RJEC 969910—8/10/2014—MOHANRAJ.D—496300———Style 2

Journal of Economic Methodology, 2014
Vol. 00, No. 0, 1–21, http://dx.doi.org/10.1080/1350178X.2014.969910

The *consistency* and *ecological rationality* approaches to normative bounded rationality

Nathan Berg*

Department of Economics, University of Otago, PO Box 56, Dunedin, New Zealand

(*Received* ■; *accepted* ■)

This paper focuses on tacit versus explicit uses of plural performance metrics as a primary methodological characteristic. This characteristic usefully distinguishes two schools of normative analysis and their approaches to normative interpretations of bounded rationality. Both schools of thought make normative claims about bounded rationality by comparing the performance of decision procedures using more than one performance metric. The consistency school makes tacit reference to performance metrics outside its primary axiomatic framework, but lexicographically promotes internal axiomatic consistency as the primary, and in most cases sufficient, normative outcome with which to undertake welfare comparisons. The consistency school's axiomatization program, in both neoclassical and behavioral forms, pre-commits to welfare interpretations that follow a hierarchy of rationalities based on the stringency of restrictions that different axiomatizations impose on choice data. In contrast, the ecological rationality school explicitly adopts multiple, domain-specific performance metrics, reflecting the view that adequate descriptions of well-being are irreducibly multivariate (i.e., non-scalar).

[AQ1]

1. Introduction

One strand in the bounded rationality literature focuses almost exclusively on internal consistency as the litmus test for rationality, which I refer to as the *consistency school* of normative bounded rationality research. In contrast, the ecological rationality school applies multiple performance metrics, explicitly delimiting the class of decision-making environments in which a particular combination of normative criteria is specifically rather than universally relevant. Consistency axioms include: transitivity of preferences required for decisions over vector-valued elements of a choice set to be representable as maximization of a scalar-valued objective function; the Savage axioms, required to guarantee that choices over random payoff distributions are representable as expected utility maximization; the Bayesian and Kolmogorov axioms, which require internal consistency of conditional and unconditional beliefs (with respect to the definitions of conditional and unconditional probability) – without requiring subjective beliefs to be accurate with respect to objective probability distributions; and the assumption of time consistency in dynamic choice models, a common assumption in intertemporal choice models that justifies exponential discounting and has evolved into a rationality axiom among bounded rationality researchers who propose quasi-hyperbolic discounting as a model of myopia and imperfect willpower.[1] The consistency school's move into behavioral economics includes empirical studies identifying empirical inconsistencies,

*Email: nathan.berg@otago.ac.nz

where observed choice behavior is incompatible with standard rational choice models, and a very active program of theoretical modeling based on new axiomatizations and their weaker restrictions on observable choice data. Frequently motivated as finally providing formal or rigorous underpinnings for previous models of bounded rationality, these attempts to axiomatize bounded rationality can be described as seeking to rationalize particular deviations from standard axiomatic rationality (e.g., regularities in choice data that appear anomalous through the lens of rational choice theory) by postulating weaker (i.e., less restrictive) lists of bounded rationality axioms.

This paper seeks to describe a new taxonomy of normative methodology consisting of two distinct approaches to bounded rationality. The consistency school of bounded rationality includes many well-known economists who describe their work as 'behavioral' while defining rationality just as neoclassical economics do: solely as internal logical consistency. Names such as Kahneman, Thaler, Diamond, and numerous others are associated with this strand of bounded rationality, targeting deviations from consistency as the primary phenomenon to be analyzed and, in so doing, maintaining consistency as the central normative criterion for rationality.

In its formalizations of normativity, the consistency school claims that logical consistency provides an exhaustive characterization of rationality and is singular as the ultimate normative concern. Even when obvious multiplicities of performance metrics present themselves as intuitively relevant for evaluating well-being, consistency is applied lexicographically, trumping other measures that could be used to compare how well different decision procedures perform. For example, although 'rational' (Nash) versus 'boundedly rational' or 'irrational' (non-Nash) strategies by experimental participants are commonly reported empirical outcomes, behavioral game theorists rarely report average or cumulative payoffs comparing the two. It often turns out, however, that less consistent decision procedures in strategic settings (e.g., non-Nash versus Nash strategies) and in games against nature (e.g., time-inconsistent versus time-consistent) earn more money.

In this case, we have two distinct normative measures. The first is a discrete, binary outcome: axiomatically inconsistent versus consistent. The second normative measure, cumulative or average earnings, is a non-axiom-based performance metric. Although earnings may be more relevant for real-world normative analysis, the primary normative yardstick for assessing rationality in much of the bounded rationality literature focuses solely on axiomatic consistency.

Given observed choice data from one group (A) that satisfy a particular characterization of (bounded) rationality and another group (B) whose choice data do not, normative analysis of the consistency school tells us that the normative variable of interest is the discrete outcome of belonging to group A or B. But what if both groups earn roughly the same amounts of money? Or what if the irrational group B earns more? Would our normative analysis then be better served by analyzing correlates and predictors of group-A versus group-B status, which may not in fact matter very much, or the correlates and predictors of high versus low earnings?

Justifying the consistency school's perspective by affirming the real-world relevance of axiomatic consistency (i.e., the normative relevance of group-A versus group-B status), would seem to require having access to plural normative metrics (e.g., comparing decision procedures that conform with bounded rationality axiomatizations of varying stringency in other normative units of performance, such as wealth, health, or happiness).[2] In many well-known games such as prisoner's dilemma, centipede, and the trust game, the so-called irrational non-Nash strategies (which sub-optimally forgo higher individual payoffs conditional on the other player's strategy) achieve higher individual and aggregate payoffs, through the joint interaction of non-best-response strategies. If decision procedures that violate more stringent axiomatic requirements of bounded rationality generate more money, health, or happiness than those that conform to more stringent axiomatic characterizations of rationality, then we face an interesting tension among plural normative metrics. Evaluating which metric provides more compelling prescriptive notions of 'ought' is an important question. Does it make sense to ignore this tension and compare decision procedures solely by consistency, applying it lexicographically as the fundamental methodological prior of normative economic theory?

When bounded rationality investigations report rankings of people's rationality according to differing degrees of conformity or non-conformity with axioms that guarantee best-response strategies – without reporting the realized payoffs associated with each discrete category of axiomatic rationality – critical information is lost. Ranking performance by payoffs (and investigating factors that influence them) in the actual environment that participants face – where others are not necessarily playing best-response strategies – might provide more important normative information about well-being. If people's choice data fail the test of consistency according to an axiomatic notion of bounded rationality but, in so doing, achieve higher payoffs in units of another normative performance metric such as dollars, then the normative appeal of axiomatic bounded rationality may be limited. Moreover, empirical characterizations of people's rationality in terms of a spectrum of conformity over nested sets of axiomatic rationalities (e.g., Manzini & and Mariotti, 2010) are likely to provide incomplete, if not distorted, information about how well people's repertoires of behavioral rules actually produce well-being.

By defining rationality and bounded rationality in terms of different degrees of stringency with respect to logical consistency, these axiomatic characterizations of bounded rationality stake their normative claims on the idea that it is the extent to which choice data are internally consistent that exhaustively characterizes a decision procedure's performance. The consistency school's justifications for its formalizations of bounded rationality (expressed as weakened, or less stringent, sets of consistency requirements, within which perfect rationality is nested) turn out, however, to depend on auxiliary normative performance metrics that contradict the axiomatization project's claims of sufficiency and exhaustiveness. Methodological contradictions among the normative claims in axiomatizations and the auxiliary normative performance metrics they tacitly introduce, such as in money-pump arguments and in the nested hierarchies of axiomatic characterizations of bounded rationality, are analyzed in detail in later sections.

A second school of thought in the taxonomy of normative approaches to bounded rationality is referred to as the *ecological rationality* school. In contrast to the consistency school, the ecological rationality school embraces and makes explicit use of a plurality of normative criteria. The justification for pluralism with respect to normative standards follows from the observation that well-being is a multi-dimensional phenomenon and therefore its characterization requires multiple measures. Pluralistic description and prescription regarding how well different decision procedures perform reflects the multiple and oftentimes incommensurable components of well-being (i.e., those that cannot be traded off against one another and compressed to a universal scalar-valued performance unit).

