correlations between nudge treatment effect and price elasticity, while assuming that the nudge is 80% effective in reducing the behavioral bias. Panel b) reports welfare effects for different values of nudge effectiveness, while assuming a correlation between nudge treatment effect and price elasticity of 50%. Dashed lines indicate confidence intervals.

Figure 18: Welfare Effects in the Market for Household Electricity, Corrected for Publication Bias

(a) Fixing $\theta = 0.8$ (b) Fixing Correlation $= 0.5$

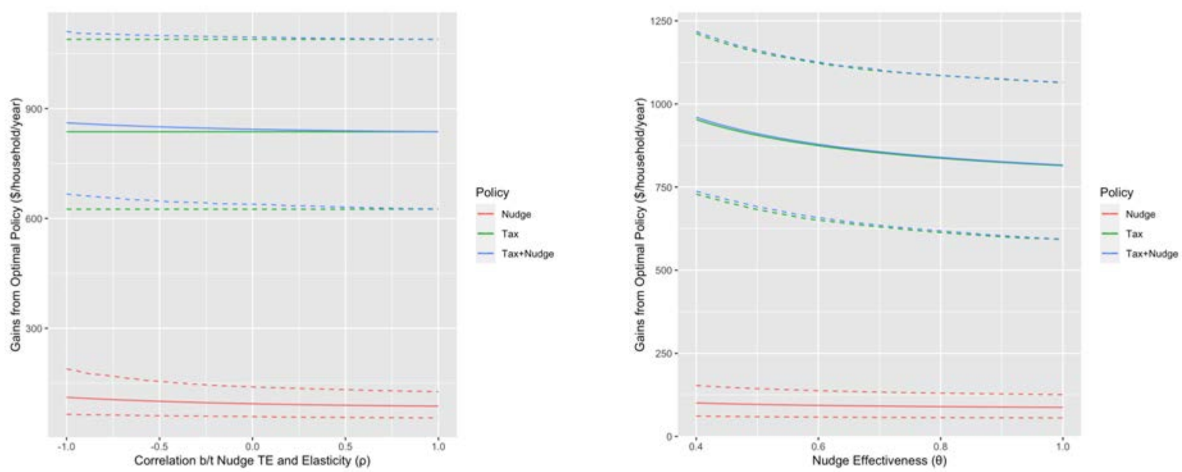

Notes: The figures illustrates welfare effects of nudges, optimal taxes, and a combination of the two policies in the market for household electricity, correcting for publication bias. Panel a) reports welfare effects for different correlations between nudge treatment effect and price elasticity, while assuming that the nudge is 80% effective in reducing the behavioral bias. Panel b) reports welfare effects for different values of nudge effectiveness, while assuming a correlation between nudge treatment effect and price elasticity of 50%. Dashed lines indicate confidence intervals.

Table 5: Optimal Taxes, Corrected for Publication Bias

	Cigarettes (per pack)	Influenza Vaccines (per vaccine)	Electricity (per kWh)
EPM of behavioral bias ($\theta = 1$)	\$1.78 (\$0.43)	-\$49 (\$78)	\$0.01 (\$0.003)
Optimal isolated tax ($\theta = 0.8$)	\$2.91 (\$0.54)	-\$215 (\$97)	\$0.20 (\$0.003)
Optimal tax with nudge ($\theta = 0.8$)	\$1.13 (\$0.11)	-\$165 (\$19)	\$0.19 (\$0.001)
Generalized targeting ratio of nudge ($\theta = 0.8, \rho = 0.5$)	1.42 (0.58)	0.72 (1.8)	0.086 (0.029)

Notes: This table reports the equivalent price metric in each market, as well as the size of the optimal tax with and without nudge, correcting for publication bias. The last row indicates the generalized targeting ratio, as defined in Equation (23) as $\frac{\theta(2-\theta)\mathrm{sd}_W(b_0)}{|(1-\theta)\mathbb{E}_W[b_0]+\mathbb{E}[\xi]|}$