Northwestern | INSTITUTE FOR POLICY RESEARCH

IPR Working Paper Series

WP-23-21

Optimal Paternalism in a Population With Bounded Rationality

Charles F. Manski
Northwestern University and IPR

Eytan Sheshinski
Hebrew University

Version: June 12, 2023

DRAFT

Abstract

A central mission of public economics has been to determine policies that optimize utilitarian welfare. To improve the realism of policy evaluation, it is desirable to enrich understanding of behavior and develop methods of analysis that use the enriched understanding to assess policies. Attention has been given to recognition that behavior may be boundedly rational, but little has been done to draw the implications of this work for mechanism design. Behavioral economists have suggested that planners should limit the choices available to individuals or should frame the options in a manner thought to influence choice, but the discussion has commonly been verbal and casual. Manski and Sheshinski formally consider a planner who has the power to design a discrete choice set from which individuals will choose. They suppose that individuals may be boundedly rational and, hence, may not maximize utility. Their concern is realistic settings in which persons have heterogeneous preferences and may vary in how their choices deviate from utility maximization. The researchers find that optimal paternalism is subtle. The policy that most effectively constrains or influences individual choices depends on the distribution of preferences and the choice probabilities measuring the extent to which persons behave suboptimally, conditional on their preferences.

This paper developed from research initiated in the early 2000s by Eytan Sheshinski, with the title "Socially Desirable Limits on Individual Choice." Presentations by Sheshinski of this original work have been made in various settings, including in a seminar at the University of Bonn in 2004 and as the Richard Musgrave Lecture in Munich in April 2012. At various stages, the authors have benefitted from the comments of Peter Diamond and Joram Mayshar.

"We're thinking maybe it's time you started getting some religious instruction. There's Catholic, Protestant, and Jewish—any of those sound good to you?"

1. Introduction

A central mission of research in public economics has been to determine policies that optimize utilitarian welfare, recognizing that policy choice affects individual behavior. To ease analysis, economists studying mechanism design have maintained simplifying assumptions about behavior. For example, the classic Mirrlees (1971) study of optimal income taxation assumed that individuals maximize static deterministic utility when choosing labor supply. It assumed that individuals have homogeneous consumption-leisure preferences and are heterogeneous only in ability, hence wage. Analysis of optimal taxation is complex even when these simplifying assumptions are maintained. Limited empirical understanding of actual individual preferences and labor supply has impeded application of the theory (Manski, 2014).

To improve the realism of policy evaluation, it is desirable to enrich empirical understanding of behavior and develop methods of theoretical analysis that use the enriched empirical understanding to assess alternative policies. There is huge scope for advancement relative to the present state of knowledge.

Theoretical study of utilitarian policy choice began in the 1700s, was formalized in the first half of the 1900s, and continued to develop steadily through the 1970s, but it then more or less reached a plateau. The subject has received disturbingly little attention recently. The field of public economics has become dominated by empirical analysis of randomized and quasi-randomized experiments, using statistical methods that do not study the structure of human behavior and do not enable prospective evaluation of proposed policies.

Empirical study of the structure of individual behavior has progressed in fields other than public economics, including industrial organization and behavioral economics. Attention has been given to understanding heterogeneity in preferences, measurement of expectations in environments of choice under uncertainty, and recognition that behavior may be boundedly rational in various respects. However, little has been done to draw the implications of this empirical work for mechanism design.

The concern of this paper is optimization of utilitarian welfare in populations with bounded rationality. Behavioral economists have suggested that social planners should limit the choice options available to individuals to ones deemed beneficial from a utilitarian perspective or, less drastically, should frame the options in a manner thought to influence choice in a positive way. Thaler and Sunstein (2003) evocatively wrote that such policies express "libertarian paternalism." However, here and in their subsequent influential work on "nudges" (Thaler and Sunstein, 2008), their discussion has been verbal and casual, rather than formal and careful. A new consensus report by a National Academies committee on *Policy Impact and Future Directions for Behavioral Economics* is similarly verbal and casual (National Academies of Sciences, Engineering, and Medicine, 2023). We find this frustrating.