In the ecological rationality school, multiple normative criteria are required to characterize the rationality of decisions, inferences, and institutions, depending on the environment or *decision domain*. Prime examples of these multiple normative criteria for describing and comparing the ecological rationality of two decision strategies or two decision-makers' choice data would include dollar payoffs, life expectancy, health outcomes, self-reported happiness, accuracy of beliefs with respect to objective frequencies, and many others, depending on context. Ecological rationality is a matching concept that does not universally apply the label 'rational' to a procedure for making decisions or inferences. Rather, ecological rationality requires a good-enough match between a decision procedure and the environment in which it is used. Characterizations of ecological rationality require a description of decision procedures, the decision environment, and performance metrics in units of measure whose levels can be compared. When a decision procedure is well matched to an environment, where 'well matched' is defined as achieving good-enough levels on the performance metrics relevant to that environment, then the pair (*decision procedure*, *environment*) is classified as ecologically rational.[3]

The primary differences between the approaches in these two schools of normative methodology are: (1) explicit use of multiple performance metrics in normative analysis; (2) a view about whether multiple normative criteria weaken or strengthen economic theory; and (3) the kinds of real-world problems to which each school's methods can be applied.

One way to evaluate the question of how these two schools of normative methodology succeed in their common goal of real-world relevance is the extent to which researchers are able to exploit the respective multiplicities of performance metrics that both schools use – tacitly in the case of the consistency school and explicitly in the case of the ecological rationality school. The criterion of real-world relevance cannot be easily dismissed as ad hoc, by virtue of the fact that both schools claim real-world relevance and improved veridicality of their models' assumptions as justifications for their methodological choices.[4]

2. Money-pump arguments and normative claims about bounded rationality

Grüne-Yanoff (2004) argues that rationality axioms based on consistency axioms such as transitivity do not follow as a necessary implication of broader notions of optimization. Therefore, he disagrees with those who interpret observed violations of axioms such as WARP, SARP, and GARP (in experimental settings where behavioral economists frequently test the power of theories of bounded rationality to explain observed violations) as evidence against optimization, sometimes referred to as the neoclassical maximization [AQ2] hypothesis.

In contrast, I argue that normative analysis can do away with the maximization hypothesis and not automatically conclude that consumers are irrational. Optimization is not required for describing what people do when they adaptively seek to make improvements in their own well-being. Violation of the maximization hypothesis is to be expected among adaptive agents who seek to: learn more about their own goals; expend costly effort to change their goals and preferences; explore which actions are available, sometimes making important discoveries that shift observed choice patterns; and receive new information about the payoffs that those actions are expected to achieve in a fundamentally non-static environment. Grüne-Yanoff (2004) seems to imply that rationality requires maximization, but that the objective being maximized changes frequently (e.g., beliefs may change and some choices involve indifference versus incommensurability that standard axiomatic formulations do not distinguish between); therefore, violation of consistency axioms is not evidence against maximization. Although I agree with the conclusion of this statement, my position is that rationality requires adaptation and experimentation, and that violation of consistency axioms is not evidence against the hypothesis that people are purposefully pursuing adaptive improvements in their well-being (which is not the same thing as maximization but shares with Grüne-Yanoff the normative view that inconsistency does not imply pathological choice).

The failures of the revealed preference project in economics that Grüne-Yanoff (2004) observes suggest that normative analysis faces a still-open question: how to describe the decision process that people use to purposefully improve well-being, whether that process is maximization (which it could be in many instances), adaptation and satisficing, or something else altogether. Observed violations of transitivity tell us only that the decision process we are trying to observe is not maximization of a static preference relation. Ruling out static optimization based on observed choice data that violate preference axioms effectively rules out only a tiny subset in the universe of purposeful decision processes, and this finding does not seem to have helped economics learn very much about the question of how people decide and what influences those decisions.[5]

The case of prospect theory illustrates how the methodological commitment to the hypothesis that economic agents must be maximizing some objective function (if only we could discover which one) serves to circumscribe rather than expand economists' investigation of how people make choices and respond to changes in policy variables that are common targets of normative analysis. From prospect theory, we learn that a determined maximizer *could possibly* exhibit Allais' paradox by maximizing an objective function with asymmetric psychological values assigned to positive and negative changes in monetary payoffs and a weighting scheme based on non-linear transformation of probabilities. As a thought experiment, prospect theory suggests a new possibility that was surprising in light of early interpretations of Allais' paradox (and other violations of expected utility theory) as evidence of failure to maximize anything at all. Allais' paradox simply demonstrates that a decision-maker does not choose risky gambles by maximizing expected utility. Prior to prospect theory, there were few, if any, *alternative* maximization stories (i.e., with an objective function that differed from that of expected utility theory) that could rationalize Allais' paradox. Prospect theory raised interesting new questions regarding the normative interpretation of reference points, asymmetry in the subjective evaluation of gains and losses, and whether the probability-weighting function represented perceptual distortions akin to optical illusions.

Despite these insights from prospect theory (similar to other famous *how-possibly* explanations,[6] e.g., Schelling's work revealing a surprising mechanism capable of generating neighborhood segregation without intergroup animus), critics such as Gigerenzer question what we learn about the way human minds work by continuing down the path of repairing broken constrained optimization models with modifications of functional forms and new psychological parameters in the constraint set. Instead, he argues for abandoning constrained optimization and developing more veridical models of decision processes. Gigerenzer is concerned that attempts to rationalize so-called anomalies by introducing more flexible objective functions or weaker sets of bounded rationality axioms – while hanging onto the core methodological tenet of constrained optimization – distract from more important normative analysis. He argues that ecological rationality helps us better understand how high-stakes decisions are actually made and how attempts to design the decision-making environment are likely to affect the pluralistic behavioral decision rules and outcome measures needed to evaluate domain-specific performance.

Gigerenzer criticizes bounded rationality models that add new parameters to make the objective functions and constraint sets used in standard decision models more flexible. Greater flexibility leads to improved statistical fit, but without necessarily revealing mental process. Introducing weakened axioms that rationalize larger sets of choice data (and therefore 'explain' anomalous behavioral patterns) is, in Gigerenzer's view, a similarly mechanical and unrevealing exercise.

One question that appears to be infrequently asked is what justification there is, in terms of measures of performance and well-being, for applying the more-is-better notion to axiomatic consistency. Those who propose a hierarchy of rationalities ranked by stringency of axiomatic consistency, implying that people whose choice data conform to more stringent axiomatic consistency are somehow better off, face criticisms from those who see little evidence or theoretical justification for prescriptive interventions to 'de-bias' irrational people (irrational, because they violate consistency axioms; e.g., Jolls & Sunstein, 2006; Jolls, Sunstein, & Thaler, 1998). Is there compelling evidence that de-biasing campaigns, or interventions that would cause people's choice data to conform with more stringent sets of rationality axioms, would improve well-being?

2.1. *Tacit multiplicity of normative yardsticks used to motivate rationality axioms*

The standard money-pump argument attempts to justify the claim that rational decision-makers should satisfy the transitivity axiom. Common to justifications for axiomatic definitions of both bounded and neoclassical rationalities, the money-pump argument demonstrates how these justifications rely on at least one auxiliary normative performance metric. The relevance of this for bounded rationality axiomatizations that do *not* require transitivity is that multiple normative criteria – outside the axioms themselves – are used in arguing for why an axiom is appealing. It is therefore instructive to re-examine the multiplicity of normative yardsticks used in justifying transitivity, because this multiplicity reappears tacitly in proposals for hierarchies of bounded rationalities based on the stringency of axioms (even those that allow for intransitivity). The money-pump argument uses money as an external normative measure in support of being consistent with transitivity. Similarly, those who advocate hierarchies of axiomatic rationality must somehow link positions on this ranking of axiomatic rationalities to at least one external metric of performance. For example, those whose choice data exhibit intransitive cycles but not menu effects (Manzini & Mariotti, 2010) demonstrate a greater degree of axiomatic consistency without achieving transitivity. The proposal that these relative positions in a hierarchy of axiomatic bounded rationalities convey information relevant to welfare economics implies that these ranked positions can be linked to some other measure of well-being external to the axioms themselves. Without the link from conformity with an axiomatized rationality to an external performance metric, these rankings in the hierarchy of rationalities may not be normatively relevant.

