The axiomatic foundation of logit and its relation to behavioral welfare

Yves Breitmoser

Humboldt University Berlin

28 May 2016

The axiomatic foundation of logit and its relation to behavioral welfare

Yves Breitmoser*
Humboldt University Berlin

May 28, 2016

Abstract

Multinomial logit is the canonical model of discrete choice and widely used to analyze preferences and welfare. Its axiomatic foundation is incomplete: binomial logit is *assumed*; independence of irrelevant alternatives extends logit to multinomial choice. The present paper provides a self-contained foundation, showing that the axiom "binomial choice is logit" is behaviorally founded by "narrow bracketing" and "no systematic mistakes" (e.g. default effects). This in turn allows me to drop the no-mistakes axiom, yielding generalized logit accounting for systematic mistakes axiomatically consistently. The results position logit in the "mistakes-debate" in behavioral welfare and clarify the foundation for the functional form.

JEL–Code: D03
Keywords: logit, axiomatic foundation, discrete choice, utility estimation, welfare

*I thank Friedel Bolle, Nick Netzer, Martin Pollrich, Sebastian Schweighofer-Kodritsch, Felix Weinhardt, Georg Weizsäcker and audiences at the BERA workshop in Berlin and at THEEM 2016 in Kreuzlingen for many helpful comments. Financial support of the DFG (project BR 4648/1–1) is greatly appreciated. Address: Spandauer Str. 1, 10099 Berlin, Germany, email: yves.breitmoser@hu-berlin.de, Telephone/Fax: +49 30 2093 99408/5619.

1 Introduction

Many economic studies analyze preferences or welfare, be it empirically, experimentally, or theoretically. Under the classical assumption that agents maximize utility, preferences and welfare can be inferred from choice by revealed preference. Behavioral analyses, however, have found choice to exhibit systematic biases (e.g. default effects) and to be stochastic—subjects' choices are incoherent and do not reliably reveal preference. This has inspired an enormous literature of models on stochastic choice to control for noise in preference estimation (Luce, 1959) and welfare analyses (Small and Rosen, 1981; Hanemann, 1984).

The central model in this literature is multinomial logit proposed by McFadden (1974).[1] McFadden (2001) attributes its appeal to the availability of "fully consistent descriptions" in terms of both an axiomatic foundation and a random utility representation. While there are many models with random utility representations, the axiomatic foundation of logit makes it special—logit appears to require little more than Luce's independence of irrelevant alternatives (IIA). To provide just one example of this perception, Fudenberg and Strzalecki (2015) assume "IIA (so that static choice is logit)" (p. 654). That is, while IIA does not suffice to derive multinomial logit, the difference between IIA and multinomial logit is perceived to be minor.

This perception is wrong. One of McFadden's axioms implicitly assumes that binomial choice is logit. IIA merely provides the extension of logit to multinomial choice. The present paper shows that the axiom "binomial choice is logit" is behaviorally founded in axioms on narrow bracketing and absence of systematic mistakes (i.e. choice biases such as default effects). This relates logit to recent discussions in behavioral economics and behavioral welfare, and perhaps most importantly, the self-contained foundation opens the door to study generalized models without assuming absence of systematic mistakes. I show that dropping this axiom yields a surprisingly tractable extension of multinomial logit, which in turn shows how to account for choice biases such as default effects in an axiomatically consistent manner.

This is discussed in detail in the concluding section. As point of departure, let us review the received axiomatic foundation (McFadden, 1974). Consider a decision maker

[1] The long list of studies studying preferences using multinomial logit includes analyses of risk preferences (Holt and Laury, 2002; Goeree et al., 2003), social preferences (Cappelen et al., 2007; Bellemare et al., 2008), and preferences and demand functions of consumers (McFadden, 1980; Berry et al., 1995).

(DM) with utility function u. The probability that DM chooses option x from a finite option set B is denoted as $\Pr(x|u,B)$. Assume all choice probabilities are positive and satisfy independence of irrelevant alternatives (IIA); formal statements follow below. Positivity and IIA imply that DM's choice probabilities can be represented as

$$\Pr(x|u,B) = \frac{\exp\{V(x,y|u)\}}{\sum_{x'\in B}\exp\{V(x',y|u)\}} \tag{1}$$

for some function V, given an arbitrary benchmark option $y \in B$. The relation of V to DM's true utility function, which is u by assumption, is unrestricted. A number of existing studies argue that V can be considered to represent DM's utility function, but this relation is formally unfounded. McFadden (1974) obtains Eq. (1) by defining $V(x,y|u)$ to be the log-odds of the choice between x and y,

$$V(x,y|u) = \log\left(\frac{\Pr(x|u,\{x,y\})}{\Pr(y|u,\{x,y\})}\right). \tag{2}$$

