

Deliberately Stochastic[†]

By SIMONE CERREIA-VIOGLIO, DAVID DILLENBERGER,
PIETRO ORTOLEVA, AND GIL RIELLA*

We study stochastic choice as the outcome of deliberate randomization. We derive a general representation of a stochastic choice function where stochasticity allows the agent to achieve from any set the maximal element according to her underlying preferences over lotteries. We show that in this model stochasticity in choice captures complementarity between elements in the set, and thus necessarily implies violations of Regularity/Monotonicity, one of the most common properties of stochastic choice. This feature separates our approach from other models, e.g., Random Utility. (JEL D80, D81)

A robust finding in the study of individual decision making is the presence of stochastic, or random, choice: when subjects are asked to choose from the same set of options multiple times, they often make different choices.¹ An extensive literature has documented this pattern in many experiments, in different settings and with different populations, both in the lab and in the field. It often involves a significant fraction of the choices, even when subjects have no value for experimentation (e.g., when there is no feedback), or when there are no bundle or portfolio effects (e.g., when only one choice is paid).² It thus appears incompatible with the typical assumption in economics that subjects have a complete and stable preference

*Cerrei-a-Vioglio: Department of Decision Sciences, Università Bocconi, Via Sarfatti, 25 Milano, Italy, and IGIER (email: simone.cerreia@unibocconi.it); Dillenberg: Department of Economics, University of Pennsylvania, Philadelphia, PA 19104 (email: ddill@sas.upenn.edu); Ortoleva: Department of Economics and Woodrow Wilson School, Princeton University, Princeton, NJ 08544 (email: pietro.ortoleva@princeton.edu); Riella: Escola de Políticas Públicas e Governo and Brazilian School of Public and Business Administration Getúlio Vargas Foundation, SGAN, Quadra 602, Brasília/DF/Brazil, 70830-051 (email: gil.riella@fgv.br). Jeffrey Ely was the coeditor for this article. We thank Eddie Dekel, Kfir Eliaz, Mira Frick, Faruk Gul, Ryota Iijima, Peter Klibanoff, Jay Lu, Efe Ok, Wolfgang Pesendorfer, Marciano Siniscalchi, and Tomasz Strzalecki for their useful comments. The coeditor and four anonymous referees provided valuable comments that improved the paper significantly. Cerrei-a-Vioglio gratefully acknowledges the financial support of ERC grant SDDM-TEA, Ortoleva gratefully acknowledges the financial support of NSF grant SES-1763326, and Riella gratefully acknowledges the financial support of CNPq of Brazil, grant 304560/2015-4. Part of this work was done while Dillenberg was visiting the economics department at Princeton University. He is most grateful to this institution for its hospitality. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

[†]Go to <https://doi.org/10.1257/aer.20180688> to visit the article page for additional materials and author disclosure statements.

¹To avoid confusion, these terms are used to denote two different phenomena: (i) *one* person faces the same question *multiple times* and gives different answers; (ii) *different* subjects answer the same question only *once*, but subjects who appear similar, given the available data, make different choices. In this paper we focus on the first one.

²The pattern of stochastic choice was first reported in Tversky (1969). A large literature followed. Focusing on choices between risky gambles (as in our model), see Camerer (1989); Starmer and Sugden (1989); Hey and Orme (1994); Ballinger and Wilcox (1997); Hey (2001); Regenwetter, Dana, and Davis-Stober (2011); Regenwetter and Davis-Stober (2012); and Agranov and Ortoleva (2017).

ranking over the available alternatives and consistently choose the option that maximizes it.³

A large body of theoretical work has developed models to capture stochastic behavior. Most of these models can be ascribed to one of two classes. First, models of “random utility/preferences,” according to which subjects’ answers change because their preferences change stochastically.⁴ Second, models of “bounded rationality,” or “mistakes,” according to which subjects have stable and complete preferences, but may fail to always choose the best option and thus exhibit a stochastic pattern.⁵

While according to the interpretations above the stochasticity of choice happens involuntarily, a third possible interpretation is that stochastic choice is a *deliberate* decision of the agent: she may *choose* to report different answers from the same menu. The goals of this paper are to develop axiomatically a model in which stochastic choice follows this interpretation, and to identify whether, and how, such a model of deliberate randomization generates different behaviors and can be distinguished from other models of stochastic choice.

A small existing literature has suggested why subjects may wish to report stochastic answers. Machina (1985) notes that this is precisely what the agent may wish to do if her preferences over lotteries or acts are *convex* (i.e., quasiconcave in probability mixtures), which implies affinity toward randomization between equally good options. Crucially, convexity is a property shared by many existing models of decision making under risk, and it captures ambiguity aversion in the context of decision making under uncertainty. Convexity of preferences also has experimental support (Becker, DeGroot, and Marschak 1963, Sophor and Narramore 2000). Different reasons for stochastic choice to be deliberate were suggested by Marley (1997) and Swait and Marley (2013), who follow lines similar to Machina (1985); Dwenger, Kübler, and Weizsäcker (2018), who suggest it may be due to a desire to minimize regret; and Fudenberg, Iijima, and Strzalecki (2015), who connect it to uncertain taste shocks. In Section IV we discuss these papers in detail.

Recent experimental evidence supports the interpretation of stochastic choice as deliberate. Agranov and Ortoleva (2017) show how subjects give different answers also when the same question is asked three times in a row and subjects are aware of the repetition; they seem to explicitly choose to report different answers.⁶ Dwenger, Kübler, and Weizsäcker (2018) find that a large fraction of subjects choose lotteries between available allocations, indicating an explicit preference for randomization. They also show similar patterns using the data from a clearinghouse for university admissions in Germany, where students must submit multiple rankings of the universities they would like to attend. These are submitted at the same time, but

³This behavior is obviously consistent whenever there are multiple alternatives that maximize preferences and the individual uses different rules to break indifferences.

⁴Thurstone (1927), Luce (1959), Harsanyi (1973), Falmagne (1978), Cohen (1980), Barberá and Pattanaik (1986), McFadden and Richter (1991), Loomes and Sugden (1995), Clark (1996), McFadden (2006), Gul and Pesendorfer (2014), Ahn and Sarver (2013), Fudenberg and Strzalecki (2015).

⁵Models of this kind appear in economics, psychology, and neuroscience, including the well-known Drift Diffusion model: among many, Busemeyer and Townsend (1993); Harless and Camerer (1994); Hey and Orme (1994); Camerer and Ho (1994); Wu and Gonzalez (1996); Ratcliff and McKoon (2008); Gul, Natenzon, and Pesendorfer (2014); Manzini and Mariotti (2014); Woodford (2014); Fudenberg and Strzalecki (2015); Baldassi et al. (2018). For surveys, Ratcliff and Smith (2004), Bogacz et al. (2006), Johnson and Ratcliff (2013).

⁶In a survey conducted at the end of the experiment, most subjects report choosing different answers deliberately. These results hold true also in robustness tests with unusually high stakes.

only one of them (chosen randomly) matters. They find that a significant fraction of students report inconsistent rankings, even when there are no strategic reasons to do so. A survey among applicants supports the interpretation that these random allocations are chosen intentionally, and show that they are correlated with an explicit preference for randomization.⁷

We develop axiomatically a general model of stochastic choice over lotteries as the outcome of a deliberate desire to randomize. We aim to capture and formalize the intuition of Machina (1985) that such inclination may be a rational reaction if the underlying preferences over lotteries are (at least locally) convex. We consider a stochastic choice function over sets of lotteries over monetary outcomes, which assigns to any menu a probability distribution over its elements. We focus on lotteries not for technical reasons, but because we are interested in linking stochastic choice to features of preferences over lotteries in general, and violations of Expected Utility in particular. The focus on lotteries over *monetary* outcomes is inessential, as similar results are available for arbitrary prize spaces. We confine attention to monetary prizes for simplicity and since most of our leading examples use them (see Section I).

We begin our analysis with a representation theorem: we show that a rationality-type condition on stochastic choice, reminiscent of known acyclicity conditions, guarantees that it can be represented as if the agent had a preference relation over the final monetary lotteries, and chose the *optimal* mixing over the available options. In this model the stochasticity has a purely *instrumental* value for the agent: she does not value the randomization per se, but rather because it allows her to obtain the lottery over final outcomes she prefers. Implicit in our approach is that agents evaluate mixtures of lotteries by looking at the distribution over final outcomes they induce.