According to the money-pump argument, intransitive people suffer because they are willing to make sequences of trades that leave them with no money. It is entirely reasonable to investigate correlations between different decision-making procedures and the levels of wealth they produce. Wealth, however, bears no logical connection to transitivity (or, as it turns out, to other consistency axioms). As I will argue in the next section, greater wealth does not imply greater degrees of internal consistency, and internal consistency does not imply greater wealth. The textbook labor–leisure tradeoff – based squarely on the rational-choice model and conforming perfectly to rational preference axioms – teaches, after all, that utility maximization does not imply money maximization.

The axiomatic approach to rationality, whether orthodox rational choice or its bounded rationality variants, is only interesting insofar as it predicts or correlates with well-being and performance measures that we care about. We could instead study what influences the well-being and performance measures we care about more directly, however, without checking for consistency with a set of axioms. If checking for axiomatic consistency provided a useful shortcut to predicting normative outcomes we care about, then the ecological rationality school and presumably many others would use it. These tests of axiomatic consistency do not, however, provide reliable information about normative outcomes that people and scientists who study them typically care about.

There are many possible mechanisms that can, in theory, reward intransitivity. For example, non-static payoff environments may give inconsistent decision-makers an advantage in discovering new information and opportunities. In environments where decision-makers do not know the payoffs associated with all elements of their choice set (perhaps because the environment is occasionally shocked in ways that shift the mapping from actions into payoffs), intransitive cycling may help detect shocks and diversify risk (Bookstaber & Langsam, 1985).

2.2. Logical inconsistency of wealth as an auxiliary performance metric in the money-pump argument

Transitivity is neither sufficient nor necessary for wealth. If *X* prefers *A* (earn nothing and live in mom's basement) over *B* (working minimum-wage and living alone), and if *X* prefers *B* over *C* (having a million dollars and living with someone he hates), then *X* may be perfectly transitive (i.e., *A* preferred over *C*), yet one would consistently observe *X* choosing *A* and consequently having no money. Transitive people, according to textbook labor–leisure tradeoffs, can consistently choose to have no money. Having no money does not imply that one has been money-pumped.

We also observe many intransitive choices among people with good incomes. Intransitive agents may be more willing to (perhaps inconsistently) choose higher-risk endeavors with higher expected returns. Or perhaps highly consistent individuals consistently save less, supply less labor, make impatient time tradeoffs in favor of cash flows with lower present value, or are consistently more risk-averse, therefore accumulating less wealth.

The previous paragraphs argued that: (1) transitivity does not imply positive wealth and (2) intransitivity does not imply having zero wealth. Therefore, the degree to which one is transitive and the extent of one's money holdings are logically unrelated. Although wealth can be an important normative metric for evaluating how well different decision procedures perform, the degree to which people conform to consistency axioms such as transitivity is manifestly not a welfare measure. Using wealth as an implicit performance metric in support of accepting transitivity as a rationality axiom, as money-pump arguments do, is logically inconsistent.

Hierarchies of rationality based on different sets of axioms with different degrees of stringency face a wholly analogous problem for those who propose using them to make normative comparisons. Just as we lack theoretical and empirical links that reliably associate degrees of conformity with the transitivity axiom to wealth, proponents of new axiomatizations of bounded rationality who argue for their normative interpretation face the same problem. Does conforming or diverging from more and less stringent axiomatic formulations of bounded rationality provide any interesting information about well-being, or show us how people ought to make decisions?

2.3. Characterizations of bounded rationality using consistency axioms do not provide the units of measure needed for meaningful normative analysis

There is another more subtle methodological contradiction in the money-pump argument's appeal to wealth as an auxiliary performance metric. The problem is more general: in arguing why an axiom should be regarded as such (i.e., assented to without evidence, or accepted without explicit testing), it is necessary that those justifications appeal to other normative criteria to avoid circularity. But the methodological premise in axiomatization programs is the logical sufficiency of the axioms as a characterization of rationality. If the compelling normative principle is, for example, wealth, then why not simply study the correlates of high-wealth-producing decision procedures and rank those procedures according to the wealth they produce? An even more serious problem is that consistency axioms, in addition to being logically unrelated to wealth, are (as far as the available evidence has shown) at best only weakly – and sometimes inversely – related to a broad spectrum of important performance metrics that empirical investigators of well-being have identified.

Consider, for example, accumulated wealth, lifespan, self-reported happiness, rates of illness, and measures of social well-being (Bruni & Porta, 2007) such as the number of people who can be counted on to provide shelter or lend a car in the event of an emergency. One advantage of these non-consistency-based performance metrics is that they are measured in freestanding units that are easy to interpret. They also facilitate straightforward interpersonal comparison. Unlike consistency-based norms that can rationalize any single choice analyzed in isolation and impose internal restrictions only on sets of two or more choices, non-consistency-based performance metrics provide normative scales that can be applied to single choices, acts, and inferences.

Consistency norms can facilitate interpersonal comparison (e.g., the fraction of observed choices or number of days when an individual's choice data satisfy transitivity). They are rarely used this way in welfare economics, however. It is far from clear that doing so would reveal anything interesting about well-being or economic performance. Regarding accumulated wealth on standard risk and time preference decision tasks in the experimental laboratory, my own work (Berg, Biele, & Gigerenzer, 2013; Berg, Eckel, & Johnson, 2014) reveals the opposite of what the hierarchy of rationalities view, using consistency-based definitions of rationality and bounded rationality, would predict. Consistent risk and time preferences are negatively correlated with cumulative payoffs in experimental decision tasks. Those who conform to expected utility theory's axioms do so by consistently avoiding risk and earning lower-than-average expected returns. Those who conform to time consistency tend to be consistently impatient, sacrificing cash flows with larger present value in favor of smaller earlier payments. Consistent Bayesian beliefs about cancer risks correlate with less accurate subjective beliefs (Berg, Biele, & [AQ3] Gigerenzer, 2008).

Although axiomatic formulations of bounded rationality adopt less stringent requirements of internal consistency, internal consistency remains the sole arbiter of normative evaluation in this research program (Manzini & Mariotti, 2007, 2010, 2012). Rather than adopting consistency as the singular standard of rationality (even though this singular standard shows up in different forms across multiple attempts to axiomatize bounded rationality), those working on normative implications of bounded rationality may discover that consistency has less to do with performance and well-being than what is promised by the axiomatization program's claims to provide exhaustive characterizations of rationality.

The methodological regularity of the consistency school using multiple normative criteria informally, while claiming that only one criterion – internal consistency – is needed to formally define bounded rationality, can be stated more generally. We observe proponents of axiomatic bounded rationality arguing for the intuitive appeal of consistency axiom α (asking us to assent to and therefore regard α as axiomatic) by claiming: agents who violate α (e.g., transitivity) will be worse off according to the auxiliary normative performance-metric μ (e.g., wealth). To avoid the tautology, agents who violate α will perform poorly according to the standard of not violating α, axiomatic characterizations of rationality *must* invoke and make reference to multiple normative concepts.

If I am correct that a non-tautological argument in favor of axiomatic rationality *requires* external evaluation by making reference to at least one auxiliary normative criterion μ, then the goal of providing rigorous foundations and exhaustive characterizations of bounded rationality by introducing the weakened axiomatization α' cannot succeed. By succeed, I mean succeed methodologically according to its own criterion of rigorously and exhaustively characterizing rationality as a set of allowable behavioral patterns (realizations of choice data) consistent with α'. Conforming to α' is justified because, relative to irrationality (i.e., not conforming to any set of axioms on the list of axiomatizations considered), it improves performance according to μ. The external or auxiliary metric μ is not part of the normative framework logically implied by α', however. How, then, can α' be regarded as exhaustive, complete, and self-contained as a normative characterization of bounded rationality?

Regarding the relationship between bounded and neoclassical rationality axioms, one observes that axiomatizations of bounded rationality typically nest perfect rationality as a special case. For example, the sequential choice axioms α' of Manzini and Mariotti (2007, 2010), when strengthened by the additional requirement of transitivity, contain neoclassical preferences α as a subset in terms of rationalizable choice data. Choice data that satisfy neoclassical preferences necessarily satisfy Manzini and Mariotti's axioms. Similarly, prospect theory α', when restricted to linear probability weighting and no loss aversion (i.e., $w(p) = p$ and $\lambda = 1$, in Tversky and Kahneman's, 1992, notation) with a globally concave value function, contains risk-averse expected utility preferences α as a special case. Expected utility theory with risk-neutral preferences, once again, recovers an earlier normative standard of expected value maximization. What ties these models together is the mathematical operation of weighted averaging, which may have little to do with the mental processes actually used to make important decisions over risky lotteries.