In order to pin down the relationship between V and u, McFadden uses Axiom 3 (page 110) which assumes additive separability of V as follows:[2]

$$V(x,y|u) = u(x) - u(y). \tag{3}$$

McFadden refers to Axiom 3 as "Irrelevance of Alternative Set Effect", but as Eq. (1) holds independently of which $y \in B$ is chosen as benchmark option, this irrelevance obtains in any case. To understand Axiom 3, recall that V are log-odds. Using V's definition, Axiom 3 turns out to be equivalent to assuming

$$\frac{\Pr(x|u,\{x,y\})}{\Pr(y|u,\{x,y\})} = \exp\{u(x)-u(y)\} \quad \Leftrightarrow \quad \Pr(x|u,\{x,y\}) = \frac{\exp\{u(x)\}}{\exp\{u(x)\}+\exp\{u(y)\}}.$$

The latter equation is the definition of binomial logit. That is, the assumed "Irrelevance of Alternative Set Effect" is equivalent to assuming that binomial choice is logit. Given this structure of binomial choice, IIA ensures that multinomial choice is multinomial logit. In turn, the core assumption, that binomial choice is logit, is not axiomatically founded and the assumptions actually underlying logit analyses are unclear.

By identifying and correcting this gap in logit's axiomatic foundation, the present

[2]Note that McFadden's notation differs slightly, as the primitives of his analysis are individual attributes rather than individual utilities.

paper shows that logit actually is based on the assumptions that choice satisfies a form of narrow bracketing and does not reflect systematic mistakes.[3] This in turn allows us to drop the no-mistakes assumption, to obtain a generalized logit model showing how to account for systematic mistakes, and connects logit to two lively discussions in behavioral economics and behavioral welfare economics.

To clarify the relation to the literature, let me first address the no-mistakes assumption, which is at the center of debate in behavioral welfare economics (Gul and Pesendorfer, 2007; Bernheim and Rangel, 2007; Kőszegi and Rabin, 2007). In the words of Kőszegi and Rabin (2008b, p. 193), choice exhibits systematic biases if DMs "take actions that they would not take could they fully assess the distribution of all relevant consequences of those actions". Prevalence of systematic biases requires the distinction of decision utility, which rationalizes choice, and true utility, which represents the decision maker's actual well-being. Welfare then equates with well-being and is consequently not revealed by choice (Kőszegi and Rabin, 2008a,b). Since the distinction of decision utility and true utility is not identified based on choice data, equally feasible approaches are to postulate that the alleged mistakes are founded in the true utility function (Gul and Pesendorfer, 2001, 2008) or to weaken the notion of revealed preference (see e.g. Bernheim and Rangel, 2009; Bernheim, 2009).

As shown below, logit's position in this debate is clear, equality of true utility and decision utility is assumed and systematic biases are thus assumed away. That is, using logit assumes that choice biases due to defaults, positioning, round numbers, or left-most digits (all of which will be discussed below) are rooted in the true utility function. As I also show, analysts need not take this position. Dropping the no-mistakes axiom yields a generalized form of logit that, by nature of the dropped axiom, allows to account for systematic mistakes without rooting them in the true utility function. Perhaps surprisingly, these approaches are almost equivalent: If we assume choice biases are not rooted in the true utility function, then their behavioral foundation shows up as a summand next to the true utility function. Specifically, in generalized logit, DM with true utility u chooses $x \in B$ with probability

$$\Pr(x|u,B) = \frac{\exp\{\lambda \cdot u(x) + w(x)\}}{\sum_{x' \in B} \exp\{\lambda \cdot u(x') + w(x')\}}. \tag{4}$$

[3] Here, narrow bracketing requires choice probabilities to be invariant to level shifts of utility, which holds true if each choice task is considered in isolation. For related behavioral analyses, see e.g. Read et al. (1999) and Rabin and Weizsäcker (2009). Further, in this introduction, I skip the technical axioms "positivity" and "richness" which are standard and ensure uniqueness of the logit representation.

If one assumes that choice biases represent deviations from utility maximization, their behavioral foundation shows up as additively separable components next to utilities. The simplicity of this (axiomatically derived) approach toward systematic mistakes is of course convenient, and with hindsight little surprising, as it once more shows that the distinction of true utility and decision utility is not identified based on choice data alone: the two approaches to either include $w(x)$ in the true utilities or treat it as an additively separable component are behaviorally equivalent. This finding directly relates to Gul and Pesendorfer (2001), who essentially argue to include properties of the option set in the utility function. The behavioral equivalence of the two approaches does not imply their normative equivalence, of course. As for normative analyses, we need to know if $w(x)$ is part of the true utility or not, which consequently requires additional sources of information (for discussion, see Benkert and Netzer, 2015).

The generalized logit model obtained here resembles the generalized logit model that Matejka and McKay (2015) obtained by analyzing rational inattention in discrete choice. The formal definitions even appear to be identical. This is in some sense reassuring, although the models actually are fundamentally different: the generalization of logit obtained here satisfies IIA while Matejka and McKay's does not. Further, the secondary components $w(x)$ obtained here are shown not to relate to payoffs or utilities, and by their compatibility with IIA they cannot relate to payoff relations across options in any way. Hence, they represent choice biases not related to utilities, many of which have been identified as the aforementioned presentation effects due to e.g. defaults or positioning. In contrast, the foundation of Matejka and McKay's model in rational inattention implies that it captures similarity effects, i.e. payoff similarity across options, and thus violates IIA (as they discuss in detail).[4] That said, the similarity of how logit is to be generalized in order to capture choice biases due to either presentation effects or similarity effects is striking. In both cases, additional summands next to the utility functions appear, which reaffirms this "additive" representation of choice biases and additionally complies with the general result that inequality of true utility and decision utility is not identified based on choice data alone.