A rare expression of the type of research that we think desirable has been given by O'Donoghue and Rabin (2003), who began their article as follows (p. 186):

> "The classical economic approach to policy analysis assumes that people always respond optimally to the costs and benefits of their available choices. A great deal of evidence suggests, however, that in some contexts people make errors that lead them not to behave in their own best interests. Economic policy prescriptions might change once we recognize that humans are humanly rational rather than superhumanly rational, and in particular it may be fruitful for economists to study the possible advantages of *paternalistic policies* that help people make better choices.
>
> We propose an approach for studying optimal paternalism that follows naturally from standard assumptions and methods of economic theory: Write down assumptions about the distribution of rational and irrational types

> of agents, about the available policy instruments, and about the government's information about agents, and then investigate which policies achieve the most efficient outcomes. In other words, economists ought to treat the analysis of optimal paternalism as a mechanism-design problem when some agents might be boundedly rational."

We present an analysis of the type sought by O'Donoghue and Rabin, addressing a different class of policy problems than they did.

We consider a social planner who has the power to design a discrete choice set from which individuals will choose. We suppose that there is no social cost to offering larger choice sets. Hence, classical utilitarian welfare economics recommends that the planner should offer the largest choice set possible. We depart from the classical setting by supposing that individuals may be boundedly rational and, hence, may not choose options that maximize utility. In such settings, it may be optimal for the planner to constrain the choice set to prevent persons from choosing inferior actions or, less drastically, to frame the choice set in a manner that influences behavior.

Example: Some public social security systems and private pensions have an early eligibility age at which a person can start receiving a pension, with less than full benefits. This age differs widely across countries. In the US, partial benefits are obtainable at age 62 and full benefits later (historically at age 65, in process of advancing to 67). The UK has had a single State Pension Age determining eligibility for full benefits (historically at age 65, in process of advancing to 67), with no option of earlier retirement with lower benefits. Imposing a constraint on the earliest age for eligibility hurts workers who would sensibly stop working before this age due to health and other personal circumstances. On the other hand, it prevents people from retiring too early, reflecting shortsightedness. Setting an early eligibility age should strike a balance between these considerations. ■

The planner's problem is straightforward if all members of the population have the same known preferences. Then the optimal paternalistic policy calls on the planner to determine the population-wide best option and mandate it. Our concern is more realistic settings in which persons have heterogeneous preferences and, additionally, may vary in how their choices deviate from utility maximization. We find that optimal paternalism is subtle in these settings. The policy that most effectively constrains or influences individual choice depends both on the distribution of preferences in the population and on the choice probabilities measuring the extent to which persons behave suboptimally, conditional on their preferences.

We show how the interaction of population preferences and choice probabilities determines the utilitarian welfare of policies that seek to ameliorate bounded rationality. The prevailing practice in empirical research in behavioral economics has been to perform experiments that present various options to subjects, framed in particular ways, and observe the choices that subjects make. Research of this type usually does not seek to measure the preferences of subjects and, hence, cannot study the interaction of choices and preferences, which is essential to evaluation of utilitarian welfare. Section 2 makes this central point in a concise, abstract manner. Our formal policy analysis covers both mandates that constrain choice and nudges that aim to influence choice.

Section 3 considers policy choice when individuals are boundedly rational in a specific way, this being that they measure utility with additive random error and maximize mismeasured rather than actual utility functions. Studying this type of bounded rationality yields more detailed analysis and enables presentation of an instructive numerical illustration. When the errors in utility mismeasurement are independent and identically distributed, we obtain a lower bound on the welfare achieved by a policy that constrains the choice set. When the errors have the Type I extreme value distribution, choice probabilities have the multinomial logit form. When the scale parameter of the error distribution is invariant across policies, this parameter succinctly characterizes the degree of rationality in the population. A numerical example shows that the optimal paternalistic policy varies in a subtle way with the degree of rationality.

The sensitivity of optimal policy to the fine structure of preferences and bounded rationality, about which little is known empirically, motivates us in the concluding Section 4 to caution against premature implementation of policies that attempt to ameliorate bounded rationality by constraining or influencing choice behavior.

2. Policy Choice in Abstraction

2.1. Setup and Findings

Let J denote the population of concern to a utilitarian planner. Let C denote a maximal finite choice set that the planner may make available to each member of J. Let each individual $j \in J$ have a utility function $u_j(\cdot)$: $C \rightarrow R$ expressing the person's preferences. Let $u_j^* \equiv \max_{c \in C} u_j I$.

To formalize utilitarian welfare, consider J to be a probability space with distribution P(j) and let $P[u(\cdot)]$ denote the population distribution of utility functions. Let utility functions be interpersonally comparable. Then the idealized optimum utilitarian welfare, if all persons maximize utility, is $E(u^*)$.