Another example is the quasi-hyperbolic beta-delta function (Phelps & Pollak, 1968), which – as interpreted by David Laibson and behavioral economists advocating its use as a utility function that captures time inconsistency – plays the role of the weakened axiom (or more flexible functional form) α'. Used in the bounded rationality literature to represent bounded willpower, it shows up as a technical generalization of a mathematically convenient discounting model (rather than a veridical description of people's mental processing when facing intertemporal tradeoffs), which contains standard exponential discounting α as a special case (when the hyperbolic discounting parameter $\beta = 1$).[7]

Non-circular justifications of bounded rationality axioms (α') and rational choice axioms nested as special cases (α) must, in general, refer to auxiliary metrics of performance (μ). As with the standard definition of rational preferences as completeness and transitivity, the consistency axioms used to characterize bounded rationality make logical errors in both directions: these axiomatizations are both too strong and too weak. They are too strong in that they rule out inconsistent behavior that nevertheless achieves high performance according to μ. And they are too weak because choices that score badly according to μ are permitted as satisfying rationality defined in terms of axiomatic consistency. If μ is the performance metric that speaks to economists' intuition, then the methodological question that remains unanswered in the consistency school's axiomatization program is: why not directly study high-μ versus low-μ behavior and its correlates?

The axiomatization program in bounded rationality seeks to universally and exhaustively characterize a hierarchy of rationalities – not by directly ranking levels of performance μ but by the stringency (i.e., smallness) of allowably consistent actions. Three nested axiomatizations of (bounded) rationalities, $\alpha \Rightarrow \alpha' \Rightarrow \alpha''$, is interpreted as: choice data that satisfies α are more rational than those that satisfy α', which are more rational than those satisfying α''). There are more sets of actions (choice datasets) rationalizable by α' (α'') than by α (α'). Because the set of α-rationalizable choice patterns is more stringent, we are asked to interpret choice data in this smaller set to be more rational than choice data consistent only with α' but violating α. Choice data consistent with α'' but violating α' is less rational still.

There is no indication in Manzini and Mariotti (2010, 2012), for example, of where this descending hierarchy of rationalities stops, except by the following ad hoc rule. Choice data that are not rationalizable by the least stringent list of axioms that the authors happen to include in the considered set of axiomatizations are regarded as irrational. Thus, if only two flavors of bounded rationality, α' and α'', are present in this hierarchy, then α'' serves as the de facto boundary separating bounded rationality from irrationality. This boundary is a theoretical artifact of the modelers' choice of which list of axiomatic definitions to include in their analysis.

If this hierarchy of rationalities provided a compelling normative standard or pointed toward a new way of doing welfare economics, then one would expect economists to have accumulated a large body of empirical evidence linking intransitivity, expected utility violations, non-Bayesian beliefs, and non-Nash play – all frequently studied forms of inconsistency – to substantially diminished well-being. We await the arrival of such evidence.

Axioms themselves are not the target of the criticism above. The relevance of normative analysis rests on objects analyzed by axioms and the units of measure that can be associated with them. For example, Sen's (1985, 1991) and Sugden's (2004, 2008) normative analyses axiomatize rankings of choice *sets* rather than *choices* themselves. These authors' axiomatizations of partial orders on choice *sets* contribute substantively to welfare economics precisely because a plurality of normative concepts is explicitly applied.

3. Contrasts between consistency and ecological rationality schools

The methodological value of the ecological rationality school's approach to normative analysis is explained in Gigerenzer and Selten (2001), Smith (2003), Berg (2003, 2010, in press-a), Gigerenzer (2004), Berg and Gigerenzer (2006, 2007, 2010), Gigerenzer, Todd and The ABC Group (1999), Todd, Gigerenzer, and The ABC Research Group (2012), and
[AQ5] Hertwig, Hoffrage, and The ABC Research Group (2013). These authors argue that Herbert Simon's normative view of bounded rationality is misinterpreted by the consistency school.

According to one widespread interpretation of Herbert Simon (e.g., Jolls et al., 1998), bounded rationality can be understood in terms of three challenges to neoclassical assumptions of unboundedess: (1) unbounded cognitive capacity, (2) unbounded willpower or self-control, and (3) unbounded self-interest. The consistency school interprets departures from assumptions (1) and (2) as irrationality, maintaining the assumption, as neoclassical economists do, that full rationality requires (1) and (2). According to this view, individuals who fail (1) and (2) are pathological because they cannot achieve as high a level in their respective objective functions as unboundedly rational individuals can.

Simon (1978) writes that the very notion of cognitive capacity – and the associated notion of cognitive limitation – is not meaningfully defined without making reference to the environment in which decisions are made. If, for example, food is uniformly distributed on a plane representing an organism's choice set over which it searches for food, then vision and memory are unlikely to influence nutritional performance or be of any benefit in terms of improving the organism's search procedure. Because random search achieves nutritional targets just as well as those that use vision and memory, there is no sense in which improved vision and memory should be required for assessing the rationality of the decision-maker and its search procedures in that environment. Human contexts where forgetting and ignoring information can be beneficial appear in the ecological rationality approach to modeling bounded rationality (Berg, Abramczuk, & Hoffrage, 2013; Berg & Hoffrage, 2008, 2010; Berg, Hoffrage, & Abramczuk, 2010).

Ecologically rational behavior is required to be well matched or well calibrated to the environment in which it is used. This matching concept that defines ecological rationality permits *some* generalizations of the form: 'decision procedure *D* performs well according to the performance metric μ in the set of environments *E*.' Universality of a single normative concept is explicitly rejected, however, in the ecological rationality school, because performance is quantified in units of measure that apply to a specific set of environments. Analysis of ecological rationality requires one to carefully circumscribe the range of environments in which a given behavioral strategy performs adequately or otherwise.

Gilboa (in press) writes in support of pluralistic approaches rather than the *one-axiom-fits-all-contexts* approach to normative analysis characteristic of both neoclassical and behavioral economics. Similarly, Gilboa, Postlewaite, and Schmeidler (2009, p. 288) write: [AQ6]

> We reject the view that rationality is a clear-cut, binary notion that can be defined by a simple set of rules or axioms. There are various ingredients to rational choice. Some are of internal coherence, as captured by Savage's axioms. Others have to do with external coherence with data and scientific reasoning. [W]e should be prepared to have conflicts between the different demands of rationality. [...] But the quest for a single set of rules that will universally define the rational choice is misguided.

3.1. Bounded rationality's move from purely descriptive to explicitly normative

Tversky and Kahneman (1986) argued for a research program that maintains strict separation between normative and descriptive analysis. Contemporary behavioral economics enthusiastically undertook this program in the 1980s and 1990s, whose ground rules held that anomalous descriptive findings should not raise doubts about the normative authority of neoclassical rationality axioms. Thaler (1991) described the research program in behavioral economics explicitly in this way, going to great pains to reassure unconvinced readers that behavioral economics posed no threat to the neoclassical normative framework and, in fact, had nothing to add to normative economics since the singular normative standard of adherence to consistency axioms had already reached a state of perfection (Berg, 2003).

Tversky and Kahneman (1986), in the conclusion of their article, however, suggested a role for anomalous behavioral findings to influence policy, by helping those who deviate from the orthodox normative model to better conform. After Thaler's (1991) article – as the descriptive finding became more widely accepted that consistency axioms frequently failed empirical tests and top-ranked economics journals began publishing more behavioral papers – Thaler changed his view regarding the limited normative scope of behavioral economics. He went on to argue for enlarging the behavioral program (which was previously limited to purely descriptive phenomena) by launching a new normative research program. The new normative research program he advocated, however, was to adhere strictly to the consistency school's approach to bounded rationality, applying the descriptive finding of widespread irrationality (i.e., violations of consistency axioms) to produce prescriptive policy advice. He has been joined by many other behavioral economists since.

Rather than a sea change that opened normative inquiry to a wider range of outcomes and measurement techniques, behavioral economists' vast empirical literature on 'biases' and 'deviations' from axiomatic norms of rationality – expected utility violations, preference reversals, time inconsistency, and non-Nash play in laboratory games – appears to have hardened the normative authority of neoclassical rationality based on consistency axioms. Axiomatic rational choice models may be descriptively wrong, the thinking goes, but they nevertheless provide reliable guidance about what people ought to do.