The second major discussion addressed here concerns the choice between "fully

[4]Matejka and McKay (2015) discuss this illuminatingly. Briefly, a DM not knowing the utilities of his options needs to sample utilities which is costly. Assuming he does so rationally, i.e. choosing optimally which options to sample and when to stop given the costs of sampling, yields the generalized model. The choice of options to sample and when to stop depends on the similarity of the options' utilities, which in turn implies that the generalized logit model relaxes IIA and captures similarity effects.

specified" structural models such as logit and "robust" or "non-parametric" methods such as least squares. This discussion has been led in many fields, e.g. in empirical work (Keane, 2010; Rust, 2010), in analyses of auctions (Bajari and Hortacsu, 2005), and in analyses of risk and social preferences. The general argument in favor of least-squares methods appeals to their alleged robustness, which presumably results from either weaker or less demanding assumptions, in contrast to the seemingly specific functional form of logit. As shown here, this specific functional form follows from relatively innocent axioms: IIA, narrow bracketing, and "decision utility = true utility".[5] IIA does not affect the functional form at all and narrow bracketing is equally assumed in least squares analyses. Further, logit's "decision utility = true utility" can be dropped without obstructing its functional form, while least squares essentially replaces it by assuming (i) utility itself is not choice relevant and (ii) choice probabilities depend on the options' distances to the utility maximizer. Neither utility irrelevance (i) nor the presentation effect (ii) has been observed in the vast literature on choice biases. Why should it be plausible to make these particular assumptions about discrete choice instead of logit's "decision utility = true utility"? Even aside from plausibility, logit's axioms are not formally stronger than those underlying least squares, which renders the robustness of least squares in analyses of discrete choice rather questionable.

Section 2 derives the main result and Section 3 further discusses the result and routes for future work.

2 The axiomatic foundation of logit and generalized logit

A decision maker (DM) has the task to choose an option $x \in B$ from a finite set $B \subseteq X$. DM's true utility is captured by $u : X \to \mathbb{R}$, and given u, the probability that DM chooses $x \in B$ is $\Pr(x|u,B)$. The axiomatic foundation of logit proposed here builds on the two received axioms.

Axiom 1 (Positivity)**.** $\Pr(x|u,B) > 0$ for all $x \in B \subseteq X$

Axiom 2 (Independence of Irrelevant Alternatives, IIA)**.** For all $B, B' \subseteq X$,

$$\frac{\Pr(x|u,B)}{\Pr(y|u,B)} = \frac{\Pr(x|u,B')}{\Pr(y|u,B')} \qquad \text{for all } x,y \in B \cap B'.$$

[5] Again, to be clear, positivity and richness axioms are additionally used to obtain uniqueness.

Luce et al. (2008) provide an axiomatic foundation of Luce's choice axiom, which in turn implies IIA. Note that IIA is used here only for the purpose of axiomatizing logit. Analysts may of course substitute their favorite alternative for it to study models relaxing IIA.

Positivity allows to define IIA in the usual manner and in conjunction with IIA, it implies that a value function $v : X \to \mathbb{R}$ exists such that $\Pr(x|u,B) = v(x|u)/\sum_{x' \in V} v(x|u)$. The relation of utility u and value v is essentially unrestricted, as noted in the Introduction. One may define a function $V : X \to \mathbb{R}$ such that $V(x) = \log\{v(x)\}$ and refer to $V(x)$ as the stimulus associated with x. Formally, though, V represents log-odds and is therefore not related to DM's true utility function in an explicit manner.

The third standard axiom is richness and ensures that the logit representations of choice probabilities are uniquely adequate given the remaining axioms. Without richness, alternative representations may be viable depending on coarseness of the choice environment. The richness assumption invoked here is satisfied if e.g. utility is continuous in x and the set of potential options X, from which the experimenter designs the experiment by selecting B, is convex.

Axiom 3 (Richness)**.** There exist $x, x' \in X$ such that $u(x) \neq u(x')$. For all $x, x' \in X$ and all $\lambda \in [0,1]$, there exists $x'' \in X$ such that $u(x'') = \lambda\, u(x) + (1-\lambda)\, u(x')$.

To introduce the novel axioms, recall that logit choice probabilities are invariant to adding an arbitrary constant to all utilities (Block and Marschak, 1960). It is natural to include this distinctive property of logit in any axiomatic foundation. To provide intuition, level invariance is satisfied if DM considers the analyzed choice tasks in isolation, i.e. independently of his level of utility outside the experiment. Following Read et al. (1999), I refer to it as narrow bracketing.