The problem is that, having bounded rationality, individuals may not maximize utility. Let the planner choose among a set S of policies that may constrain or influence choice behavior. Suppose that, with policy s, person j chooses $c_j(s) \in C$. For each $I \in C$, let $P[c(s) = i|u(\cdot)]$ denote the fraction of persons with utility function $u(\cdot)$ who would choose option I under policy s. The utilitarian welfare achieved by this policy is

$$(1) \qquad E\{u[c(s)]\} = \int \sum_{I \in C} u(i) \cdot P[c(s) = i|u(\cdot)]dP[u(\cdot)].$$

The optimal feasible welfare is achieved by a policy that solves the problem $\max_{s \in S} E\{u[c(s)]$.

Observe that the value of $E\{u[c(s)]\}$ depends on both the preference distribution $P[u(\cdot)]$ and the conditional choice probabilities $P[c(s)|u(\cdot)]$. It is revealing to consider the regret of a policy, its degree of sub-optimality, relative to the idealized optimum utilitarian welfare $E(u^*)$. The regret of policy s is

$$(2) \qquad E(u^*) - E\{u[c(s)]\} = \int \sum_{I \in C} [u^* - u(i)]P[c(s) = i|u(\cdot)dP[u(\cdot)].$$

For each utility function u(·) and alternative I, $[u^* - u(i)]P[c(s) = i|u(\cdot)]$ is the degree of sub-optimality of I multiplied by its choice probability. Thus, the regret of policy s is a weighted average of the multiplicative interactions of choice probabilities for alternatives and their degrees of sub-optimality.

2.2. Mandates and Nudges

In the Introduction, we distinguished between policies that constrain and influence individual choices. Both types are encompassed in the above general setup. A choice-constraining policy limits the effective choice set to some $C(s) \subset C$, implying that $P[c(s) = i|u(\cdot)] = 0$ for all $i \notin C(s)$ and all u(·). Such a policy is a mandate if C_s is singleton. Mandating alternative i yields welfare E[u(i)]. Hence, the optimal mandate is to an alternative i^m that solves the problem $\max_{i \in C} E[u(i)]$.

The optimal mandate yields the idealized optimal utilitarian welfare if all members of the population have the same preferences. It yields lower welfare when preferences are heterogeneous. This holds by Jensen's Inequality, which implies that $\max_{i \in C} E[u(i)] \leq E[\max_{i \in C} u(i)]$, the inequality being strict when preferences are heterogeneous. The optimal mandate is best for persons who most prefer alternative i^m, but not for those with other preferences.

A choice-influencing policy enhances the prominence of a specified alternative but does not prevent persons from choosing other alternatives. Behavioral economists have sought to enhance prominence by specifying an alternative to be the "default option," by placing it first in the ordering of alternatives, by associating it with favorable images, and in other ways. Whatever mechanism is used, the objective is to increase the probability with which persons choose this alternative.

When considering choice-influencing policies, a behavioral economist may recommend enhancing the prominence of i^m, the optimal mandate. Such policies have been called nudges. In general, the impact of a nudge on choice behavior may depend not only on the manner in which the policy frames the choice set C, but also on the joint distribution of preferences and forms of bounded rationality in the population. It is impossible to evaluate nudges in abstraction. One must consider the context.

3. Policy Choice with Additive Error in Utility Measurement

In this section we consider policy choice when individuals are boundedly rational in a specific way. We assume that they measure utility with additive error and maximize mismeasured rather than actual utility functions. We do not assert that this type of bounded rationality is prevalent in actual populations. We study it because the idea is easy to understand and because it enables us to apply findings on choice probabilities developed in the literature analyzing random utility models.

3.1. Choice Probabilities Generated by Random Utility Models

Let policy s constrain choice to a choice set C(s), which may be any non-empty subset of C. We assume that, under policy s, person j mismeasures the utility of each $c \in C(s)$ as $u_j(c) + \varepsilon_j(c, s)$, chooses an alternative $c_j^{\#}(s) \equiv \text{argmax}_{\ c \in C(s)}\ u_j(c) + \varepsilon_j(c, s)$, and thus achieves utility $u(c_j^{\#})$. For simplicity, we assume that the conditional error distribution $P[\varepsilon(c, s), c \in C(s)|u(\cdot)]$ is continuous. This implies that $c_j^{\#}(s)$ is unique for almost every person j. Hence, choice probabilities are well-defined, with

$$(3) \quad P[c^{\#}(s) = i|u(\cdot)] \ = \ P[u(i) + \varepsilon(i, s) \geq u(c) + \varepsilon(c, s), \text{all } c \in C(s)|u(\cdot)], \quad i \in C(s).$$