Critics of the consistency school's normative behavioral economics include some who have contributed substantially to the axiomatization literature in economics, such as Gilboa and Sugden. In his article titled, 'Why incoherent preferences do not justify paternalism,' Sugden's (2008) argues for challenging the normative status of axiomatic rationality. He suggests that we can accept the descriptive validity of data showing violations of consistency axioms without accepting the common normative interpretation that such violations motivate new rationalizations for paternalistic policies (Sugden, 2004).

When normative theory and observed behavior come into conflict, behavioral economics typically follows the research program laid out in Tversky and Kahneman (1986) by unequivocally attributing error to the agent who violates consistency. That is not the only valid deduction, however, based on conflict between consistency-based normative theory and observed behavior. One might instead conclude that those normative principles previously thought to have prescriptive value are simply incomplete, or perhaps have a more limited range of applicability.

3.2. Optical illusions are not a good analogy for violations of consistency axioms

Tversky and Kahneman (1986) put forward optical illusions as an analogy for behavioral anomalies (i.e., choice data that violate standard choice axioms). Optical illusions are deviations between perceived versus objectively measured distances (or visual information). The implication of the analogy is that the axiomatic foundation of normative decision theory is as objectively grounded as the measure of physical distance.

Thaler (1991, p. 138) writes,

> It goes without saying that the existence of an optical illusion that causes us to see one of two equal lines as longer than the other should not reduce the value we place on accurate measurement. On the contrary, illusions demonstrate the need for rulers!

Yet, in documenting that observed behavior deviates from the assumptions (and predictions) of axiomatic rationality, there is, in fact, no analog in decision theory to the straight lines of objectively equal length. Unlike straightforward geometric verification of equal lengths against which incorrect perceptions can be compared, the fact that human decisions do not satisfy rationality axioms in no way implies an illusion or mistake.

Normative analysis of bounded rationality would benefit from pursuing new normative criteria that classify decision procedures in ways that help assess whether they are well matched to the environments in which they are used according to the principle of ecological rationality (Gigerenzer & Selten, 2001; Smith, 2003). Instead of new normative criteria and measures of performance, the consistency school maintains axiomatic consistency as the reference point against which deviations comprise the normative outcome.

4. Case study: Manzini and Mariotti's consistency approach to bounded rationality

Manzini and Mariotti (2007, 2010, 2012) provide an opportunity to demonstrate methodological contrasts between the consistency and ecological rationality schools. Despite these authors' stated goal of bridging the respective normative perspectives of these schools, I will argue that these three papers of Manzini and Mariotti (MM) can be classified unambiguously within the consistency school.

MM undertake to axiomatize boundedly rational sequential choice by specifying different sets of axioms that rationalize different choice procedures. These procedures include the Rational Shortlist Method, Sequential Rationalizable Choice, and Categorize Then Choose. Common to these distinct choice procedures is that they first shrink the choice set to a smaller consideration set before finally arriving at a choice. In Sequential Rationalizable Choice, the orderings (possibly partial) used in pre-choice stages, referred to as *rationales*, represent a sequence of reasoning that permits intransitive cycles. Violations of invariance are not permitted, however. Violations of invariance occur when the inclusion of a strictly dominated element in the choice set reverses the ranking of two others. A hierarchy then follows in which perfect rationality requires transitivity, bounded rationality allows some violations (i.e., intransitive cycles) but not others (violation of invariation, referred to as menu effects). Irrationality is a residual category indicated by choice data that are not rationalizable either by perfect or bounded rationality axioms.

MM (2010) propose that the axiomatizations which they consider can provide ‘rigorous underpinnings’ for non-compensatory or lexicographic heuristics, such as Take-The-Best, in Gigerenzer, Todd and The ABC Group (1999). MM (2010, 2012) also claim that their axiomatic characterizations of bounded rationality help advance welfare economics. Although there is much to admire technically and in MM's (2012) frankness regarding gaps between their axiomatizations of bounded rationality and veridical mental processes, they are explicit in insisting on consistency as the sole normative criterion of importance. Consequently, their claims about bridging gaps toward the ecological rationality approach and advancing welfare economics appear invalid.

The Take-The-Best heuristic makes a binary inference by considering a vector of cues or signals, one at a time following a sequence that is particular to the context in which it is used. Each cue may or may not be decisive in pointing toward an inference, a feature that is well captured by the partial orders that MM use as rationales. When a cue is decisive, this single-reason prediction supersedes or trumps all subsequent cues in the sequence. As soon as one cue makes an inference, the heuristic stops, and all subsequent cues are ignored, which Gigerenzer refers to as one-reason decision-making. The theory that ignoring information can improve performance draws on theoretical, agent-based, and experimental studies (Baucells, Carrasco, & Hogarth, 2008; Berg & Hoffrage, 2008, 2010; Berg et al., 2010; Bookstaber & Langsam, 1985; Kameda, Tsukasaki, Hastie, & Berg, 2011). Although MM and Gigerenzer agree that limiting the information required to make choices and inferences is intuitively appealing, the normative criteria they use are in conflict.

Within the context of the consistency school, MM no doubt make substantive contributions by identifying an axiomatic framework that can account for (i.e., provide a characterization of) decision procedures which proceed in steps, reducing the size of an unmanageably large choice set before arriving at a final decision. Many decisions studied in the social sciences follow this structure: Maslow's hierarchy of needs; home buyers who impose non-compensatory restrictions to shrink the feasible set of affordable homes down to a smaller consideration set; Yee, Dahan, Hauser, and Orlin's (2007) analysis of how consumers [AQ7] choose mobile phones; or Berg's (2014) analysis of business owners' location choice.

Gigerenzer's approach is to investigate how the mind might work (i.e., heuristics) and then undertake normative analysis by testing which classes of environments a heuristic performs well in. In contrast, MM are primarily interested in the gap between what the heuristic user is allowed to do according to a particular set of axioms versus what an omniscient, perfectly rational agent would have done, which produces a labeling scheme describing a ranked hierarchy of axiomatic rationalities. The main finding in MM (2007) is to characterize a 'family of boundedly rational choice procedures that can account for these observed anomalies' (p. 1824).

The word 'anomalies' is important for identifying MM's normative model, which is most explicit in MM (2012). Anomalies refer to choice data that reveal inconsistencies (i.e., violations of rational choice axioms). By their account, it is the gap between observed human behavior and the norm of axiomatic consistency that makes the phenomenon interesting and, as it turns out, pathological: MM (2010) repeatedly label inconsistency as irrational, categorizing several distinct forms of alleged irrationality with respect to neoclassical consistency. MM's stark characterization of inconsistency as pathological reveals that their core normative standard is rationalizability with respect to a set of consistency axioms. Despite Manzini and Mariotti's sympathy for Gigerenzer, Selten, and the work of others in the ecological rationality school, the normative and descriptive content of MM fall squarely in the consistency school.

Like other rationalizability and representation theorems in the consistency school's axiomatization program, the axioms themselves offer no *how-actually* explanations. Instead, the *characterizations* that rationalizability and representation theorems provide are demonstrations from a thought experiment: if decision-makers conformed to a set of axioms, then the following restrictions on observed choice data would be satisfied. MM (2007) are correct that their model makes testable predictions, where prediction means: *if* people conform to the axioms, *then* certain inconsistencies should not be observed. A multi-category empirical test of this form is the main investigation in MM (2010).

In contrast to MM's use of the term, Gigerenzer reserves the term *prediction* to mean out-of-sample predictive accuracy. Gigerenzer's heuristics predict specific choices and inferences, whereas MM (2012) classify the observation space into *subsets* – not specific choices – that are consistent with different axiomatic characterizations of rationality, irrationality, or intermediate forms of quasi-consistency labeled boundedly rational.[8]

The opening lines of Manzini and Mariotti (2010) state: 'If people are irrational, how are they irrational? And how can we describe their behavior and perform welfare analysis?' Unfortunately, MM's theory makes no out-of-sample predictions about well-being or performance (i.e., no predicted difference among rational, boundedly rational, and irrational people's dollars, years of life, happiness, or percentage-point deviations measuring inverse objective accuracy). One is therefore hard-pressed to see how their model – whose 'description' of behavior consists of labeling choice data according to the sets of axioms they satisfy – helps us 'perform welfare analysis.'