Axiom 4 (Narrow bracketing)**.** For all $r \in \mathbb{R}$ and $x \in B \subseteq X$, $\Pr(x|u,B) = \Pr(x|u+r,B)$

Narrow bracketing implies that binomial choice probabilities can be expressed as functions of utility differences (Block and Marschak, 1960), which is satisfied for the large class of strong utility models (including for example probit).[6] The novel observation made here is that in conjunction with richness (and IIA), narrow bracketing yields a

[6]Alternatively, invoking scale invariance as opposed to level invariance would lead us to so-called strict utility models, see e.g. (Luce and Suppes, 1965), the best known instance of which is $\Pr(x|u,B) = (u(x))^t/\sum_y (u(y))^t$.

generalized logit representation of choice probabilities. Specifically, given the earlier axioms, narrow bracketing implies the value function $v(x) = \exp\{\lambda \cdot u(x) + w(x)\}$ where $w(x)$ does not depend on the utilities as such. This representation of choice probabilities coincides with the "generalized logit" form derived from rational inattention in discrete choice (Matejka and McKay, 2015), as indicated in the Introduction.

Definition 1 (Generalized logit)**.** Consider a DM with utility u. The choice probabilities are generalized logit if there exist $\lambda \in \mathbb{R}$ and $w : X \to \mathbb{R}$ such that for all finite $B \subseteq X$ and all $x \in B$,

$$\Pr(x|u,B) = v(x) / \sum_{x' \in B} v(x') \qquad \text{with } v(x) = \exp\{\lambda \cdot u(x) + w(x)\}. \tag{5}$$

Generalized logit allows the non-utility components $w(x)$ to be option specific, i.e. behaviorally relevant. Matejka and McKay consider a DM not knowing his utilities and sampling utilities to learn about the optimal option. In this context, the non-utility component of the choice propensities relate prior beliefs about optimal actions and the informational strategy of DM, which essentially depends on the perceived similarity of options. This implies that Matejka and McKay's model relaxes IIA. In the present context, the non-utility components of choice propensities, $w(x)$, relate to the presentation of the choice task as I discussed below.

If $w(x) = \text{const}$, these components cancel out and choice probabilities are logit.

Definition 2 (Logit)**.** Consider a DM with utility u. The choice probabilities are logit if there exists $\lambda \in \mathbb{R}$ such that for all finite $B \subseteq X$ and all $x \in B$,

$$\Pr(x|u,B) = v(x) / \sum_{x' \in B} v(x') \qquad \text{with } v(x) = \exp\{\lambda \cdot u(x)\}. \tag{6}$$

Clearly, any axiom bridging generalized logit and logit needs to rule out presentation effects by setting $w(x) = \text{const}$. Any such axiom is equivalent to assuming that choice depends solely on true utility, i.e. that decision utility equals true utility.[7] This effect can be achieved in various ways. For the purpose of the present analysis, I invoke permutation invariance—permuting choice probabilities is equivalent to permuting utilities—but for clarity, I refer to the axiom as "decision utility = true utility".

[7]Recall that the true utility is $u(x)$ by assumption. Correspondingly, and taking generalized logit with value function $v(x) = \exp\{\lambda \cdot u(x) + w(x)\}$ as given, let me define the decision utility to be $u(x) + w(x)/\lambda$.

Axiom 5 (Decision utility = true utility)**.** Given any bijective function $f: X \to X$ and any $B \subseteq X$, $\Pr(f(x)\,|\,u, B) = \Pr\left(x\,|\,u \circ f, B\right)$ for all $x \in B$.

Now we are equipped to establish the result.

Theorem 1. *Consider a decision maker with utility u, in a choice environment satisfying richness with choice probabilities* $\Pr(\cdot)$ *satisfying positivity.*

1. $\Pr(\cdot)$ *has a logit representation* $\Leftrightarrow$ $\Pr(\cdot)$ *satisfies Axioms 2, 4, and 5.*

2. $\Pr(\cdot)$ *has a generalized logit representation* $\Leftrightarrow$ $\Pr(\cdot)$ *satisfies Axioms 2 and 4.*

Proof. I proof that the axioms imply logit and generalized logit representation, respectively. The other direction, that the two representations satisfy the respective axioms, is immediate. To simplify notation, define $u_x = u(x)$ for all $x \in X$. I say that u_x is interior if there exist x', x'' such that $u(x') < u_x < u(x'')$.

Step 1: By positivity and IIA, for any $B \subseteq X$,

$$\Pr(x|u,B) = \frac{V\left(u_x, u_y, \{x,y\}\right)}{\sum_{x' \in B} V\left(u_{x'}, u_y, \{x', y\}\right)} \qquad \text{for all } x, y \in B, \tag{7}$$

with $V\left(u_x, u_y, \{x,y\}\right) := \Pr(x|u,\{x,y\})/\Pr(y|u,\{x,y\})$. The argument, building on IIA, is well-known, see e.g. McFadden (1974, p. 109). Note that $V(u_y, u_y|\{y,y\}) = 1$.