Inserting these choice probabilities into (1) yields the welfare achieved by policy s, which is

$$(4) \quad E\{u[c(s)]\} \ = \ \int \sum_{i \in C(s)} u(i)\cdot P[u(i) + \varepsilon(i, s) \geq u(c) + \varepsilon(c, s), \text{all } c \in C(s)|u(\cdot)]dP[u(\cdot)].$$

Equation (3) provides a standard random-utility model interpretation of bounded rationality (McFadden, 1974; Manski, 1977). The values of the choice probabilities (3) are determined by the conditional error distribution $P[\varepsilon(c, s), c \in C(s)|u(\cdot)]$. In the absence of restrictions on this distribution, any choice probabilities are possible. Hence, assuming that a random utility model expresses bounded rationality does not, per se, yield restrictions on E{u[c(s)]}. Some knowledge of the error distributions is necessary. The discussion below considers successively stronger levels of knowledge.

3.2. Simple Scalability

A lower bound on E{u[c(s)]} emerges if, conditional on each utility function u(·), the error components ε(c, s), c ∈ C(s) are known to be independent and identically distributed (i.i.d.). We do not assume knowledge of the specific distribution here. Indeed, it may possibly vary with u(·). The i.i.d. assumption expresses the idea that individuals make "white-noise" errors in utility measurement. The error distribution may vary across persons j and policies s. However, errors do not vary systematically across alternatives. The i.i.d. assumption is generally not appropriate when studying policies that generate nudges, which asymmetrically influence evaluation of different choice options.

Manski (1975) showed that, when errors are conditionally i.i.d., choice probabilities are related to utility functions by a set of inequalities called *simple scalability*. For each utility function u(·) and alternative pair (a, b) ∈ C(s) x C(s),

$$\text{(5a)}\quad u(a) > u(b) \Leftrightarrow P[u(a) + \varepsilon(a, s) \geq u(a) + \varepsilon(c, s), \text{ all } c \in C(s)|u(\cdot)]$$
$$> P[u(b) + \varepsilon(b, s) \geq u(c) + \varepsilon(c, s), \text{ all } c \in C(s)|u(\cdot)],$$

$$\text{(5b)}\quad u(a) = u(b) \Leftrightarrow P[u(a) + \varepsilon(a, s) \geq u(a) + \varepsilon(c, s), \text{ all } c \in C(s)|u(\cdot)]$$
$$= P[u(b) + \varepsilon(b, s) \geq u(c) + \varepsilon(c, s), \text{ all } c \in C(s)|u(\cdot)].$$

Let $u^{mean}(s)$ denote the mean value of utility within set C(s); that is, $u^{mean}(s) \equiv [1/|C(s)|] \sum_{c \in C(S)} u(c)$. It follows from (5a)-(5b) that

$$(6) \quad u^{mean}(s) \ \le \ \sum_{i \in C(s)} u(i) \cdot P[u(i) + \varepsilon(i, s) \ge u(c) + \varepsilon(c, s), \text{ all } c \in C(s)|u(\cdot)].$$

Hence,

$$(7) \quad \int u^{mean}(s) dP[u(\cdot)] \ \le \ E\{u[c(s)]\}.$$

This bound varies across policies that constrain choice to different subsets of C. However, among policies that constrain choice to the same subset of C, the bound does not vary across policies that seek to influence utility measurement in different ways.

3.3. Multinomial Logit Choice Probabilities

A substantial strengthening of the knowledge considered thus far supposes that errors have a type 1 extreme-value distribution, also called the Gumbel distribution. Assume that $\varepsilon(c, s)$, $c \in C(s)$ are independent and have the common distribution function $P[\varepsilon(c, s) \le t] = \exp[-e^{-q(s)t}]$. Here q(s) is a positive scaling factor whose value may vary with s. Then the conditional choice probabilities have the multinomial logit form (McFadden, 1974):

$$(8) \quad P[c^{\#}(s) = i|u(\cdot)] \ = \ P[u(i) + \varepsilon(i, s) \ge u(c) + \varepsilon(c, s), \text{ all } c \in C(s)|u(\cdot)] \ = \ e^{q(s)u(i)}/\sum_{c \in C(s)} e^{q(s)u(c)}, \ \ i \in C(s).$$