Suppose we had datasets for every person in a population measuring which bucket of axiomatic rationality they fall into. What could an applied business decision-maker do with that? What policy question facing an organization, city, or state, or what theoretical issue in welfare economics about efficiency of market outcomes, could be advanced by having data tagging each individual's conformity with sets of axioms in MM (i.e., each person's choice data labeled *rational*, *boundedly rational*, or *irrational*)?

One is left to wonder, because the word ‘welfare’ appears only in a footnote and the concluding paragraph of MM (2010): ‘To conclude, we believe that the [model] offers one possible solution to the hard problem of welfare analysis in the context of boundedly rational choice.’

Perhaps, MM have in mind that their characterization of bounded rationality (violating transitivity while satisfying other weaker consistency axioms) could be used to make people’s choice data conform to weaker standards of bounded rationality, although no justification as to why this would be individually or collectively desirable is provided. Or perhaps their categorization of choice data into different buckets, labeled as distinct types of irrationality, is an implicit suggestion for others to quantify welfare by assigning smaller welfare losses to those whose data exhibit less severe inconsistencies (i.e., better versus worse flavors of irrationality).

Despite MM’s professed interest in non-compensatory heuristics and their commendable discovery of a weaker set of consistency axioms that rationalize some of these heuristics, the over-riding normative criterion in their analysis is no different than neoclassical economics: internal consistency. MM do not embrace a pluralistic toolkit of normative metrics that reveal how well decision procedures perform in specific environments – in units of welfare that connect to recognizable (popular or scientific) conceptions of thriving and living well. A final problem with MM’s labeling scheme and hierarchy of rationalities is the suggestion that people who conform to stricter consistency axioms are more rational and therefore enjoy superior well-being. The link between subject and predicate in this proposition appears without substantiation.

In Lipsey and Lancaster’s (1956–1957) theory of the second best, the ‘negative corollary’ states: when moving from an outcome with $V > 1$ violations of the constraints required for Pareto optimality to a new outcome where one of the previously unfulfilled constraints is now satisfied (i.e., $V - 1 > 0$ constraints violated), a Pareto improvement need not follow.

By analogy, everyone may be worse off when ascending the hierarchy of imperfect rationalities (Güth & Kliemt, 2001). MM provide no theoretical link explaining why conforming to increasingly stringent axioms of bounded rationality makes us individually or collectively better off.

4.1. Example of wrong inferences and policies based on using the wrong mental model

Figure 1 presents a simplified non-compensatory decision tree for choosing a mobile phone (oversimplified for the purpose of illustration) to demonstrate several points in favor of non-compensatory or lexicographic decision-tree models that are fundamental to Gigerenzer’s methodological point of view. The example draws on Yee et al.’s (2007) study of shopper behavior when choosing mobile phones from a website featuring 100 mobile phones that can be compared along 16 hedonic features. The combinatorics of pairwise rankings for all 100 phones along 16 features is an overwhelming task: $16 \times 100!/(2! \times 98!) = 79{,}200$ pairwise comparisons. To economize on time and cognitive effort, it stands to reason – according to both Manzini and Mariotti (2007) and Gigerenzer, Todd, and The ABC Group (1999) – that a smart shopping strategy *could* proceed by imposing a few non-compensatory restrictions that effectively shrink the consideration set from 100 phones to a smaller, more manageable, set.

In the model shown in Figure 1, the non-compensatory restriction imposed at the first stage concerns weight: if a phone is too heavy, it is excluded from consideration and its other features (e.g., price) do not influence choice whatsoever. Other possible non-

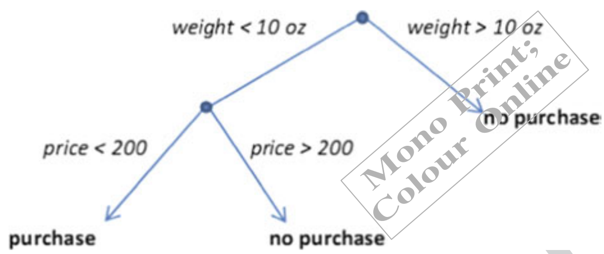

Figure 1. Decision tree for purchasing a mobile phone (or reducing a large feasible set to a smaller consideration set).

compensatory threshold conditions that could be included in a more veridical model might include phones' memory, speed, camera features, color, and network and contract details.

Table 1 supposes that the feasible set contains only four phones. The phone in the first row is light enough (9 ounces $<$ 10 ounces) and therefore makes it to the next branch of the decision tree. But this phone turns out to be too expensive (\$250 $>$ \$200) resulting in a 'no purchase' decision. The phone in the second row is both lightweight and inexpensive enough, resulting in 'purchase' decision. The third and fourth phones listed in Table 1 are much cheaper, but – and here is the point that distinguishes non-compensatory models from virtually all choice models based on tradeoffs – the price variation on phones listed in rows 3 and 4 play absolutely no role, because these phones are too heavy and therefore eliminated from the consideration set at an earlier stage. The phones in rows 3 and 4 violate the weight threshold, and the decision tree therefore discards them from the consideration set irrespective of price.

When an economist encounters choice data generated by people employing a non-compensatory decision procedure, but mistakenly uses the standard toolkit of choice models, then their models are likely to make wrong inferences about the effects of phone characteristics on purchase decisions, no matter how flexible the functional form. The reason the inferences are wrong is because the model of mind is wrong. Models that allow all features of the phone to trade off against each other wrongly assume that the mind pays attention to and integrates variation among all right-hand-side features. To demonstrate this point, consider what the statistician who estimates a compensatory probabilistic choice model based on the choice data in Table 1 (generated by the non-compensatory decision process in Figure 1) 'learns' from his mis-specified model. A regression of purchase decisions on price and weight (linear probability of purchase conditional on price and weight) based on Table 1 data is:

$$\text{Probability of purchase} = -2.509 + 0.201 \text{ ounces} + 0.008 \text{ price}.$$

The mis-specified compensatory regression model implies that the consumer is more likely to purchase a phone when it is heavier and more expensive. Mis-specified estimates

Table 1. Cues and purchase outcomes according to the non-compensatory decision.

Weight in ounces	Price in dollars	Purchase decision ($y = 1$ if yes, $y = 0$ if no)
2	250	0
9	199	1
11	50	0
12	0	0

pointing in the wrong direction are qualitatively the same for probit, logit, and virtually any other probability model based on a nonlinear transformation of a linear index in phone features, because the linear index makes the model compensatory, whereas the true mental process is non-compensatory. Statistical significance is not at issue either, because if consumers make non-compensatory choices based on the decision tree in Figure 1, then scaling up the number of observed purchasers (essentially replicating the data in Table 1 any number of times) will result in an arbitrarily high degree of statistical significance for the mis-specified model. More generally, estimating a compensatory model when the data-generating process is in fact non-compensatory, as this stylized example shows, can lead to mistaken inferences – both theoretical *and* empirical – and high-cost policy mistakes.

5. Conclusion

Economists working on models of bounded rationality in both the consistency and ecological rationality schools draw inspiration from Herbert Simon, although the emphasis they place on consistency versus adaptation and the normative criteria they use to evaluate what it means to make good decisions are different. Bounded rationality in the consistency school shares with neoclassical economics the central methodological importance of modeling behavior as if it were a solution to a well-defined optimization problem based on a scalar-valued objective function subject to constraints. Bounded rationality in the consistency school also shares with neoclassical economics an a priori or definitional commitment to internal logical consistency as *the* singular normative criterion than trumps all others. By including a larger number of free parameters than the neoclassical models nested within them as a parameter restriction, models of bounded rationality in the consistency school would appear to allow for a wider range of descriptive possibilities. Agents whose boundedly rational behavior is specified as an optimal choice rule that depends on ‘psychological’ parameters measuring bias with respect to the neoclassical ideal of perfect consistency can be sorted into classes of behavior that violate perfect neoclassical consistency in various ways or by varying degrees. The singular normative standard of consistency in bounded rationality models that allow descriptively for different degrees of conforming to consistency axioms is the same normative standard that guides much of the revealed preference literature and neoclassical economics more generally.