Next, we show that our model has some stark implications. Possibly the most well-known property of stochastic choice, widely used in the literature, is Regularity (also called Monotonicity): it posits that the probability of choosing p from a set cannot decrease if we remove elements from it. It is often seen as the stochastic equivalent of independence of irrelevant alternatives (IIA), and it is satisfied by many models in the literature, most prominently, models of Random Utility, albeit it is well known that it is often empirically violated. We show that our model of deliberate stochastic choice will *necessarily* lead to some violations of Regularity (unless the stochastic choice is degenerate, i.e., there is no stochasticity). Intuitively, our agent may choose from a set A two options that, together, allow her to “hedge.” But this holds only if they are *both* chosen: they are complementary to each other. If either option is removed from A , the possibility of hedging may disappear and the agent no longer has incentive to pick the remaining one, which in turn generates a violation of Regularity. The key observation is that the agent considers all the elements chosen *as a whole*, for the general hedging they provide *together*. By contrast, Regularity is based on the assumption that the appeal of each option is independent from the other options present in the menu or in

⁷Other papers have documented a desire to randomize. To our knowledge, the first was Rubinstein (2002), who confronts subjects with the task of making multiple choices in a row from a fixed set of alternatives. Each choice resembles the task of betting on the color of a ball drawn from an urn with a known distribution of colors. He documents a strong tendency to give different answers, even when it leads to dominated choices. Sandroni, Ludwig, and Kircher (2013) consider a version of the dictator game in which dictators can choose between 7.5 euros for themselves and 0 to the recipient, 5 to both, or a lottery between them. About one-third of the subjects choose to randomize. Similarly, Miao and Zhong (2018) find that substantial proportions of subjects in dictator games chose to randomize between allocations.

the choice. Thus, a violation of Regularity is an essential feature of the hedging behavior that we aim to capture, as we formally show.⁸

That our model is inconsistent with Regularity has direct implications on its relation with existing models, most prominently models of Random Utility. Since it is well known that the latter must satisfy Regularity, the only behavior that can be represented by both models is one that can always be described as if the agent had only one utility and randomization occurred only in the case of indifference, a degenerate random utility. In other words, the conceptual difference between the two models is reflected in a substantial behavioral difference, via the property of Regularity, which is very easily testable in experiments. Since also the models in Fudenberg, Iijima, and Strzalecki (2015) satisfy Regularity, the same relation holds between our model and theirs.

The remainder of the paper is organized as follows. Section I presents the general Deliberate Stochastic Choice model. Section II establishes that our model is incompatible with Regularity and studies its other behavioral implications. Section III discusses the relation with the existing literature. All proofs appear in the Appendices.

I. A General Model of Deliberate Stochastic Choice

A. Framework and Foundations

Let $[w, b] \subseteq \mathbb{R}$ be a nontrivial interval of monetary prizes and let Δ be the set of lotteries (Borel probability measures) over $[w, b]$, endowed with the topology of weak convergence. We use x, y, z and p, q, r for generic elements of $[w, b]$ and Δ , respectively. Denote by $\delta_x \in \Delta$ the degenerate lottery (Dirac measure at x) that gives the prize $x \in [w, b]$ with certainty. If p and q are such that p strictly first-order stochastically dominates q , we write $p \succ_{FOSD} q$.

Denote by \mathcal{A} the collection of all finite and nonempty subsets of Δ . For any $A \in \mathcal{A}$, $\text{co}(A)$ denotes the convex hull of A , that is, $\text{co}(A) = \{\sum_j \alpha_j p_j : p_j \in A \text{ and } \alpha_j \in [0, 1], \sum_j \alpha_j = 1\}$.

The primitive of our analysis is a *stochastic choice function* ρ over \mathcal{A} , i.e., a map ρ that associates to each $A \in \mathcal{A}$ a probability measure $\rho(A)$ over A . For any stochastic choice function ρ , $A \in \mathcal{A}$, and $p \in A$, $\text{supp}_\rho(A)$ denotes the support of $\rho(A)$, and we write $\rho(A)(p)$ to denote the probability ρ assigns to p in menu A .

As a final bit of notation, since $\rho(A)$ is a probability distribution over lotteries, thus a compound lottery, we can compute the induced lottery over final monetary outcomes. Denote it by $\overline{\rho(A)} \in \Delta$, that is

$$\overline{\rho(A)} = \sum_{q \in A} \rho(A)(q)q.$$

By construction, the convex hull of a set A , $\text{co}(A)$, will also correspond to the set of all monetary lotteries that can be obtained by choosing a specific ρ and computing the distribution it induces over final prizes.⁹

⁸In Section II we explore other implications of the model, beyond Regularity. For example, we show that it necessarily implies a version of stochastic intransitivity.

⁹That is, by construction $\text{co}(A) = \{p \in \Delta : p = \overline{\rho(A)} \text{ for some stochastic choice function } \rho\}$.

We can now discuss our first axiom. Our goal is to capture behaviorally an agent who is *deliberately* choosing her stochastic choice function following an underlying preference relation over lotteries. When asked to choose from a set A , she considers all lotteries that can be obtained from A by randomizing: using our notation above, she considers the whole $\text{co}(A)$, and the lottery $\rho(A)$ can be seen as her “choice.”

Our axiom is a rationality-type postulate for this case. Consider two sets A_1 and A_2 , and suppose that $\overline{\rho(A_2)} \in \text{co}(A_1)$. This means that the lottery chosen from A_2 could be obtained also from A_1 . Standard rationality posits that the “choice” from A_1 , $\overline{\rho(A_1)}$, must then be at least as good as anything that can be obtained from A_2 . Since we do not observe preferences, we cannot impose this; but at the very least we can say that there cannot be anything in A_2 that strictly first-order stochastically dominates $\overline{\rho(A_1)}$. This is the content of our axiom, extended to any sequence of length k of sets.

AXIOM 1 (Rational Mixing): For each $k \in \mathbb{N} \setminus \{1\}$ and $A_1, \dots, A_k \in \mathcal{A}$, if

$$\overline{\rho(A_2)} \in \text{co}(A_1), \dots, \overline{\rho(A_k)} \in \text{co}(A_{k-1}),$$

then $q \in \text{co}(A_k)$ implies $q \not\succ_{FOSD} \overline{\rho(A_1)}$.

Rational Mixing is related to conditions of rationality and acyclicity typical in the literature on revealed preference with limited observations, along the lines of Afriat’s condition and the Strong Axiom of Revealed Preferences (see, e.g., Chambers and Echenique 2016). Intuitively, the ability to randomize allows the agent to choose any option in the convex hull of all sets; thus, it is as if we could only see the choices from convex sets, and posit a rationality condition for this case.

Note that Rational Mixing implicitly (i) includes a form of coherence with strict first-order stochastic dominance, and (ii) assumes that the agent cares only about the *induced distribution over final outcomes*, rather than the *procedure* in which it is obtained. That is, for the agent the stochasticity is *instrumental* to obtain a better distribution over final outcomes, rather than being valuable per se. This implies a form of reduction of compound lotteries, which we will maintain throughout.

B. Deliberate Stochastic Choice Model

DEFINITION 1: A stochastic choice function ρ admits a Deliberate Stochastic Choice representation if there exists a complete preorder (a transitive and reflexive binary relation) \succsim over Δ such that

(i) For every $A \in \mathcal{A}$,

$$\overline{\rho(A)} \succsim q \text{ for every } q \in \text{co}(A);$$

(ii) For every pair $p, q \in \Delta$, $p \succ_{FOSD} q$ implies $p \succ q$.

THEOREM 1: *A stochastic choice function ρ satisfies Rational Mixing if and only if it admits a Deliberate Stochastic Choice representation.*

A Deliberate Stochastic Choice model captures a decision maker who has a preference relation \succsim over lotteries and chooses deliberately the randomization that generates the optimal mixture among existing options. This procedure is most prominent in regions where \succsim is strictly convex and, in particular, if there exist some $p, q \in \Delta$ and $\alpha \in (0, 1)$ such that $\alpha p + (1 - \alpha)q \succ p, q$. When faced with the choice from $\{p, q\}$, she would strictly prefer to randomly choose rather than to pick either of the two options. The stochasticity is thus an expression of the agent's preferences.

Prominent examples of preferences with this property are the Rank Dependent Expected Utility (RDU) model with the common inverse S-shaped probability weighting function (overweight small probabilities and underweight large probabilities, where the former implies convexity),¹⁰ and strictly convex versions of Quadratic Utility (Chew, Epstein, and Segal 1991).¹¹ Another example is when \succsim follows the Cautious Expected Utility model of Cerreia-Vioglio, Dillenberger, and Ortoleva (2015); these preferences are convex (albeit not strictly convex everywhere), and may exhibit stochastic choice. In Cerreia-Vioglio et al. (2018) we show that the corresponding Cautious Stochastic Choice model tightly links stochasticity of choice and the Certainty Bias, as captured by both Allais paradoxes (Common-Ratio and Common-Consequence effects), one of the most prominently observed violations of Expected Utility. Both stem from the presence of multiple utilities and the use of caution.¹² The general class of convex preferences over lotteries was studied in Cerreia-Vioglio (2009).

On the other hand, the model does not restrict preferences to have any region of strict convexity. It also permits indifference to randomization, e.g., when \succsim follows Expected Utility, or, more generally, when it satisfies Betweenness;¹³ or even aversion to randomization, e.g., if \succsim is RDU with convex probability weighting function: in these cases the agent has no desire to mix and we should observe no stochasticity (except for indifferences).