Hence, the welfare achieved by policy s is

$$(9) \quad E\{u[c(s)]\} \ = \ \int \sum_{i \in C(s)} u(i) \cdot \left[e^{q(s)u(i)}/\sum_{c \in C(s)} e^{q(s)u(c)}\right] dP[u(\cdot)].$$

In this setting, q(s) concisely measures the success of policy s in inducing individuals to measure utility correctly within the constrained choice set C(s). Thus, q(s) quantifies the *degree of rationality* in the population, under policy s. As $q(s) \to \infty$, the spread of the error distribution decreases and the choice probability for the alternative that maximizes utility increases to one. As $q(s) \to 0$, the spread of the error distribution increases and the choice probabilities for all alternatives converge to $1/|C(s)|$. For each alternative $i \in C(s)$, the partial derivative of the choice probability with respect to q(s) is

$$(10) \quad \partial P[c^{\#}(s) = i|u(\cdot)]/\partial q(s) \; = \; P[c^{\#}(s) = i|u(\cdot)]\{u(i) - v[s, u(\cdot)]\},$$

where $v[s, u(\cdot)] \equiv \sum_{i \in C(s)} u(i)\cdot P[c^{\#}(s) = i|u(\cdot)]$ is expected utility for persons with utility function u(·). Thus, increasing q(s) raises the choice probabilities of alternatives whose utility is higher than expected utility and vice-versa.

3.4. A Numerical Example Showing the Subtlety of Optimal Choice-Constraining Policy

In Section 3.3, policy s was described by two factors, the set C(s) constraining individual choice and the degree of rationality q(s) achieved by the policy. We now specialize further, considering a set S of policies that yield the same value of q(s), now labeled q, and that differ only in their choice-constraining sets C(s). A policy cannot exclude every option in C, so there exist $2^{|C|} - 1$ such policies.

The welfare yielded by policy s is

$$(10) \quad E\{u[c(s)]\} \; = \; \int \sum_{i \in C(s)} u(i)\cdot \left[e^{q\cdot u(i)}/\sum_{c \in C(s)} e^{q\cdot u(c)}\right] dP[u(\cdot)].$$

Thus, the optimal choice-constraining policy is a function of q and P[u(·)]. Analytical determination of a policy that maximizes welfare does not seem feasible, but numerical calculation of welfare is possible given a specification of q and P[u(·)].

A numerical example demonstrates that optimal policy choice is subtle. The example is based on the famous Hotelling (1929) model of choice when individuals and stores are located on a line. Let C contain three alternatives (potential stores), each identified by a location x_i, i = 1, 2, 3 on a line. Let J contain three individuals, each residing at a location θ_j, j = 1, 2, 3 on this line. Let the utility of alternative *i* to person j be $u_i(\theta_j) = -(x_i - \theta_j)^2$. Thus, due to transportation costs, utility decreases with the distance of individual j's location, θ_j, from store x_i. By construction, preferences are *single-peaked*.

In our example, we specify $x_1 = 0.5$, $x_2 = 1$, $x_3 = 1.6$ and $\theta_1 = -0.5$, $\theta_2 = 1$, $\theta_3 = 2$. This yields the utility values shown in Figure 1:

Figure 1

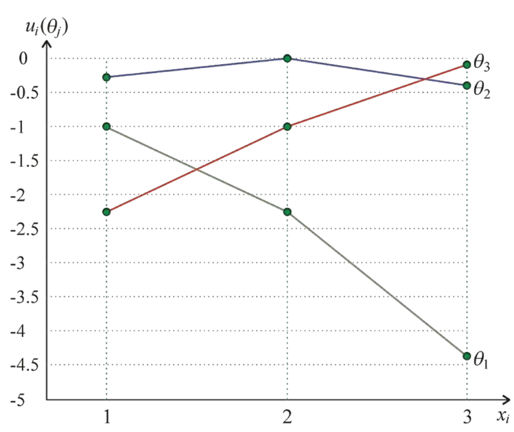

The seven possible constrained choice sets are {1}, {2}, {3}, {1, 2}, {2, 3}, {1, 3}, and {1, 2, 3}. The corresponding social welfare functions are denoted W^1, W^2, W^3, $W^{1,2}$, $W^{2,3}$, $W^{1,3}$, and $W^{1,2,3}$, respectively. Figure 2 plots each of these welfare functions against different levels of *q*. For each *q*, the optimum choice-set corresponds to *the outer envelope* of these plots.