In contrast, the ecological rationality school explicitly uses multiple normative criteria and is consequently more eclectic.[9] There is a difference between explicitly embracing multiple normative criteria (fitting the definition of the ecological rationality school) versus implicit or tacit application of multiple normative criteria (which does not fit the definition of ecological rationality and is characteristic of analysis relating to rationality axioms in the consistency school). This distinction has substantive implications and is not merely stylistic.

The consistency school argues that individuals whose choice data conform to axiomatic rationality, denoted α (representing a short list of axioms such as transitivity and completeness), are most rational; those whose choice data conform to a less stringent, or boundedly rational, set of axioms, denoted α', are moderately rational; and those whose choice data violate α' (necessarily violating α as well, because the sets of choices it allows are nested in that of α') are irrational. By applying this formulation to welfare economics, the consistency school's approach to bounded rationality tacitly assumes that there are auxiliary performance metrics (wealth, in the case of money-pump arguments) that depend on this axiomatic hierarchy of rationalities. An example of tacitly referring to multiple normative metrics while insisting there is only one singularly important metric among them is demonstrated by researchers who measure degrees of conformity with axioms on the one hand, and wealth on the other, but then characterize rationality solely by the axiomatic standard – even though conformity with axioms and wealth could are substantively different metrics that may, in principle, exhibit positive, negative or zero correlation. I presented such examples in Manzini and Mariotti's work, which characterizes bounded rationality as an intermediate degree of conformity with respect to axiomatic rationality.

The implicitness or non-specificity regarding the question of which performance metrics really matter leaves open a wide interpretive gap regarding how, for example, lawmakers and regulators might (mis-)apply this framework based on a hierarchy of axiomatic rationalities. Are we to believe (without evidence) that raising a population's rate of conformity with consistency axioms would lead to greater wealth or other substantive improvements in well-being? Unfortunately, normative assertions based on axiomatic characterizations of bounded rationality appear to rest only on vague suggestions that rankings of different individuals' degrees of conformity with lists of axioms (based on differing degrees of internal consistency) can tell us how to do welfare economics. The opacity of suggested and implied mappings from any proposed hierarchy of axiomatic rationalities into meaningful performance metrics (like wealth or health) winds up obscuring how changes in the environment (especially policy changes which are the object of normative analysis) might affect the individual and aggregate outcomes that determine well-being or social welfare (measured as multivariate outcomes comprised of multiple indicators or performance metrics in the ecological rationality school).

Absent strong evidence that axiomatic characterizations of rationality based on internal consistency can tell us what we want to know about well-being, our normative analysis would do better (from the view of the ecological rationality school) to put aside the axioms and instead directly study those individual and aggregate outcomes that matter for well-being (possibly multivariate characterizations thereof). If conformity with consistency axioms were strongly correlated with other performance metrics (e.g., wealth, health or accuracy) in a particular class of decision tasks (where logical consistency is highly rewarded), then axiomatic rationality might, in principle, serve as one component used to evaluate ecological rationality. Ecological rationality attempts to understand which factors influence the multiple decision processes that produce well-being. As uncontroversial as that perhaps sounds, it appears that disagreement on the singularity of consistency versus the normative pluralism of ecological rationality may likely continue to influence and circumscribe the ways in which economists characterize bounded rationality and the normative interpretations given to them.

Notes

1. Some observers distinguish axioms made for technical convenience (e.g., continuity of preference orderings) from what are – arguably – more substantive requirements of rationality (e.g., ruling out intransitive cycles). This distinction is not clear-cut, however, according to Gigerenzer's (1991) tools-to-theory heuristic, describing how ideas that first appear in the social sciences as technical tools, such as linear regression and Bayes' rule, later reappear as veridical descriptions of mental or cognitive process. The tools-to-theory heuristic is an apt description of shifts that took place during the last half of the twentieth century in the interpretation of preference axioms. Axiomatizations of rational preferences over commodity bundles (or risky payoff lotteries) were first introduced as technical requirements guaranteeing that a preference ordering could be represented as utility (or expected utility) maximization. These technical requirements subsequently evolved into broad normative interpretations for how well-functioning minds ought to make decisions, with the novel implication that experimental tests of consistency with respect to preference axioms could be interpreted as a litmus test for individual rationality.

2. Analogous calls for improved units of measure in evaluating well-being appears in Karabell's (2014) critique of GDP and its numerous shortcomings for making scalar-valued comparison across countries and through time.
3. It is worth emphasizing that the criterion of achieving good-enough levels on the performance metrics relevant to a particular environment easily accommodates multivariate characterizations of well-being. For example, suppose that good-enough levels of both material wealth and health are joint requirements for success. If a choice falls below an appropriately calibrated threshold of health, then no improvements in income can satisfy this environment-specific definition of success, because it incommensurably requires minimum levels of both wealth and health. Similarly, improvements in health cannot offset insufficient wealth. This example underscores that 'good-enough' or satisficing profiles of vector-valued normative metrics are perfectly compatible with normative analysis in the ecological rationality school based on notions of adaptive success that depend incommensurably on multiple performance metrics.
4. Real-world relevance is an explicit motivation that appears repeatedly in the writing of many, but not all, researchers working on models of bounded rationality. Rubinstein (2001), for example, rejects the notion that the success of theoretical models should be assessed in terms of real-world relevance or applicability to public policy.
5. For example, despite decades of research, do we know enough to feel confident that monetary policy pursuing zero or negative interest rates will stimulate economic growth? While this might appear to be a non sequitur given that microeconomic choice models are the focus in most of the literature on revealed choice, the methodological commitment to maximization of a stable objective function (whether neoclassical or behavioral) profoundly influences (or rationalizes) a wide range of high-stakes policy decisions. If economists considered a broader universe of behavioral decision processes, then prediction might become more difficult. But doing so would also expand consideration of unintended consequences (allowing, for example, the possibility that agents abruptly shift their response rules to policy tools) to help avoid policy mistakes and perhaps discover simpler approaches that would achieve policy makers' goals – by better understanding how consumers and investors actually make decisions, what they respond to in particular environments, and what they ignore.
6. In his defense of non-representationality, Grüne-Yanoff (2013) distinguishes *how-possibly* from *how-actually* explanations. He argues that both classes of explanations, each with their respective role in generating meaningful learning opportunities, can function productively – even complementarily – in theorizing and model making.
7. See Rubinstein (2003) for more detailed critique and a procedural alternative that he advocates as intuitive and reasonable although it violates multiple consistency axioms.
8. Anand (1993) argues that choices made from choice sets with different numbers of elements (e.g., MM's choice sets of size 2, 3, and 4) cannot be compared and that inconsistent rankings which depend on size of the choice set are to be expected and should not be regarded as irrational.
9. Names such as Gigerenzer, Selten, and Vernon Smith can easily be associated with the ecological rationality school, although the taxonomy's boundaries can be interpreted more broadly. For example, institutional economists such as Veblen (and many since) drew on criticisms of rational choice made by psychologists to suggest methodological changes in the normative framework underlying neoclassical economics. Similarly, evolutionary economists seeking to explain innovation, creativity and entrepreneurship undertake to expand their normative analysis beyond that of neoclassical consistency as the singular definition of rationality. Many economists explicitly embrace multiple normative metrics applying them specifically in well-defined domains rather than universally, which places their normative framework squarely in the ecological rationality school, even though such authors do not apply that label to describe their own work: for example, Witt (2003), Weizsäcker (2005), Schubert [AQ17] (2012), and Wegner (2009). A similarly rich normative literature interrogating competing notions of rationality appears in the Austrian economics literature (e.g., Caplan, 2000).

References

Anand, P. (1993). The philosophy of intransitive preference. *The Economic Journal*, *103*, 337–346.

Backhouse, R. E. (1998). If mathematics is informal, then perhaps we should accept that economics [AQ8] must be informal too. *The Economic Journal*, *108*, 1848–1858.

Baucells, M., Carrascom, J. A., & Hogarth, R. M. (2008). Cumulative dominance and heuristic performance in binary multi-attribute choice. *Operations Research*, *56*, 1289–1304.

Berg, N. (2003). Normative behavioral economics. *Journal of Socio-Economics*, *32*, 411–427.

Berg, N. (2010). Behavioral economics. In R. C. Free (Ed.), *21st century economics: A reference handbook* (Vol. 2, pp. 861–872). Los Angeles, CA: Sage.

[AQ9] Berg, N. (in press-a). Gerd Gigerenzer. In M. Altman (Ed.), *Encyclopedia of behavioral economics*.