Step 2: Then, by Axiom 4 and richness $\Rightarrow$ there exists a function $\tilde{V}: \mathbb{R} \times X^2 \to \mathbb{R}$ such that $V\left(u_x, u_y, \{x,y\}\right) = \tilde{V}\left(u_x - u_y, \{x,y\}\right)$.
Axiom 4 holds for all $r \in \mathbb{R}$, which implies

$$\frac{d\Pr(x|u+r,B)}{dr} = 0 \qquad \text{for all } r \in \mathbb{R} \text{ and } B \subseteq X. \tag{8}$$

Thus, $dV\left(u_x + r, u_y + r, \{x,y\}\right)/dr = 0$ for all r, which implies $\partial V/\partial u_x = -\partial V/\partial u_y$, and by (7) it holds for all u_y, which implies $\partial^2 V/\partial u_x \partial u_y = 0$. Hence, functions V_d, f and g exist such that

$$V\left(u_x, u_y, \{x,y\}\right) = V_d\left(f(u_x) - g(u_y), \{x,y\}\right) \qquad \text{with } f'(u_x) = g'(u_y) \tag{9}$$

for all x, y. By richness, $dV\left(u_x + r, u_y + r, \{x,y\}\right)/dr = 0$ also holds for all interior u_x, u_y, which implies $f'(u_x) = g'(u_y) = \text{const}$ and a function $\tilde{V}$ exists such that

$$V\big(u_x, u_y, \{x,y\}\big) = \tilde{V}\big(u_x - u_y, \{x,y\}\big) \tag{10}$$

Step 3: Then, by richness, $V\big(u_x, u_y, \{x,y\}\big) = \exp\{\lambda \cdot (u_x - u_y) + d(x,y)\}$.
Note that given $x \in X$, Eq. (7) holds true for all $y \in X$, and by richness, this implies that there exists $y \in X$ such that

$$\frac{d}{dr} \frac{V\big(u_x, u_y + r, \{x,y\}\big)}{V\big(u_x, u_y + r, \{x,y\}\big) + 1} = 0 \qquad \text{for } r \approx 0. \tag{11}$$

Hence, $V(u_x, u_y + r|B) = V(u_x, u_y|B) \cdot f(r)$ for all interior u_x, u_y and $r \approx 0$, for some function $f : \mathbb{R} \to \mathbb{R}$ with $f(0) = 1$. By (10), this implies $\tilde{V}(u_x - u_y - r|B) = \tilde{V}(u_x - u_y|B) \cdot f(r)$ for all $r \approx 0$, and differentiating we obtain $\tilde{V}'(u_x - u_y|B) = \tilde{V}(u_x - u_y|B) \cdot f'(0)$. Solving this differential equation, $\tilde{V}(u_x - u_y|B) = \exp\{\lambda \cdot (u_x - u_y) + d\}$ using $\lambda := f'(0)$, i.e.

$$V\big(u_x, u_y, \{x,y\}\big) = \exp\{\lambda \cdot (u_x - u_y) + d(x,y)\}. \tag{12}$$

Step 4: Then, (7) holds true for $V\big(u_x, u_y, \{x,y\}\big) = \exp\{\lambda\, u_x + w(x)\}$ for some $w(x)$.
Using V as characterized in the previous step, u_y cancels out and we obtain

$$\Pr(x|u,B) = \frac{\exp\{\lambda \cdot u_x + d(x|y)\}}{\sum_{x' \in B} \exp\{\lambda \cdot u_{x'} + d(x'|y)\}}. \tag{13}$$

Since the choice probabilities are independent of the option $y \in X$ chosen as benchmark (Step 1), functions w, w_2 exist such that $d(x,y) = w(x) + w_2(y)$. Thus, $w_2(y)$ cancels out, confirming "$\Leftarrow$" in statement 2 of the Theorem.

Step 5: Then, by Axiom 5, (7) holds true for $V\big(u_x, u_y, \{x,y\}\big) = \exp\{\lambda\, u_x\}$.
Axiom 5 implies $w(x) = w(y)$ for all $x, y \in X$. For contradiction, fix x, y such that $w(x) \neq w(y)$ and fix a bijection $f : X \to X$ such that $f(x) = y$. By assumption, Axiom 5 yields $\Pr(f(x)|u,B) = \Pr\big(x\,|\,u \circ f, B\big)$, i.e.

$$\frac{\exp\{\lambda \cdot u(f(x)) + w(f(x))\}}{\sum_{x' \in B} \exp\{\lambda \cdot u(f(x')) + w(f(x'))\}} = \frac{\exp\{\lambda \cdot u(f(x)) + w(x)\}}{\sum_{x' \in B} \exp\{\lambda \cdot u(f(x')) + w(x')\}}$$

This contradicts $w(x) \neq w(x')$ if $B = \{x,y\}$, given $f(x) = y \neq x$. Hence, $w(x) = w(y)$ for all $x, y \in X$, confirming "$\Leftarrow$" in statement 1 of the Theorem. $\square$

3 Discussion

Multinomial logit is widely used to estimate utility and demand functions. McFadden (2001) argues that its appeal relates to its lean axiomatic foundation, which appears to rest largely on independence of irrelevant alternatives (IIA). The appearance is wrong, as a closer look at the axioms in McFadden (1974) has revealed. Axiom 3 implicitly assumes that binomial choice is logit, while IIA serves to extend logit to multinomial choice. The present paper provides the first self-contained axiomatic foundation of logit, uncovering two hidden assumptions: "narrow bracketing" and "true utility = decision utility". These findings have several notable implications, which are discussed in this concluding section: They clarify the behavioral foundation of logit's specific functional form, relate logit to the central debate in behavioral welfare economics, and suggest routes for future work by including presentation effects in generalized logit analyses. The implications are discussed in this order.