Figure 2

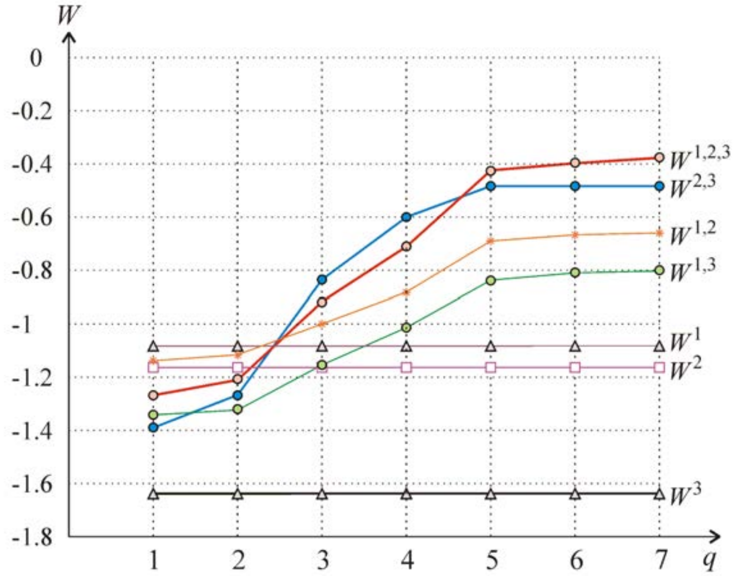

Observe that the optimal constrained choice set has a single alternative (W^1) at low values of q and includes all alternatives ($W^{1,2,3}$) at high values of q. Of particular interest is the fact that the ordering of welfare across choice sets is not nested, with reswitching occurring as q rises. For example, choice set {1, 2} outperforms set {2, 3} when q is smaller than about 2.5, but the welfare ordering reverses when q is larger. Choice set {1, 2, 3} outperforms set {2, 3} when q is smaller than about 2.2, the welfare ordering reverses for q between 2.2 and about 4.8, and then reverses again for q above 4.8.

This is only a example, but it suffices to demonstrate the subtlety of optimal policy choice. We find that, even in a highly simplified environment assuming multinomial logit choice probabilities, the welfare ordering of different choice-constraining policies is rather sensitive to the degree of rationality in the population. Empirical measurement of the degree of rationality would be a challenging task.

4. Conclusion

The optimal paternalistic policy for a population with homogeneous preferences is a mandate. However, we showed in Section 2 that the optimal policy in a heterogeneous population may be complex, involving either or both a constraint on choice and a nudge to influence choice. The detailed analysis of Section 3 showed that optimal paternalistic policy depends on the fine structure of the population distribution of preferences and bounded rationality. Thus, optimal policy may be highly context specific.

Unfortunately, detailed empirical knowledge of population distributions of preferences and bounded rationality is rare. Hence, we express caution against premature implementation of policies that attempt to ameliorate bounded rationality by constraining or influencing choice behavior. In the absence of firm empirical understanding of population behavior, such policies may do more harm than good.

References

Hotelling, H. (1929), “Stability in Competition,” *The Economic Journal*, 39, 41-57.

Manski, C. (1975), “Maximum Score Estimation of the Stochastic Utility Model of Choice,” *Journal of Econometrics*, 3, 205-228.

Manski, C. (1977), “The Structure of Random Utility Models,” *Theory and Decision*, 8, 229-254.

Manski, C. (2014), “Identification of Income-Leisure Preferences and Evaluation of Income Tax Policy,” *Quantitative Economics*, 5, 145-174.

McFadden, D. (1974), “Conditional Logit Analysis of Qualitative Choice Behavior,” In P. Zarembka, ed., *Frontiers in Econometrics*, New York: Academic Press.

Mirrlees, J. (1971), “An Exploration in the Theory of Optimum Income Taxation,” *Review of Economic Studies*, 38, 175-208.

National Academies of Sciences, Engineering, and Medicine (2023), *Policy Impact and Future Directions for Behavioral Economics*, Washington, DC: The National Academies Press.

O’Donoghue, T. and M. Rabin (2003), “Studying Optimal Paternalism, Illustrated by a Model of Sin Taxes,” *American Economic Review* (*Papers and Proceedings*), 93, 186-191.

Thaler, R. and C. Sunstein (2003), “Libertarian Paternalism,” *American Economic Review* (*Papers and Proceedings*), 93, 175-79.

Thaler, R. and C. Sunstein (2008), *Nudge: Improving Decisions About Health, Wealth, and Happiness*, New Haven: Yale University Press.