Note also that the model puts no restriction on the way the agent resolves indifferences: when multiple alternatives maximize the preference relation, any could be chosen. Although it is a typical approach not to rule how indifferences

¹⁰ If we order the prizes in the support of a finite lottery p , with $x_1 < x_2 < \dots < x_n$, then the functional form for RDU is $V(p) = u(x_n)f(p(x_n)) + \sum_{i=1}^{n-1} u(x_i) \left[f(\sum_{j=i}^n p(x_j)) - f(\sum_{j=i+1}^n p(x_j)) \right]$, where $f: [0, 1] \rightarrow [0, 1]$ is strictly increasing and onto and $u: [w, b] \rightarrow \mathbb{R}$ is increasing. Regions where $f(l) > l$ imply *affinity* to randomization.

¹¹ For example, such preferences over lotteries are convex if they are represented by the product of two positive expected utility functionals, that is, $V(p) = E_p(u) \times E_p(v)$, where u and v are positive, continuous, and strictly increasing.

¹² Preferences admit a Cautious Expected Utility representation if there exists a set \mathcal{W} of strictly increasing and continuous (Bernoulli) utility functions over monetary outcomes, such that the value of any lottery p is given by $V(p) = \inf_{v \in \mathcal{W}} v^{-1}(E_p(v))$. That is, the agent has a set of utility functions over outcomes, and evaluates each lottery p by first computing the certainty equivalent of p with respect to each possible function in the set, and then picking the smallest one. Cerreia-Vioglio et al. (2018) show that as long as there are finitely many utilities, agents have a strict desire to randomize if and only if they violate Expected Utility in line with the Certainty Bias.

¹³ Betweenness requires that $p \sim q \Rightarrow \alpha p + (1 - \alpha)q \sim q$ for all $p, q \in \Delta$, $\alpha \in (0, 1)$. It is satisfied by any Expected Utility maximizer. Utility theories with the Betweenness property were studied by Dekel (1986) and Chew (1989). See also an earlier axiomatization developed by Fishburn (1983).

are resolved, following it may lead to discontinuities.¹⁴ This implies that, for each preference relation \succsim , there is potentially more than one stochastic choice function derived from it, depending on how indifferences are resolved. In general, we say that a stochastic choice function $\hat{\rho}$ is consistent with a preference relation \succsim if and only if for each $A \in \mathcal{A}$, $\hat{\rho}(A) \succsim q$ for all $q \in \text{co}(A)$.

As we pointed out in the introduction, Theorem 1 does not require the lotteries to be over monetary outcomes; they could be over an arbitrary set of prizes, as long as there is some natural dominance relation (a partial order) \triangleright over the space of lotteries that one can use to replace $>_{FOSD}$ in the statement of the Rational Mixing axiom. This can be a generalization of the concept of first-order stochastic dominance, or any other partial order that satisfies the additional property that $p \triangleright q$ implies that $p \triangleright \alpha p + (1 - \alpha)q$ for every $\alpha \in [0, 1)$ and for every pair p and q .

Remark 1: Our framework implicitly assumes that we observe the stochastic choice function ρ for all sets in \mathcal{A} . This is very demanding, and a natural question is what tests are required if we observe only limited data. In fact, the Rational Mixing axiom is necessary and sufficient in *any* dataset that includes all doubletons $\{p, q\}$ such that $p >_{FOSD} q$. Consider any $\mathcal{B} \subseteq \mathcal{A}$ such that $\{p, q\} \in \mathcal{B}$ whenever $p >_{FOSD} q$ and denote by $\rho_{\mathcal{B}}$ the restriction of ρ on \mathcal{B} . Then, $\rho_{\mathcal{B}}$ satisfies Rational Mixing if and only if $\rho_{\mathcal{B}}$ admits a Deliberate Stochastic Choice representation. (The proof follows exactly the same steps as the proof of Theorem 1.)

Remark 2: The preference relation \succsim in a Deliberate Stochastic Choice model need not admit a utility representation, e.g., if it is lexicographic.¹⁵ As usual, to guarantee the existence of a utility representation we need \succsim to be continuous.¹⁶ We call this case a *Continuous* Deliberate Stochastic Choice model. Proposition 3 in Appendix A gives an axiomatic characterization of it, obtained by strengthening Rational Mixing with a continuity requirement. In words, consider the binary relation R , defined as pRq if $p = \overline{\rho(A)}$ for some A such that $q \in \text{co}(A)$. Rational Mixing simply posits that the transitive closure of R is consistent with $>_{FOSD}$. To obtain a continuous representation, we need to extend this requirement to the closed

¹⁴ While with choice correspondences the continuity of the underlying preference relation implies continuity of the choice correspondence (i.e., satisfies the closed graph property), here it is as if we observed also the outcome of how indifference is resolved (which may be stochastic). This will necessarily imply discontinuities of ρ , following standard arguments. An alternative, although significantly less appealing, approach would be to consider a stochastic choice *correspondence*, which could be fully continuous.

¹⁵ For example, consider the binary relation \succsim_L defined by $p \succsim_L q$ if either $E_p(x^3) > E_q(x^3)$ or $E_p(x^3) = E_q(x^3)$ and $E_p(x) \geq E_q(x)$. This binary relation is complete, transitive, and satisfies Independence, but fails Continuity. Note that for any (finite) menu of lotteries $A \in \mathcal{A}$, $\arg \max(\succsim_L, A) \neq \emptyset$. Define ρ to be such that $\rho(A)(q) = 1/|\arg \max(\succsim_L, A)|$ for all $q \in \arg \max(\succsim_L, A)$. Since \succsim_L satisfies Independence, $\overline{\rho(A)} \succsim_L q$ for all $q \in \text{co}(A)$, and since \succsim_L strictly preserves first-order stochastic dominance, all the requirements of Theorem 1 are met. Yet clearly \succsim_L does not admit a utility representation.

¹⁶ We say that a relation \succsim is continuous if and only if for each $q \in \Delta$ the sets $\{p \in \Delta : p \succsim q\}$ and $\{p \in \Delta : q \succsim p\}$ are closed with respect to the topology of weak convergence.

transitive closure of R (i.e., the minimal continuous and transitive relation that contains R).

II. Regularity and Deliberate Randomization

In this section we study the relation of the Deliberate Stochastic Choice model with the Regularity property we mentioned in the introduction. As we explained there, Regularity, also called Monotonicity, is a well known and extensively used property in the stochastic choice literature. Formally, we have the following definition.

AXIOM 2 (Regularity): *For each $A, B \in \mathcal{A}$ and $p \in A$, if $A \subseteq B$, then $\rho(B)(p) \leq \rho(A)(p)$.*

Intuitively, Regularity states that if we remove some elements from a set, the probability of choosing the remaining elements cannot decrease. Conceptually, it is related to notions of independence of irrelevant alternatives applied to a stochastic setting: the removal of any element, chosen or unchosen, cannot “hurt” the chances of choosing any of the remaining ones. In other words, the attractiveness of an option should not depend on the availability of other ones. Crucially, the property of Regularity is one of the characterizing features of models of Random Utility. We should note that despite its widespread use and normative appeal, substantial experimental evidence has been collected that shows how Regularity is violated: see Rieskamp, Busemeyer, and Mellers (2006) for a survey.

To analyze this property in our model, it will be useful to formally define when is the individual exhibiting stochastic choice. Recall that in our model this may happen either when there is a genuine desire to randomize, or in the case of indifferences. We say that an agent exhibits a non-degenerate stochastic choice function if stochasticity is present beyond indifference: if we can find some p and q such that the agent randomizes between them and also does so when either is made a “little bit worse” by mixing with δ_w (the worst possible outcome).

DEFINITION 2: *A stochastic choice function ρ is non-degenerate if there exist $p, q \in \Delta$ with $|\text{supp}_\rho(\{p, q\})| \neq 1$ and $\lambda \in (0, 1)$ such that*

$$|\text{supp}_\rho(\{\lambda p + (1 - \lambda)\delta_w, q\})| \neq 1 \quad \text{and} \quad |\text{supp}_\rho(\{p, \lambda q + (1 - \lambda)\delta_w\})| \neq 1.$$

Endowed with these definitions, we have the following theorem.

THEOREM 2: *Let ρ be a stochastic choice function that admits a Continuous Deliberate Stochastic Choice representation \succsim . The following statements are equivalent:*

- (i) ρ is non-degenerate;
- (ii) ρ and any other $\hat{\rho}$ consistent with \succsim violates Regularity;

- (iii) \succsim has a point of strict convexity, that is, there exist $p, q \in \Delta$ and $\lambda \in (0, 1)$ such that

$$\lambda p + (1 - \lambda)q \succ p, q.$$

Theorem 2 shows that the Deliberate Stochastic Choice model *must* lead to violations of Regularity, unless the Stochastic Choice is degenerate (no stochasticity, or purely to break indifferences). That is, whenever there is even a single instance in which the individual makes a non-degenerate stochastic choice, then there necessarily is some other instance where her choices violate Regularity. In fact, the result is stronger: violations of Regularity and stochasticity imply one another; and both occur if and only if the underlying preferences have points of strict convexity. Without them, under the model there should never be any stochasticity (except possibly in the case of indifference), as the agent does not have a desire to randomize; but with them, we also have violations of Regularity.

One important implication of the result above is to distinguish our model from models that satisfy Regularity, which include most popular models such as Random Utility and Luce's (1958) model. Thus, the desire to randomize not only is conceptually different, but also leads to a behavior that violates the core property of many models in the literature.