The specific functional form of multinomial logit has repeatedly been questioned by applied and behavioral economists. While it is typically accepted that choice is stochastic, the notion that the underlying utility perturbations have specifically an extreme value distribution, as opposed to a normal distribution or any other one, seems more difficult to accept. In light of the axiomatic foundation derived here, the underlying assumptions can now be discussed one by one. To begin with, positivity and IIA do not affect the functional form of choice propensities. Similarly, richness does not affect the functional form, as it is merely an assumption about the environment from which the analyst or experimenter constructs the budget set. Narrow bracketing postulates that the utility level realized outside the analyzed choice tasks does not affect behavior. This assumption is made almost universally in behavioral and empirical analyses (including least-squares analyses), but its validity should probably be scrutinized by future work. This promises to be a hard task, however. The most progressive existing work allows utility parameters to depend on demographic variables (see e.g. Bellemare et al., 2008, in behavioral economics), which of course does not violate narrow bracketing. Finally, the axiom "true utility = decision utility" is vividly debated in behavioral welfare economics, as reviewed in the Introduction (for an overview of the debate, see e.g. Kőszegi and Rabin, 2008b). Regardless of one's position in this debate, dropping the axiom does not fundamentally affect the functional form. Dropping it yields a generalized form of logit where the decision maker (DM) with utility function u chooses $x \in B$ with probability

$$\Pr(x|u,B) = \frac{\exp\{\lambda \cdot u(x) + w(x)\}}{\sum_{x' \in B} \exp\{\lambda \cdot u(x') + w(x')\}}. \tag{14}$$

Before I turn to discussing the non-utility related component $w(x)$, let us appreciate the difference to least squares approaches. Given utility function u, parameter vector p, and set of observations O, take any observation $o \in O$ and let $o^\star(u,p)$ denote the model's prediction given u,p in the respective choice task. The sum of squared deviations of observations from predictions is

$$SSq(p|u,O) = \sum_{o \in O} (o - o^\star(u,p))^2 = \log \prod_{o \in O} \exp\{(o - o^\star(u,p))^2\}.$$

Aside from the missing normalization toward choice probabilities, this fits generalized logit with the additional assumptions that true utility itself is choice-irrelevant ($\lambda = 0$) and that discrete choice is presentation dependent in the specific sense that the quadratic distance to the utility maximizer counts, $w(x) = (x - x^\star(u,p))^2$. These assumptions are not supported by behavioral or normative results (in the context of discrete choice). The other assumptions, narrow bracketing and IIA, are maintained in least squares analyses. Narrow bracketing is maintained in the sense that, when utility parameters are estimated, outside endowments are generally neglected in utility functions. IIA is satisfied if the utility maximizer $o^\star$ is taken to be the utility maximizer over the global environment X (which implies that it is independent of the specific budget set B imposed in the analyzed choice task). Thus, the difference between the functional forms of logit and least squares are specifically in the assumptions on the relation of decision utility and true utility. Generalized logit obtains without any restriction on the relation of true utility and decision utility and therefore imposes strictly weaker assumptions than least squares—despite the seemingly specific functional form.

Now, let us turn to the implications of dropping the axiom "true utility = decision utility". Dropping the axiom implies that decision utility regarding $x \in X$ consists of two components, the true utility, which is $u(x)$ by assumption, and a second component $w(x)$. The existing literature suggests a clear interpretation of $w(x)$. Following Kőszegi and Rabin (2008b), the difference between true utility and decision utility represents choice biases that prevent subjects from maximizing their true utility. Various such biases

Figure 1: Dictator game transfers with manual and graphical (“slider”) choice entering

(a) Andreoni and Miller (2002) (b) Fisman et al. (2007)

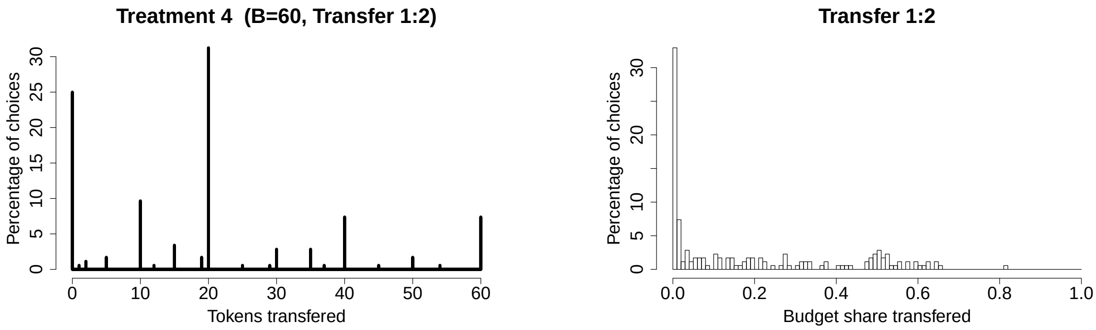

Note: In a dictator game (DG), the dictator is endowed with a budget of *B* tokens, of which he may transfer any number to an anonymous recipient. The above histograms show the distribution of choices for specific tasks in two seminal DG experiments. In these cases, for each token given up by the dictator, two tokens are added to the recipients account (approximately so in Fisman et al., 2007).

have been identified in the literature. DMs are inclined to select default options if they exist (McKenzie et al., 2006; Dinner et al., 2011; Spiegler, 2015), they respond to the ordering or positioning of options (Dean, 1980; Miller and Krosnick, 1998; Feenberg et al., 2015), they respond overly to the left-most digit in multi-digit comparison (Poltrock and Schwartz, 1984; Lacetera et al., 2012), and they focus on round numbers in numerical choice (Heitjan and Rubin, 1991; Manski and Molinari, 2010).