Berg, N. (in press-b). Success from satisficing and imitation: Entrepreneurs' location choice and
[AQ11] [AQ10] implications of heuristics for local economic development. *Journal of Business Research*.

Berg, N., Abramczuk, K., & Hoffrage, U. (2013). Schelling's neighborhood segregation model with FACE-recognition. In R. Hertwig, U. Hoffrage, & The ABC Research Group (Eds.), *Simple heuristics in a social world* (pp. 225–257). New York: Oxford University Press.

Berg, N., Biele, G., & Gigerenzer, G. (2013). *Does consistency predict accuracy of beliefs?: Economists surveyed about PSA*. Working Paper, University of Otago.

Berg, N., Eckel, C., & Johnson, C. (2014). *Inconsistency pays?: Time-inconsistent subjects and EU violators earn more*. Working Paper, University of Otago.

Berg, N., & Gigerenzer, G. (2006). Peacemaking among inconsistent rationalities. In C. Engel & L. Daston (Eds.), *Is there value in inconsistency?* (pp. 421–433). Baden-Baden: Nomos.

Berg, N., & Gigerenzer, G. (2007). Psychology implies paternalism?: Bounded rationality may reduce the rationale to regulate risk-taking. *Social Choice and Welfare*, *28*, 337–359.

Berg, N., & Gigerenzer, G. (2010). As-if behavioral economics: Neoclassical economics in disguise? *History of Economic Ideas*, *18*, 133–166.

Berg, N., & Hoffrage, U. (2008). Rational ignoring with unbounded cognitive capacity. *Journal of Economic Psychology*, *29*, 792–809.

Berg, N., & Hoffrage, U. (2010). Compressed environments: Unbounded optimizers should sometimes ignore information. *Minds and Machines*, *20*, 259–275.

Berg, N., Hoffrage, U., & Abramczuk, K. (2010). Fast acceptance by common experience: FACE-recognition in Schelling's model of neighborhood segregation. *Judgment and Decision Making*, *5*, 391–410.

Binder, M. (2010). *Elements of an evolutionary theory of welfare: Assessing welfare when*
[AQ12] *preferences change*. Hoboken, NJ: Taylor and Francis.

Binder, M. (2013). Should evolutionary economists embrace libertarian paternalism? *Journal of*
[AQ13] *Evolutionary Economics*. Advance online publication.

Bookstaber, R., & Langsam, J. (1985). On the optimality of coarse behavior rules. *Journal of Theoretical Biology*, *116*, 161–193.

Bruni, L., & Porta, P. L. (Eds.). (2007). *Handbook on the economics of happiness*. Gloucestershire: Edward Elgar.

Caplan, B. (2000). Rational irrationality: A framework for the neoclassical-behavioral debate. *Eastern Economic Journal*, *26*, 191–211.

Gigerenzer, G. (1991). From tools to theories: A heuristic of discovery in cognitive psychology. *Psychological Review*, *98*, 254–267.

Gigerenzer, G. (2004). Fast and frugal heuristics: The tools of bounded rationality. In D. J. Koehler & N. Harvey (Eds.), *Blackwell handbook of judgment and decision making* (pp. 62–88). Oxford: Blackwell.

Gigerenzer, G., & Selten, R. (Eds.). (2001). Rethinking rationality.. *Bounded rationality: The adaptive toolbox* (pp. 1–12). Cambridge, MA: MIT Press.

Gilboa, I., Postlewaite, A., & Schmeidler, D. (2009). Is it always rational to satisfy Savage's axioms? *Economics and Philosophy*, *25*, 285–296.

Grüne-Yanoff, T. (2004). The problems of testing preference axioms with revealed preference theory. *Analyse & Kritik*, *26*, 382–397.

Grüne-Yanoff, T. (2012). Old wine in new casks: Libertarian paternalism still violates liberal
[AQ14] principles. *Social Choice and Welfare*, *38*, 635–645.

Grüne-Yanoff, T. (2013). Appraising models non-representationally. *Philosophy of Science*, *80*, 850–861.

Güth, W., & Kliemt, H. (2001). *From full to bounded rationality: The limits of unlimited rationality*. Discussion papers, interdisciplinary research project 373: Quantification and simulation of economic processes.

The ABC Research Group, Hertwig, R., & Hoffrage, U. (2013). *Simple heuristics in a social world*. Oxford: Oxford University Press.

RJEC 969910—8/10/2014—MOHANRAJ.D—496300———Style 2

Jolls, C., & Sunstein, C. R. (2006). Debiasing through law. *Journal of Legal Studies*, *35*, 199–241.

Jolls, C., Sunstein, C. R., & Thaler, R. H. (1998). A behavioral approach to law and economics. *Stanford Law Review*, *50*, 1471–1541.

Kameda, T., Tsukasaki, T., Hastie, R., & Berg, N. (2011). Democracy under uncertainty: The wisdom of crowds and the free-rider problem in group decision making. *Psychological Review*, *118*, 76–96.

Karabell, Z. (2014). (Mis)leading indicators: Why our economic numbers distort reality. *Foreign Affairs*, March/April. Retrieved from http://www.foreignaffairs.com/articles/140749/zachary-karabell/misleading-indicators

[AQ15] Kitcher, P. (1992). The naturalists return. *Philosophical Review*, *101*, 53–114.

Lipsey, R. G., & Lancaster, K. (1956–1957). The general theory of second best. *Review of Economic Studies*, *24*, 11–32.

Manzini, P., & Mariotti, M. (2007). Sequentially rationalizable choice. *The American Economic Review*, *97*, 1824–1839.

Manzini, P., & Mariotti, M. (2010). *Revealed preferences and boundedly rational choice procedures: An experiment*. Working Paper, University College London.

Manzini, P., & Mariotti, M. (2012). Categorize the choose: Boundedly rational choice and welfare. *Journal of the European Economic Association*, *10*, 1141–1165.

Phelps, E. S., & Pollak, R. A. (1968). On second-best national saving and game-equilibrium growth. *Review of Economic Studies*, *35*, 185–199.

[AQ16] Rubinstein, A. (1998). *Modeling bounded rationality*. Cambridge, MA: MIT Press.

Rubinstein, A. (2001). A theorist's view of experiments. *European Economic Review*, *45*, 615–628.

Rubinstein, A. (2003). Economics and psychology? The case of hyperbolic discounting. *International Economic Review*, *44*, 1207–1216.

Schubert, C. (2012). Is novelty always a good thing?: Towards an evolutionary welfare economics. *Journal of Evolutionary Economics*, *22*, 585–619.

Sen, A. K. (1985). Well-being, agency and freedom: The Dewey lectures 1984. *The Journal of Philosophy*, *82*, 169–221.

Sen, A. K. (1991). Welfare, preference and freedom. *Journal of Econometrics*, *50*, 15–29.

Simon, H. A. (1978). Rationality as process and as product of thought. *The American Economic Review*, *68*, 1–16.

Smith, V. L. (2003). Constructivist and ecological rationality in economics. *American Economic Review*, *93*, 465–508.

Sugden, R. (2004). The opportunity criterion: Consumer sovereignty without the assumption of coherent preferences. *The American Economic Review*, *94*, 1014–1033.

Sugden, R. (2008). Why incoherent preferences do not justify paternalism. *Constitutional Political Economy*, *19*, 226–248.

Thaler, R. H. (1991). *Quasi rational economics*. New York: Russell Sage Foundation.

The ABC Research Group, Todd, P. M., & Gigerernzer, G. (2012). *Ecological rationality: Intelligence in the world*. Oxford: Oxford University Press.

Tversky, A., & Kahneman, D. (1986). Rational choice and the framing of decisions. *Journal of Business*, *59*, S251–S278.

Wegner, G. (2009). Substantive versus procedural liberalism: Exploring a dilemma of contemporary liberal thought. *Journal of Institutional and Theoretical Economics*, *165*, 535–557.

Weizsäcker, C. C. (2005). *The welfare economics of adaptive preferences*. Preprint, Max Planck Institute for Research on Collective Goods, No. 05/11.

Witt, U. (2003). Economic policy making in evolutionary perspective. *Journal of Evolutionary Economics*, *13*, 77–94.

Yee, M., Dahan, E., Hauser, J. R., & Orlin, J. (2007). Greedoid-based noncompensatory inference. *Marketing Science*, *26*, 532–549.