To gain an intuition, consider some p, q where preferences are strictly convex, i.e., there exists $\lambda \in (0, 1)$ such that $\lambda p + (1 - \lambda)q \succ p, q$, as in item (iii) of the theorem. For simplicity, suppose that $r = \hat{\lambda}p + (1 - \hat{\lambda})q$ is the unique \succsim -optimal mix between p and q . Let r_ε be a lottery within distance $\varepsilon > 0$ of r but strictly first-order stochastically dominated by it. First observe that p will be chosen with probability $\hat{\lambda}$ from $\{p, q, r_\varepsilon\}$: in the face of both p and q , the presence of r_ε is of no value to the agent. But suppose q is removed: then, as long as ε is small enough (r_ε is close to r), p will be chosen with very small probability from $\{p, r_\varepsilon\}$: the value of p decreases significantly without q , and the agent now puts most weight on r_ε . This pair of choices violates Regularity. The idea behind this construction is that p and q are *complementary* to each other. But if q is removed, the agent can no longer hedge between them; r_ε , which was inferior to their mixture, then becomes an attractive alternative in the face of p . Overall, the crucial aspect is that the ability of choosing *both* p and q at the same time renders them appealing, while they would not be appealing in isolation. This is a fundamental aspect of when hedging is advantageous: the *whole* set of chosen elements is relevant for the agent, for the hedging opportunities it provides. Such complementarity between alternatives violates standard independence of irrelevant alternatives arguments, according to which chosen elements should be appealing in isolation, which is also reflected in the Regularity axiom. For that reason, violations of Regularity are a "structural" feature of our model.

The result of Theorem 2 has an additional conceptual implication. From the point of view of our model, violations of Regularity should not be seen as mistakes or as forms of bounded rationality. On the contrary, our model entails a strong form of "rationality:" individuals are endowed with well-defined, stable, and monotone preferences over lotteries, and select the combination of options from all possible ones to

maximize them, reducing compound lotteries. With the exception of possibly violating Expected Utility, our agents are thus as close as possible to the standard rational economic decision maker. Our results show that despite all this, not only they may exhibit stochastic choice, but when they do they also, as a manifestation of their preferences, violate the property of Regularity, often described as the counterpart of “rationality” for stochastic choice. (This interpretation may thus be put in question.)

We conclude this section by discussing which violations of Regularity and of other properties are compatible with, or even implied by, our model. On the one hand, our model is compatible with many violations, including documented ones, as well as others that are yet to be explored experimentally, beyond those of the form identified in the constructive proof of Theorem 2. For example, consider some $p, q, r \in \Delta$ such that some mixture of p and q is ranked above r , but r is better than p and any mixture between p and r . (Such examples are easy to construct when preferences are strictly convex. In the Cautious Stochastic Choice model discussed in footnote 12, these will be present whenever the set \mathcal{W} is finite, with r being a degenerate lottery.) In this case, from the set $\{r, p\}$, p would never be chosen, but it may be chosen with strictly positive probability from $\{p, q, r\}$. Our model is also consistent with situations where r lies above the line segment connecting p and q , and violations of Regularity take the form $\rho(\{p, q, r\})(r) > \rho(\{p, r\})(r)$. Violations of this kind have been widely documented and referred to as versions of the compromise effect or of the attraction effect without dominance (see Simonson 1989; Tversky and Simonson 1993; Rieskamp, Busemeyer, and Mellers 2006; Soltani, De Martino, and Camerer 2012; see also the discussion in Natenzon 2019).

On the other hand, the Deliberate Stochastic Choice model is not compatible with violations of Regularity due to the addition of a (strictly) first-order stochastically dominated option. For example, suppose that p is chosen more frequently in $\{p, q, r\}$ than in $\{p, q\}$, where r is dominated by p but not by q . These violations of Regularity have been documented empirically, and referred to as either the asymmetric dominance effect or the attraction effect with a dominated option (see Ok, Ortoleva, and Riella 2015; Soltani, De Martino, and Camerer 2012). They are incompatible with our model because no dominated option can ever be part of an optimal mixture, thus its addition cannot modify the optimal combination: violations of this type may occur only because of indifferences. For similar reasons, our model satisfies a weaker version of Regularity that posits that choice probabilities do not decrease if we remove elements that are *never* chosen, again, except for indifferences.

Two other known properties of stochastic choice are Weak and Strong Stochastic Transitivity. Take any p, q, r with $\rho(\{p, q\})(p) \geq 0.5$ and $\rho(\{q, r\})(q) \geq 0.5$. Weak Stochastic Transitivity is satisfied if $\rho(\{p, r\})(p) \geq 0.5$; Strong Stochastic Transitivity requires $\rho(\{p, r\})(p) \geq \max\{\rho(\{p, q\})(p), \rho(\{q, r\})(q)\}$. Both properties are known to be independent of Regularity, and substantial evidence has shown that they are often violated, especially the stronger version. It is easy to construct examples of similar violations in our model (e.g., the evidence discussed in Rieskamp, Busemeyer, and Mellers 2006, p. 636); in general, our model is consistent with violations of both forms. It is worth noting that Machina (1985) already alluded to the idea that with strictly convex preferences, one may expect either version of Stochastic Transitivity to be violated, and thus, unlike violations of transitivity

of the underlying preference relation \succsim , such violations should be perceived as neither normatively disturbing nor descriptively rare. Our next result establishes that not only these violations are compatible with our model, but that any non-degenerate stochastic choice function that admits a Continuous Deliberate Stochastic Choice representation *necessarily* violates Strong Stochastic Transitivity.¹⁷

PROPOSITION 1: *Let ρ be a stochastic choice function that admits a Continuous Deliberate Stochastic Choice representation \succsim . The following statements are equivalent:*

- (i) ρ is non-degenerate;
- (ii) ρ and any other $\hat{\rho}$ consistent with \succsim violates Strong Stochastic Transitivity.

In other words, violations of Strong Stochastic Transitivity, just like those of Regularity, are core features of our model whenever the stochastic choice is non-degenerate. Indeed, an immediate corollary of the last two results is that under the continuous version of our model, detecting a violation of either of the two properties implies that a violation of the other can be found as well.

Finally, our model makes some new empirical predictions on stochastic choice via the Rational Mixing Axiom. For example, it is easy to see that the addition of an option within the convex hull of a set should not change the final distribution of prizes that the agent receives: because it does not provide any new hedging opportunity, it should not affect what final distribution the agent is able to achieve, thus the stochastic choice may change only due to indifferences.

III. Relation with Models in the Literature

Random Utility.—As we discussed, Theorem 2 can be used to easily compare the Deliberate Stochastic Choice model with models of Random Utility. Formally, we say that a stochastic choice function ρ admits a Random Utility representation if there exists a probability measure over utilities such that for each alternative s in a choice problem A , the probability of choosing s from A , $\rho(A)(s)$, equals the probability of drawing a utility function u such that s maximizes u in A .¹⁸

It is well known that a stochastic choice function that admits a Random Utility representation must satisfy Regularity. This is intuitive: if an option is the best according to one utility, its choice cannot be made less likely by removing alternatives. (In models of Random Utility, there is no complementarity between the chosen

¹⁷For some special cases, also violations of Weak Stochastic Transitivity are implied. For example, reconsider the Cautious Stochastic Choice model with a finite set \mathcal{W} . Violations of Weak Stochastic Transitivity can always be constructed whenever there are $p, q \in \Delta$ such that both $p \sim q$ and $\arg \max_{\mu \in [0,1]} V(\mu p + (1 - \mu)q)$ is unique.

¹⁸Stochastic choice functions over a finite space of alternatives that admit a Random Utility representation were axiomatized by Falmagne (1978) (see also Barberá and Pattanaik 1986). An issue arises when the utility functions allow for indifferences; assumptions are needed on how they are resolved. Two approaches have been suggested. First, to impose that the set of utility functions such that the maximum is not unique has measure zero for every choice problem, as is the case, for example, for logit or probit. Second, to impose a tie-breaking rule, that may vary with each utility, but that selects one of the maximizers coherently and independently of the choice problem (e.g., satisfying Sen's α). In what follows we assume that one of these two approaches is adopted.

elements.) But then, following Theorem 2 we have a sharp distinction between our model and models of Random Utility: the only behavior that can be represented by both models is one compatible with a *degenerate* Random Utility model, i.e., with only one utility possible, in which the agent exhibits stochastic behavior only when she is indifferent. Another immediate implication is that observing a violation of Regularity, an easily testable condition, implies that the agents' behavior cannot be represented by Random Utility, while it may be represented by the Deliberate Stochastic Choice.

Random Expected Utility.—Gul and Pesendorfer (2006) axiomatize the Random Expected Utility model, a version of Random Utility where all the utility functions involved are of the Expected Utility type. One of the conditions that characterize this model is Linearity.¹⁹

AXIOM 3 (Linearity): For each $A \in \mathcal{A}$, $p \in A$, $q \in \Delta$, and $\lambda \in (0, 1)$,

$$\rho(A)(p) = \rho(\lambda A + (1 - \lambda)q)(\lambda p + (1 - \lambda)q).$$

We now show that if ρ admits a Continuous Deliberate Stochastic Representation and in addition satisfies Linearity, then ρ is a *degenerate* Random Expected Utility model, i.e., again a model with only one utility. Formally we have the following result.