To illustrate the magnitude of such biases, Figure 1 provides histograms of transfers in dictator games from two experiments under mostly identical conditions (in a dictator game, the dictator chooses how much of his income to transfer to the recipient). Essentially, the only difference between the experiments is the user interface: the number of tokens to be transferred is entered either manually (Figure 1a) or graphically, via mouse and slider (Figure 1b). The choice tasks are otherwise virtually equivalent, but the differences in choice patterns are drastic. The manual entry of numbers induces strong round number effects, which in turn biases parameter estimates. Specifically, analyses of dictator games typically assume that true utility is CES altruism and that it equates with decision utility. CES altruism does not account for the behavioral implications of the user interface, and thus, even if true utility u indeed is CES altruism, decision utility differs. Given the above results on generalized logit, these differences are captured by the non-utility component w, and neglecting it biases estimates of true utility u. This is in principle known at least since Heitjan and Rubin (1991), but discrete choice models capturing such effects have not yet been proposed.

Generalized logit provides a versatile model to study non-utility components of choice propensities, as it is generally applicable without imposing ex-ante restrictions on the relation of true utility and decision utility. The best-known choice biases are those discussed above and result from the presentation of choice tasks. The idea that presentation affects choice is also the basis of nudging and thus behavioral welfare economics. In the context of presentation effects, it is intuitive to refer to $w(x)$ as the *focality* of option x, or as its mental luminescence. Being the default affects focality, positioning of options affects focality, the left-most digit does, as does the roundness of numbers. The focality $w(x)$ of option x may well be defined in absolute terms, independently of the focality of other objects, similarly to physical luminescence and thus honoring IIA. The behavioral impact of an option's focality depends on the focality of the other options in the budget set, of course, analogously to the utility of options, which is defined in absolute terms but can be interpreted only in relation to the utilities of other options. Thus, neither utility nor focality necessarily violate IIA.[8]

Currently, little is known about the focality-related component $w(x)$ of choice propensities predicted by generalized logit. Major biases are identified, however, and intuitive representations can be constructed for each bias. For example, $w(x)$ may be a function of x's positioning, left-most digit, level of roundness, or its relation to the default. This, in turn, illustrates the potential of further theoretical, empirical, and experimental work, which is intriguing. Generalized logit provides a first versatile model to study focality and choice biases, and at the same time, it is highly tractable and relates closely to the models used currently (multinomial logit or generalizations such as nested logit). Thus, it allows choice biases to be studied in the same manner that utility has been studied for decades. Generalized logit is equally applicable to study biases in consumer choice (e.g. left-digit bias) or strategic choice (e.g. to control for the impact of ordering of actions). Understanding choice biases, i.e. understanding how $w(x)$ depends on presentation, in turn allows to obtain unbiased estimates of the true utility $u(x)$ in behavioral economics, of the willingness to pay in consumer choice, and the depth of reasoning in strategic choice. Combined, this provides a general approach to analyze and predict nudging effects. Thus, generalized logit promises to substantially facilitate work in behavioral economics, empirical economics, and behavioral welfare.

[8]Having said that, maintaining IIA is simply one possibility here. The examples of choice biases discussed above mostly do not indicate that IIA must be relaxed, probably with the exception of "positioning" depending on context, but other biases such as similarity effects or "salience" generally require relaxing IIA (for recent discussions, see e.g. Kőszegi and Szeidl, 2013, and Matejka and McKay, 2015).

References

Andreoni, J. and Miller, J. (2002). Giving according to GARP: An experimental test of the consistency of preferences for altruism. *Econometrica*, 70(2):737–753.

Bajari, P. and Hortacsu, A. (2005). Are structural estimates of auction models reasonable? evidence from experimental data. *Journal of Political Economy*, 113(4):703–741.

Bellemare, C., Kröger, S., and Van Soest, A. (2008). Measuring inequity aversion in a heterogeneous population using experimental decisions and subjective probabilities. *Econometrica*, 76(4):815–839.

Benkert, J.-M. and Netzer, N. (2015). Informational requirements of nudging. Technical report, CESifo Group Munich.

Bernheim, B. D. (2009). Behavioral welfare economics. *Journal of the European Economic Association*, 7(2-3):267–319.

Bernheim, B. D. and Rangel, A. (2007). Toward choice-theoretic foundations for behavioral welfare economics. *American Economic Review*, 97(2):464–470.

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

Berry, S., Levinsohn, J., and Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica*, 63(4):841–890.

Block, H. D. and Marschak, J. (1960). Random orderings and stochastic theories of responses. *Contributions to probability and statistics*, 2:97–132.