PROPOSITION 2: Let ρ be a stochastic choice function that admits a Continuous Deliberate Stochastic Choice representation and satisfies Linearity. Then, there exists a continuous function $u: [w, b] \rightarrow \mathbb{R}$ such that, for any choice problem A ,

$$\text{supp}_\rho(A) \subseteq \{p \in A : E_p(u) \geq E_q(u) \ \forall q \in A\}.$$

Other Models of Deliberate Randomization.—A small existing literature has suggested models of stochastic choice as deliberate randomization. As we have discussed, our own model formalizes and extends the intuition of Machina (1985) (see also Marley 1997 and Swait and Marley 2013) in a fully axiomatic setup.²⁰ Dwenger, Kübler, and Weizsäcker (2018) propose a model in which agents choose to randomize following a desire to minimize regret. Their key assumption is that the regret after making the wrong choice is smaller if the choice is stochastic rather than deterministic. Allen and Rehbeck (2019) develop a model of deliberate randomization that includes as special cases known models of stochastic choice.

¹⁹ Given a set $A \subseteq \Delta$ and a lottery $q \in \Delta$, the set $\lambda A + (1 - \lambda)q$ denotes the set of all lotteries $p \in \Delta$ such that $p = \lambda r + (1 - \lambda)q$ for some $r \in A$.

²⁰ Machina (1985) suggests the following condition: if $A, A' \in \mathcal{A}$ are such that $\text{co}(A) \subseteq \text{co}(A')$ and $\overline{\rho(A')} \in \text{co}(A)$, then $\overline{\rho(A')} = \overline{\rho(A)}$. (This condition is related to Sen's α axiom.) While naturally related to our Rational Mixing axiom, this condition is not sufficient to characterize our model. (Unless preferences are strictly convex, it is also not necessary, because of indifferences: for example, A and A' may differ only for the inclusion of some strictly dominated option that is never chosen in either case, but the stochastic choice may not coincide as indifference may be resolved differently.)

Fudenberg, Iijima, and Strzalecki (2015) provide conditions under which stochastic choice corresponds to the maximization of Expected Utility and a perturbation function that depends only on the choice probabilities. Formally, they axiomatize a stochastic choice function ρ such that, for each choice problem A ,

$$(1) \quad \rho(A) = \arg \max_{p \in \Delta(A)} \sum_{x \in A} [p(x)u(x) - c(p(x))],$$

where $\Delta(A)$ is the set of probability measures on A , u is a von Neumann-Morgenstern utility function and $c: [0, 1] \rightarrow \mathbb{R} \cup \{\infty\}$ is strictly convex and C^1 in $(0, 1)$. They call this representation *Weak Additive Perturbed Utility* (Weak APU).²¹ Because of the strictly convex perturbation function c , this functional form gives the agent an intrinsic incentive to randomize. However, there are two important differences with our model.

A first difference is that we study a domain of menus of lotteries while Fudenberg, Iijima, and Strzalecki (2015) study menus of final outcomes. This is not a mere technical difference, as our goal is to study, in the spirit of Machina (1985), the link of stochastic choice with non-Expected Utility behavior; and (deterministic) non-Expected Utility preferences over lotteries must necessarily be present for a comparison to be possible.

A second, crucial difference between the models is that even though the model in Fudenberg, Iijima, and Strzalecki (2015) rewards probabilistic choices and this sometimes gives the individual an incentive to randomize, their model *does* satisfy Regularity (Fudenberg, Iijima, and Strzalecki 2015, p. 2386). This is a crucial conceptual difference, as it implies that their model does not include one of the main driving forces of ours, as we discussed above. It also implies that the formal relation between their model and ours is the same as with Random Utility: the only behavior compatible with both is one of an agent that exhibits stochastic choice only when indifferent.²²

The results above are summarized in Figure 1.²³

Other Related Literature.—Our paper is related to various other strands of the literature. First, it is related to models that connect violations of rationality (in the form of the Weak Axiom of Revealed Preferences or Regularity) to various forms of bounded rationality and costly information processing: among many, see for deterministic choice Manzini and Mariotti (2007), Masatlioglu, Nakajima, and Ozbay (2012), Ok, Ortoleva, and Riella (2015); for stochastic choice, Manzini and Mariotti (2014), Caplin and Dean (2015), Matějka and McKay (2015), Natenzon (2019), and

²¹Their paper also characterizes the case in which the function c satisfies the additional requirement that $\lim_{q \rightarrow 0} c'(q) = -\infty$, which they call an Additive Perturbed Utility representation.

²²An alternative way to apply their paper to the case of lotteries is, instead of using their representation theorem directly, to use a continuous version of their functional form, $\sum [p(x)u(x) - c(p(x))]$, as a representation for the preferences in our Theorem 1. This would lead to a model that is a hybrid of the two formulations.

²³For preferences that satisfy Linearity but not Regularity, suppose that facing a menu A , the agent considers two functions u and v , finds the sets $\arg \max_{p \in \text{co}(A)} E_p(u)$ and $\arg \max_{p \in \text{co}(A)} E_p(v)$, and splits the probability of choice evenly among all maximizers. This behavior satisfies Linearity, but violates Regularity whenever one adds to a set an option that is the unique maximizer for one of the utility functions for which, before the addition, there was more than one maximizer.




FIGURE 1. RELATION WITH OTHER MODELS

many references therein. As we have discussed, as opposed to these papers, in our model violations of Regularity follow even if subjects are fully rational and are fully informed about the options, as long as their underlying preferences are (at least locally) convex.

Stoye (2015) studies choice environments in which agents can randomize at will (thus restricting observability to convex sets). Considering as a primitive the choice correspondence of the agent in an Anscombe-Aumann setup,²⁴ he characterizes various models of choice under uncertainty that include a desire to randomize. Unlike Stoye, we take as a primitive the agent's stochastic choice function, instead of the choice correspondence; this not only suggests different interpretations, but also entails substantial technical differences. In addition, we study a setup with risk, and not uncertainty, and characterize the most general model of deliberate randomization given a complete preference relation over monetary lotteries.

Finally, as we have mentioned, our general representation theorem (Theorem 1) is related to the literature on revealed preference on finite datasets. By randomizing over a set of alternatives, the agent can obtain any point of its convex hull. It is as if we could only see individuals' choices from convex sets, restricting our ability to observe the entire preferences. Our problem is then related to the issue of eliciting preferences with limited datasets, originated by Afriat (1967), and for our first theorem we employ techniques from this literature. Our results are particularly related to Chambers and Echenique (2016) and Nishimura, Ok, and Quah (2017).

²⁴The paper considers also a setup with pure risk, but in that case the analysis is mostly focused on characterizing the case of Expected Utility, where there is no desire to randomize.

APPENDIX A. CONTINUOUS DELIBERATE STOCHASTIC CHOICE

In this section we extend the result of Theorem 1 to the case in which the underlying preference relation admits a representation by a continuous utility function. For this we need to strengthen the consistency condition in the Rational Mixing axiom to apply not only to the transitive closure of the relation R defined in Remark 2, but to its closed transitive closure. Formally, define the binary relation R on Δ as

$$pRq \text{ if and only if } \exists A \in \mathcal{A} \text{ such that } p = \overline{\rho(A)} \text{ and } q \in \text{co}(A).$$

Intuitively, pRq if it ever happens that p is chosen, either directly ($\{p\} = \text{supp}_\rho(A)$) or as the outcome of randomization ($p = \overline{\rho(A)}$), from a set A where q could have also been chosen ($q \in \text{co}(A)$). Denote by $\overline{\text{tran}(R)}$ the minimal continuous and transitive binary relation on Δ such that $R \subseteq \overline{\text{tran}(R)}$. We call $\overline{\text{tran}(R)}$ the closed transitive closure of R . We can now write the following postulate.

AXIOM 4 (Continuous Rational Mixing): *If $p, q \in \Delta$ are such that $p \overline{\text{tran}(R)} q$, then it cannot be the case that $q >_{\text{FOSD}} p$.*

We have the following result.

PROPOSITION 3: *A stochastic choice function ρ satisfies Continuous Rational Mixing if and only if it admits a Deliberate Stochastic Choice representation \succsim that can be represented by a continuous utility function.*

PROOF OF PROPOSITION 3:

Suppose first that ρ admits a Deliberate Stochastic Choice representation \succsim that can be represented by a continuous utility function. This implies that \succsim is a continuous preorder. Since, by the representation of ρ , $R \subseteq \succsim$, this implies that $\overline{\text{tran}(R)} \subseteq \succsim$. But then, if $p, q \in \Delta$ are such that $p \overline{\text{tran}(R)} q$, we also have that $p \succsim q$, which implies that it is not true that $q >_{\text{FOSD}} p$. That is, ρ satisfies Continuous Rational Mixing.