Cappelen, A., Hole, A., Sørensen, E., and Tungodden, B. (2007). The pluralism of fairness ideals: An experimental approach. *American Economic Review*, 97(3):818–827.

Dean, M. L. (1980). Presentation order effects in product taste tests. *The Journal of psychology*, 105(1):107–110.

Dinner, I., Johnson, E. J., Goldstein, D. G., and Liu, K. (2011). Partitioning default effects: why people choose not to choose. *Journal of Experimental Psychology: Applied*, 17(4):332.

Feenberg, D. R., Ganguli, I., Gaule, P., Gruber, J., et al. (2015). It's good to be first: Order bias in reading and citing nber working papers. Technical report, National Bureau of Economic Research, Inc.

Fisman, R., Kariv, S., and Markovits, D. (2007). Individual preferences for giving. *American Economic Review*, 97(5):1858–1876.

Fudenberg, D. and Strzalecki, T. (2015). Dynamic logit with choice aversion. *Econometrica*, 83(2):651–691.

Goeree, J. K., Holt, C. A., and Palfrey, T. R. (2003). Risk averse behavior in generalized matching pennies games. *Games and Economic Behavior*, 45(1):97–113.

Gul, F. and Pesendorfer, W. (2001). Temptation and self-control. *Econometrica*, 69(6):1403–1435.

Gul, F. and Pesendorfer, W. (2007). Welfare without happiness. *American Economic Review*, 97(2):471–476.

Gul, F. and Pesendorfer, W. (2008). The case for mindless economics. *The foundations of Positive and normative Economics: A handbook*, pages 3–42.

Hanemann, W. M. (1984). Welfare evaluations in contingent valuation experiments with discrete responses. *American journal of agricultural economics*, 66(3):332–341.

Heitjan, D. F. and Rubin, D. B. (1991). Ignorability and coarse data. *Annals of Statistics*, pages 2244–2253.

Holt, C. A. and Laury, S. K. (2002). Risk aversion and incentive effects. *American Economic Review*, 92(5):1644–1655.

Keane, M. P. (2010). Structural vs. atheoretic approaches to econometrics. *Journal of Econometrics*, 156(1):3–20.

Kőszegi, B. and Rabin, M. (2007). Mistakes in choice-based welfare analysis. *The American economic review*, 97(2):477–481.

Kőszegi, B. and Rabin, M. (2008a). Choices, situations, and happiness. *Journal of Public Economics*, 92(8):1821–1832.

Kőszegi, B. and Rabin, M. (2008b). Revealed mistakes and revealed preferences. *The foundations of positive and normative economics: a handbook*, pages 193–209.

Kőszegi, B. and Szeidl, A. (2013). A model of focusing in economic choice. *The Quarterly journal of economics*, 128(1):53–104.

Lacetera, N., Pope, D. G., and Sydnor, J. R. (2012). Heuristic thinking and limited attention in the car market. *American Economic Review*, 102(5):2206–2236.

Luce, R. (1959). *Individual choice behavior*. Wiley New York.

Luce, R. D., Ng, C., Marley, A., and Aczél, J. (2008). Utility of gambling i: entropy modified linear weighted utility. *Economic Theory*, 36(1):1–33.

Luce, R. D. and Suppes, P. (1965). Preference, utility, and subjective probability. In Luce, R. D., Bush, R. R., and Galanter, E., editors, *Handbook of Mathematical Psychology, Vol. III*, pages 252–410. Wiley, New York.

Manski, C. F. and Molinari, F. (2010). Rounding probabilistic expectations in surveys. *Journal of Business & Economic Statistics*, 28(2):219–231.

Matejka, F. and McKay, A. (2015). Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review*, 105(1):272–98.

McFadden, D. (1974). Conditional logit analysis of qualitative choice models. *Frontiers of Econometrics, ed. P. Zarembka. New York: Academic Press*, pages 105–142.

McFadden, D. (1980). Econometric models for probabilistic choice among products. *The Journal of Business*, 53(3):13–29.

McFadden, D. (2001). Economic choices. *American Economic Review*, 91(3):351–378.

McKenzie, C. R., Liersch, M. J., and Finkelstein, S. R. (2006). Recommendations implicit in policy defaults. *Psychological Science*, 17(5):414–420.

Miller, J. M. and Krosnick, J. A. (1998). The impact of candidate name order on election outcomes. *Public Opinion Quarterly*, 62:291–330.

Poltrock, S. E. and Schwartz, D. R. (1984). Comparative judgments of multidigit numbers. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 10(1):32.

Rabin, M. and Weizsäcker, G. (2009). Narrow bracketing and dominated choices. *American Economic Review*, 99(4):1508.

Read, D., Loewenstein, G., and Rabin, M. (1999). Choice bracketing. *Journal of Risk and Uncertainty*, 19(1-3):171–97.

Rust, J. (2010). Comments on: "structural vs. atheoretic approaches to econometrics" by Michael Keane. *Journal of Econometrics*, 156(1):21–24.

Small, K. and Rosen, H. (1981). Applied welfare economics with discrete choice models. *Econometrica*, 49(1):105–130.

Spiegler, R. (2015). Choice complexity and market competition. *Annual Review of Economics*.