Conversely, suppose ρ is a stochastic choice function that satisfies Continuous Rational Mixing. Pick any pair of lotteries p and q in Δ with $q >_{\text{FOSD}} p$. This implies that $q >_{\text{FOSD}} \alpha q + (1 - \alpha)p$ for every $\alpha \in [0, 1]$. By Continuous Rational Mixing, we must have that $\overline{\rho(\{p, q\})} = q$, which implies that qRp . Moreover, again by Continuous Rational Mixing, we cannot have that $p \overline{\text{tran}(R)} q$. This shows that $\overline{\text{tran}(R)}$ is an extension of the first-order stochastic dominance relation. By Levin's Theorem (see Bridges and Mehta 1995, Lemma 8.3.4), there exists a continuous function $u: \Delta \rightarrow \mathbb{R}$ such that $p \overline{\text{tran}(R)} q$ implies $u(p) \geq u(q)$, with strict inequality whenever it is not true that $q \overline{\text{tran}(R)} p$. Now we can proceed as in the proof of Theorem 1, using the preference relation the function u induces, to conclude the proof. ■

APPENDIX B. PROOF OF THE RESULTS IN THE MAIN TEXT

PROOF OF THEOREM 1:

It is clear that if ρ admits a Deliberate Stochastic Choice representation, then ρ satisfies Rational Mixing. Suppose, thus, that ρ satisfies Rational Mixing and define the binary relation R the same way it is defined in the main text (see Remark 2). Pick any pair of lotteries p and q such that $p \succ_{FOSD} q$. This implies that $p \succ_{FOSD} (\alpha p + (1 - \alpha)q)$ for every $\alpha \in [0, 1)$. Define $A_1 = \{p, q\}$ and $A_2 = \{p\}$. Notice that Rational Mixing implies that we must have $\overline{\rho(A_1)} = p$. Consequently, we have pRq . Moreover, if we have $k \in \mathbb{N}$ and A_1, \dots, A_k such that $\overline{\rho(A_1)} = q$ and $\overline{\rho(A_i)} \in \text{co}(A_{i-1})$ for $i = 2, \dots, k$, Rational Mixing implies that $p \notin \text{co}(A_k)$. This shows that we cannot have $q \text{ tran}(R) p$. We conclude that $\text{tran}(R)$ is an extension of the first-order stochastic dominance relation. Now pick any complete extension \succsim of $\text{tran}(R)$. By what we have just seen, \succsim is also an extension of the first-order stochastic dominance relation. Moreover, by definition, we have that $\overline{\rho(A)}Rq$ for every $q \in \text{co}(A)$, for every $A \in \mathcal{A}$. Consequently, we have $\overline{\rho(A)} \succsim q$ for every $q \in \text{co}(A)$, for every $A \in \mathcal{A}$. This proves Theorem 1. ■

PROOF OF THEOREM 2:

(i) implies (iii). Assume that ρ is non-degenerate. Then, there exist $p, q \in \Delta$ and $\lambda \in (0, 1)$ with $p \succsim q$ and $|\text{supp}_\rho(\{p, \lambda q + (1 - \lambda)\delta_w\})| = 2$. Since \succsim preserves strict first-order stochastic dominance and is a Deliberate Stochastic Choice representation of ρ , this implies that $p, q \succ \delta_w$ and there exists $\gamma \in (0, 1)$ with

$$\gamma p + (1 - \gamma)q \succ \gamma p + (1 - \gamma)(\lambda q + (1 - \lambda)\delta_w) \succsim p \succsim q.$$

That is, \succsim has a point of strict convexity.

(iii) implies (i). By assumption, there exist $p, q \in \Delta$ and $\gamma \in (0, 1)$ such that

$$(2) \quad \gamma p + (1 - \gamma)q \succ p, q.$$

Since \succsim is continuous, this implies that there exists $\lambda \in (0, 1)$ such that $\gamma(\lambda p + (1 - \lambda)\delta_w) + (1 - \gamma)q \succ \lambda p + (1 - \lambda)\delta_w, q$ and $\gamma p + (1 - \gamma)(\lambda q + (1 - \lambda)\delta_w) \succ p, \lambda q + (1 - \lambda)\delta_w$. Since \succsim is a Deliberate Stochastic Choice representation of ρ , this can happen only if $|\text{supp}_\rho(\{p, q\})| = |\text{supp}_\rho(\{\lambda p + (1 - \lambda)\delta_w, q\})| = |\text{supp}_\rho(\{p, \lambda q + (1 - \lambda)\delta_w\})| = 2$.

(iii) implies (ii). Assume that there exist $p, q \in \Delta$ and $\lambda \in (0, 1)$ such that $\lambda p + (1 - \lambda)q \succ p, q$. Since \succsim strictly preserves first-order stochastic dominance, we must have $p \neq \delta_w$ and $q \neq \delta_w$. By continuity of \succsim , there exist maximal and minimal α^M and α_m in $[0, 1]$ such that $\alpha^M p + (1 - \alpha^M)q \sim \alpha_m p + (1 - \alpha_m)q \succsim \alpha p + (1 - \alpha)q$ for every $\alpha \in [0, 1]$. Note that we must have $0 < \alpha_m \leq \alpha^M < 1$. By construction, $\alpha^M p + (1 - \alpha^M)q \succ \lambda p + (1 - \lambda)(\alpha^M p + (1 - \alpha^M)q)$ for every

$\lambda \in [\alpha_m, 1]$. Continuity of \succsim and the fact that \succsim strictly preserves first-order stochastic dominance now imply that there exists $\varepsilon \in (0, 1)$ such that

$$\begin{aligned} & \varepsilon \delta_w + (1 - \varepsilon)(\alpha^M p + (1 - \alpha^M)q) \\ & \succ \lambda p + (1 - \lambda)(\alpha^M p + (1 - \alpha^M)q) \\ & \succsim \lambda p + (1 - \lambda)[\varepsilon \delta_w + (1 - \varepsilon)(\alpha^M p + (1 - \alpha^M)q)] \end{aligned}$$

for all $\lambda \in [\alpha_m, 1]$. Let $r = \varepsilon \delta_w + (1 - \varepsilon)(\alpha^M p + (1 - \alpha^M)q)$ and fix any $\hat{\rho}$ consistent with \succsim . The observation above implies that $\hat{\rho}(\{p, r\})(p) < \alpha_m$. However, the definition of α_m and the fact that \succsim strictly preserves first-order stochastic dominance imply that $\text{supp}_{\hat{\rho}}(\{p, q, r\}) \subseteq \{p, q\}$ and $\hat{\rho}(\{p, q, r\})(p) \geq \alpha_m$, which is a violation of Regularity.

(ii) implies (iii). By contradiction, assume that \succsim does not have a point of strict convexity. That is, suppose that for all $p, q \in \Delta$ with $p \succsim q$, we have $p \succsim \lambda p + (1 - \lambda)q$ for every $\lambda \in [0, 1]$. Since \succsim is a complete preorder, this is equivalent to say that \succsim has convex lower contour sets. That is, the set

$$\{p \in \Delta : r \succsim p\}$$

is convex for all $r \in \Delta$. Now, let \geq be any linear order (a complete, transitive and antisymmetric binary relation) on Δ (the existence of \geq is guaranteed by the Well Ordering Principle, for example). Define \trianglerighteq to be the relation that applies \succsim and \geq lexicographically. That is, for every $p, q \in \Delta$, $p \trianglerighteq q$ if and only if $p \succ q$ or $p \sim q$ and $p \geq q$. Note that \trianglerighteq is also a linear order on Δ . Finally, let $\hat{\rho}$ be the stochastic choice function that, for each $A \in \mathcal{A}$, assigns probability 1 to the unique maximizer of \trianglerighteq in A . It is clear that $\hat{\rho}$ satisfies Regularity. Moreover, by the definition of \trianglerighteq , if $p \in A$ is such that $\{p\} = \text{supp}_{\hat{\rho}}(A)$, then $p \succsim q$ for every $q \in A$. Since \succsim has convex lower contour sets, this implies that, in fact, $p \succsim q$ for every $q \in \text{co}(A)$. That is, $\hat{\rho}$ is a stochastic choice function consistent with \succsim that satisfies Regularity, which is a contradiction. ■

PROOF OF PROPOSITION 1:

Given the equivalence result in Theorem 2, we may show that (ii) is equivalent to \succsim having a point of strict convexity.

Suppose, then, that \succsim has a point of strict convexity. That is, suppose there exist $p, q \in \Delta$ and $\lambda \in (0, 1)$ such that $\lambda p + (1 - \lambda)q \succ p, q$. Since \succsim preserves strict first-order stochastic dominance, this implies that $p, q \succ \delta_w$. Now, let $\hat{\rho}$ be any stochastic choice function consistent with \succsim . Without loss of generality, suppose that $\hat{\rho}(\{p, q\})(p) \geq 0.5$. Since \succsim preserves strict first-order stochastic dominance, we also have that $\hat{\rho}(\{q, \gamma q + (1 - \gamma)\delta_w\})(q) = 1$. Finally, by the continuity of \succsim , there exists $\gamma \in (0, 1)$ such that $\lambda p + (1 - \lambda)(\gamma q + (1 - \gamma)\delta_w) \succ p$. Since $\hat{\rho}$ is

consistent with \succsim , this implies that $\hat{\rho}(\{p, \gamma q + (1 - \gamma)\delta_w\})(p) < 1$, which violates Strong Stochastic Transitivity.

Conversely, suppose that \succsim does not have a point of strict convexity. Let $\hat{\rho}$ and \supseteq be defined as in the proof of Theorem 2. By the argument there, $\hat{\rho}$ is a stochastic choice function consistent with \succsim . Now suppose $p, q, r \in \Delta$ are such that $\hat{\rho}(\{p, q\})(p) \geq 0.5$ and $\hat{\rho}(\{q, r\})(q) \geq 0.5$. By construction, this is equivalent to $p \supseteq q$ and $q \supseteq r$, which implies that $p \supseteq r$. This now implies that $\hat{\rho}(\{p, r\})(p) = 1$. This shows that $\hat{\rho}$ satisfies Strong Stochastic Transitivity. ■

PROOF OF PROPOSITION 2:

Let \succsim be the Continuous Deliberate Stochastic Choice representation of a stochastic choice function ρ and suppose that ρ satisfies Linearity. Fix any pair of lotteries p and q such that $p \succ q$. We claim that we must have that $\rho(\{p, q\})(p) = 1$. To see that, let λ^* be the minimum value in $[0, 1]$ such that $\lambda^*p + (1 - \lambda^*)q \succsim \lambda p + (1 - \lambda)q$ for every $\lambda \in [0, 1]$. Note that $\lambda^* \neq 0$. Otherwise, we would have that $\lambda^*p + (1 - \lambda^*)q = q \succsim p$, a contradiction with $p \succ q$. If $\lambda^* = 1$, then trivially we obtain that $\rho(\{p, q\})(p) = 1$. Thus, we are left with the case $\lambda^* \in (0, 1)$. Since \succsim represents ρ , we must have that $\rho(\{\lambda^*p + (1 - \lambda^*)q, q\})(\lambda^*p + (1 - \lambda^*)q) = 1$. By Linearity, we get that $\rho(\{p, q\})(p) = 1$. Now fix any pair of lotteries p and q in Δ with $p \succsim q$. For each $\lambda \in (0, 1)$, define $p_\lambda^b = \lambda p + (1 - \lambda)\delta_b$ and $p_\lambda^w = \lambda p + (1 - \lambda)\delta_w$. Since \succsim strictly preserves first-order stochastic dominance, we have that, for any $\lambda \in (0, 1)$, $p_\lambda^b \succ q_\lambda^w$. By what we have just proved, this implies that $\rho(\{p_\lambda^b, q_\lambda^w\})(p_\lambda^b) = 1$. By Linearity, we have that, for any $\alpha \in (0, 1)$ and $r \in \Delta$, $\rho(\{\alpha p_\lambda^b + (1 - \alpha)r, \alpha q_\lambda^w + (1 - \alpha)r\})(\alpha p_\lambda^b + (1 - \alpha)r) = 1$. Since \succsim represents ρ , this implies that $\alpha p_\lambda^b + (1 - \alpha)r \succsim \alpha q_\lambda^w + (1 - \alpha)r$. Since this is true for any $\lambda \in (0, 1)$, continuity of \succsim implies that $\alpha p + (1 - \alpha)r \succsim \alpha q + (1 - \alpha)r$. We have just shown that, for any $p, q \in \Delta$ with $p \succsim q$, we have $\alpha p + (1 - \alpha)r \succsim \alpha q + (1 - \alpha)r$, for every $\alpha \in (0, 1)$ and $r \in \Delta$. Since \succsim is continuous, it is well-known that this implies that it admits an expected-utility representation. That is, there exists a continuous function $u: [w, b] \rightarrow \mathbb{R}$ such that, for every pair of lotteries p and q in Δ , $p \succsim q$ if, and only if, $E_p(u) \geq E_q(u)$. The proposition is now an immediate consequence of this observation. ■

REFERENCES

- Afriat, S. N. 1967. "The Construction of Utility Functions from Expenditure Data." *International Economic Review* 8 (1): 67–77.
- Agranov, Marina, and Pietro Ortoleva. 2017. "Stochastic Choice and Preferences for Randomization." *Journal of Political Economy* 125 (1): 40–68.
- Ahn, David S., and Todd Sarver. 2013. "Preference for Flexibility and Random Choice." *Econometrica* 81 (1): 341–61.
- Allen, Roy, and John Rehbeck. 2019. "Revealed Stochastic Choice with Attributes." Unpublished.
- Baldassi, Carlo, Simone Cerreia-Vioglio, Fabio Maccheroni, and Massimo Marinacci. 2018. "A behavioral characterization of the Drift Diffusion Model and its multi-alternative extension to choice under time pressure." Unpublished.

- Ballinger, T. Parker, and Nathaniel T. Wilcox. 1997. "Decisions, Error and Heterogeneity." *Economic Journal* 107 (443): 1090–1105.
- Barberá, Salvador, and Prasanta K. Pattanaik. 1986. "Falmagne and the Rationalizability of Stochastic Choices in Terms of Random Orderings." *Econometrica* 54 (3): 707–15.
- Becker, Gordon M., Morris H. DeGroot, and Jacob Marschak. 1963. "An Experimental Study of Some Stochastic Models for Wagers." *Behavioral Science* 8 (3): 199–202.
- Bogacz, Rafal, Eric Brown, Jeff Moehlis, Philip Holmes, and Jonathan D. Cohen. 2006. "The Physics of Optimal Decision Making: A Formal Analysis of Models of Performance in Two-Alternative Forced-Choice Tasks." *Psychological Review* 113 (4): 700–65.
- Bridges, Douglas S., and Ghanshyam B. Mehta. 1995. *Representations of Preferences Orderings*. New York: Springer-Verlag.
- Busemeyer, Jerome R., and James T. Townsend. 1993. "Decision Field Theory: A Dynamic-Cognitive Approach to Decision Making in an Uncertain Environment." *Psychological Review* 100 (3): 432–59.
- Camerer, Colin F. 1989. "Does the Basketball Market Believe in the 'Hot Hand'?" *American Economic Review* 79 (5): 1257–61.
- Camerer, Colin F., and Teck-Hua Ho. 1994. "Violations of the Betweenness Axiom and Nonlinearity in Probability." *Journal of Risk and Uncertainty* 8 (2): 167–96.
- Caplin, Andrew, and Mark Dean. 2015. "Revealed Preference, Rational Inattention, and Costly Information Acquisition." *American Economic Review* 105 (7): 2183–2203.
- Cerreia-Vioglio, Simone. 2009. "Maxmin Expected Utility on a Subjective State Space: Convex Preferences under Risk." Unpublished.
- Cerreia-Vioglio, Simone, David Dillenberger, and Pietro Ortoleva. 2015. "Cautious Expected Utility and the Certainty Effect." *Econometrica* 83 (3): 693–728.
- Cerreia-Vioglio, Simone, David Dillenberger, Pietro Ortoleva, and Gil Riella. 2018. "Cautious Stochastic Choice." Unpublished.
- Chambers, Christopher P., and Federico Echenique. 2016. *Revealed Preference Theory*. Cambridge: Cambridge University Press.
- Chew, S. H. 1989. "Axiomatic Utility Theories with the Betweenness Property." *Annals of Operations Research* 19 (1): 273–98.
- Chew, S. H., L. G. Epstein, and U. Segal. 1991. "Mixture Symmetry and Quadratic Utility." *Econometrica* 59 (1): 139–63.
- Clark, Stephen A. 1996. "The Random Utility Model with an Infinite Choice Space." *Economic Theory* 7 (1): 179–89.
- Cohen, Michael A. 1980. "Random Utility Systems—The Infinite Case." *Journal of Mathematical Psychology* 22 (1): 1–23.
- Dekel, Eddie. 1986. "An Axiomatic Characterization of Preferences under Uncertainty: Weakening the Independence Axiom." *Journal of Economic Theory* 40 (2): 304–18.
- Dwenger, Nadja, Dorothea Kübler, and Georg Weizsäcker. 2018. "Flipping a Coin: Evidence from University Applications." *Journal of Public Economics* 167: 240–50.
- Falmagne, J. C. 1978. "A Representation Theorem for Finite Random Scale Systems." *Journal of Mathematical Psychology* 18 (1): 52–72.
- Fishburn, Peter C. 1983. "Transitive Measurable Utility." *Journal of Economic Theory* 31 (2): 293–317.
- Fudenberg, Drew, Ryota Iijima, and Tomasz Strzalecki. 2015. "Stochastic Choice and Revealed Perturbed Utility." *Econometrica* 83 (6): 2371–2409.
- Fudenberg, Drew, and Tomasz Strzalecki. 2015. "Dynamic Logit with Choice Aversion." *Econometrica* 83 (2): 651–91.
- Gul, Faruk, Paulo Natenzon, and Wolfgang Pesendorfer. 2014. "Random Choice as Behavioral Optimization." *Econometrica* 82 (5): 1873–1912.
- Gul, Faruk, and Wolfgang Pesendorfer. 2006. "Random Expected Utility." *Econometrica* 74 (1): 121–46.
- Harless, David W., and Colin F. Camerer. 1994. "The Predictive Utility of Generalized Expected Utility Theories." *Econometrica* 62 (6): 1251–89.
- Harsanyi, John C. 1973. "Games with Randomly Disturbed Payoffs: A New Rationale for Mixed-Strategy Equilibrium Points." *International Journal of Game Theory* 2 (1): 1–23.
- Hey, John D. 2001. "Does Repetition Improve Consistency?" *Experimental Economics* 4 (1): 5–54.
- Hey, John D., and Chris Orme. 1994. "Investigating Generalizations of Expected Utility Theory Using Experimental Data." *Econometrica* 62 (6): 1291–1326.
- Johnson, Eric J., and Roger Ratcliff. 2013. "Computational and Process Models of Decision Making in Psychology and Behavioral Economics." In *Neuroeconomics: Decision Making and the Brain*, edited by Paul W. Glimcher and Ernst Fehr, 35–48. London: Academic Press.

- Loomes, Graham, and Robert Sugden. 1995. "Incorporating a Stochastic Element into Decision Theories." *European Economic Review* 39 (3–4): 641–48.
- Luce, R. Duncan. 1958. "A Probabilistic Theory of Utility." *Econometrica* 26 (2): 193–224.
- Luce, R. Duncan. 1959. *Individual Choice Behavior: A Theoretical Analysis*. New York: Wiley.
- Machina, Mark J. 1985. "Stochastic Choice Functions Generated from Deterministic Preferences over Lotteries." *Economic Journal* 95 (379): 575–94.
- Manzini, Paola, and Marco Mariotti. 2007. "Sequentially Rationalizable Choice." *American Economic Review* 97 (5): 1824–39.
- Manzini, Paola, and Marco Mariotti. 2014. "Stochastic Choice and Consideration Sets." *Econometrica* 82 (3): 1153–76.
- Marley, A. A. J. 1997. "Probabilistic Choice as a Consequence of Nonlinear (Sub)Optimization." *Journal of Mathematical Psychology* 41 (4): 382–91.
- Masatlioglu, Yusufcan, Daisuke Nakajima, and Erkut Y. Ozbay. 2012. "Revealed Attention." *American Economic Review* 102 (5): 2183–2205.
- Matějka, Filip, and Alisdair McKay. 2015. "Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model." *American Economic Review* 105 (1): 272–98.
- McFadden, Daniel L. 2006. "Revealed Stochastic Preference: A Synthesis." In *Rationality and Equilibrium, Studies in Economic Theory*, Vol. 26, edited by C. D. Aliprantis et al., 1–20. Berlin: Springer.
- McFadden, Daniel, and M. K. Richter. 1991. "Stochastic Rationality and Revealed Stochastic Preference." In *Preferences, Uncertainty, and Rationality*, edited by John S. Chipman, D. McFadden, and K. Richter, 161–86. Boulder, CO: Westview Press.
- Miao, Bin, and Songfa Zhong. 2018. "Probabilistic Social Preference: How Machina's Mom Randomizes Her Choice." *Economic Theory* 65 (1): 1–24.
- Natanzon, Paulo. 2019. "Random Choice and Learning." *Journal of Political Economy* 127 (1): 419–57.
- Nishimura, Hiroki, Efe A. Ok, and John K.-H. Quah. 2017. "A Comprehensive Approach to Revealed Preference Theory." *American Economic Review* 107 (4): 1239–63.
- Ok, Efe A., Pietro Ortoleva, and Gil Riella. 2015. "Revealed (P)Reference Theory." *American Economic Review* 105 (1): 299–321.
- Ratcliff, Roger, and Gail McKoon. 2008. "The Diffusion Decision Model: Theory and Data for Two-Choice Decision Tasks." *Neural Computation* 20 (4): 873–922.
- Ratcliff, Roger, and Philip L. Smith. 2004. "A Comparison of Sequential Sampling Models for Two-Choice Reaction Time." *Psychological Review* 111 (2): 333–67.
- Regenwetter, Michel, Jason Dana, and Clinton P. Davis-Stober. 2011. "Transitivity of Preferences." *Psychological Review* 118 (1): 42–56.
- Regenwetter, Michel, and Clinton P. Davis-Stober. 2012. "Behavioral Variability of Choices versus Structural Inconsistency of Preferences." *Psychological Review* 119 (2): 408–16.
- Rieskamp, Jörg, Jerome R. Busemeyer, and Barbara A. Mellers. 2006. "Extending the Bounds of Rationality: Evidence and Theories of Preferential Choice." *Journal of Economic Literature* 44 (3): 631–61.
- Rubinstein, Ariel. 2002. "Irrational Diversification in Multiple Decision Problems." *European Economic Review* 46 (8): 1369–1378.
- Sandroni, Alvaro, Sandra Ludwig, and Philipp Kircher. 2013. "On the Difference between Social and Private Goods." *B.E. Journal of Theoretical Economics* 13 (1).
- Simonson, Itamar. 1989. "Choice Based on Reasons: The Case of Attraction and Compromise Effects." *Journal of Consumer Research* 16 (2): 158–74.
- Soltani, Alireza, Benedetto De Martino, and Colin Camerer. 2012. "A Range-Normalization Model of Context-Dependent Choice: A New Model and Evidence." *PLoS Computational Biology* 8 (7): e1002607.
- Sopher, Barry, and J. Mattison Narramore. 2000. "Stochastic Choice and Consistency in Decision Making under Risk: An Experimental Study." *Theory and Decision* 48 (4): 323–49.
- Starmer, Chris, and Robert Sugden. 1989. "Probability and Juxtaposition Effects: An Experimental Investigation of the Common Ratio Effect." *Journal of Risk and Uncertainty* 2 (2): 159–78.
- Stoye, Jörg. 2015. "Choice Theory When Agents Can Randomize." *Journal of Economic Theory* 155: 131–51.
- Swait, Joffre, and A. A. J. Marley. 2013. "Probabilistic Choice (Models) as a Result of Balancing Multiple Goals." *Journal of Mathematical Psychology* 57 (1–2): 1–14.
- Thurstone, L. L. 1927. "A Law of Comparative Judgment." *Psychological Review* 34 (4): 273–86.
- Tversky, Amos. 1969. "Intransitivity of Preferences." *Psychological Review* 76 (1): 31–48.

- Tversky, Amos, and Itamar Simonson.** 1993. "Context-Dependent Preferences." *Management Science* 39 (10): 1179–89.
- Woodford, Michael.** 2014. "Stochastic Choice: An Optimizing Neuroeconomic Model." *American Economic Review* 104 (5): 495–500.
- Wu, George, and Richard Gonzalez.** 1996. "Curvature of the Probability Weighting Function." *Management Science* 42 (12): 1676–90.

This article has been cited by:

1. Pedram Heydari. 2021. Luce arbitrates: Stochastic resolution of inner conflicts. *Games and Economic Behavior* **126**, 33-74. [[Crossref](#)]
2. Carlo Baldassi, Simone Cerreia-Vioglio, Fabio Maccheroni, Massimo Marinacci, Marco Pirazzini. 2020. A Behavioral Characterization of the Drift Diffusion Model and Its Multialternative Extension for Choice Under Time Pressure. *Management Science* **66**:11, 5075-5093. [[Crossref](#)]
3. David Hobson, Matthew Zeng. 2020. Randomised rules for stopping problems. *Journal of Applied Probability* **57**:3, 981-1004. [[Crossref](#)]
4. Shaowei Ke, Qi Zhang. 2020. Randomization and Ambiguity Aversion. *Econometrica* **88**:3, 1159-1195. [[Crossref](#)]
5. Matthew Kovach, Gerelt Tserenjigmid. 2019. Behavioral Foundations of Nested Stochastic Choice and Nested Logit. *SSRN Electronic Journal* . [[Crossref](#)]
6. Alex Rees-Jones, Ran I. Shorrer, Chloe Tergiman. 2019. Correlation Neglect in Student-to-School Matching. *SSRN Electronic Journal* . [[Crossref](#)]
7. Julien Monardo. 2019. Shape Restrictions for Structural Demand Estimation: Beyond the ARUM. *SSRN Electronic Journal* . [[Crossref](#)]
8. Roy Allen, Paweł Dziwulski, John Rehbeck. 2019. Revealed Statistical Consumer Theory. *SSRN Electronic Journal* . [[Crossref](#)]
9. Sean Duffy, Steven Gussman, John Smith. 2019. Judgments of Length in the Economics Laboratory: Are There Brains in Choice?. *SSRN Electronic Journal* . [[Crossref](#)]
10. Junnan He, Paulo Natenzon. 2018. Moderate Expected Utility. *SSRN Electronic Journal* . [[Crossref](#)]
11. Vicky Henderson, David Hobson, Matthew Zeng. 2018. Cautious Stochastic Choice, Optimal Stopping and Deliberate Randomization. *SSRN Electronic Journal* . [[Crossref](#)]