Efficient Coding and Choice Behavior

Michael Woodford

Columbia University

RLDM 2015 University of Alberta June 7-10, 2015

Anomalies in Choice Behavior

• In experiments, subjects often make choices that violate normative principles of rational choice

Anomalies in Choice Behavior

- In experiments, subjects often make choices that violate normative principles of rational choice
- "Behavioral economics" literature proposes to model such choices as the result of basing choice on some transformed description of the situation, that introduces biases, rather than on a correct description

Anomalies in Choice Behavior

- In experiments, subjects often make choices that violate normative principles of rational choice
- "Behavioral economics" literature proposes to model such choices as the result of basing choice on some transformed description of the situation, that introduces biases, rather than on a correct description
- But this raises a question: why should the brain produce and use incorrect descriptions, rather than more accurate ones?

Problem

In addition to whatever you own, you have been given 1000. You are now asked to choose between (a) winning an additional 500 with certainty, or (b) a gamble with a 50 percent chance of winning 1000 and a 50 percent chance of winning nothing.

Problem

In addition to whatever you own, you have been given 1000. You are now asked to choose between (a) winning an additional 500 with certainty, or (b) a gamble with a 50 percent chance of winning 1000 and a 50 percent chance of winning nothing.

Problem

In addition to whatever you own, you have been given 2000. You are now asked to choose between (a) losing 500 with certainty, and (b) a gamble with a 50 percent chance of losing 1000 and a 50 percent chance of losing nothing.

Problem

In addition to whatever you own, you have been given 1000. You are now asked to choose between (a) winning an additional 500 with certainty, or (b) a gamble with a 50 percent chance of winning 1000 and a 50 percent chance of winning nothing.

Majority of subjects [84%] choose (a)

Problem

In addition to whatever you own, you have been given 2000. You are now asked to choose between (a) losing 500 with certainty, and (b) a gamble with a 50 percent chance of losing 1000 and a 50 percent chance of losing nothing.

Majority of subjects [69%] choose (b)

Puzzling aspects of this behavior:

In both cases, subjects are choosing between the same probability distributions over final wealth levels:

(a) initial wealth + 1500 with certainty

VS

(b) 50 percent chance of initial wealth + 1000, 50 percent chance of initial wealth + 2000 Puzzling aspects of this behavior:

In both cases, subjects are choosing between the same probability distributions over final wealth levels:

(a) initial wealth + 1500 with certainty

VS

- (b) 50 percent chance of initial wealth + 1000, 50 percent chance of initial wealth + 2000
- Explanation in prospect theory: an **"isolation effect"** is invoked to predict that subjects consider only gains or losses that result from choice of (a) or (b), in isolation from context of having a certain initial wealth

Kahneman-Tversky (1979)

Even considering gains and losses in isolation: choices (a) and
(b) have equal expected value

 moreover, risk aversion in case of first problem, but risk-seeking in second ["reflection effect"]

Kahneman-Tversky (1979)

Even considering gains and losses in isolation: choices (a) and
(b) have equal expected value

— moreover, risk aversion in case of first problem, but risk-seeking in second ["reflection effect"]

 Explanation in prospect theory: subject evaluates not average of true net gain x, but instead average of v(x), where "value function" v(x) is a nonlinear transformation

— v(x) increasing but **concave** for $x > 0 \Rightarrow$ risk-averse in domain of **gains**

— v(x) increasing and **convex** for $x < 0 \Rightarrow$ risk-seeking in domain of **losses**

Kahneman-Tversky (1979)

 Even considering gains and losses in isolation: choices (a) and (b) have equal expected value

 moreover, risk aversion in case of first problem, but risk-seeking in second ["reflection effect"]

 Explanation in prospect theory: subject evaluates not average of true net gain x, but instead average of v(x), where "value function" v(x) is a nonlinear transformation

— v(x) increasing but **concave** for $x > 0 \Rightarrow$ risk-averse in domain of **gains**

— v(x) increasing and **convex** for $x < 0 \Rightarrow$ risk-seeking in domain of **losses**

• But why such a nonlinear transformation?

Woodford (Columbia)

A Proposed Explanation

- Possible explanation, analogous to explanation of some kinds of perceptual biases:
 - decisions are based on a **noisy mental representation** of the decision situation, rather than a precisely correct representation
 - the representation is imprecise because of the **finite processing capacity** of the circuits devoted to this task

A Proposed Explanation

- Possible explanation, analogous to explanation of some kinds of perceptual biases:
 - decisions are based on a **noisy mental representation** of the decision situation, rather than a precisely correct representation
 - the representation is imprecise because of the **finite processing capacity** of the circuits devoted to this task
 - "biases" may actually represent **Bayes-optimal** judgments, conditional on the imperfect representation that is available as a basis for the judgments

A Proposed Explanation

- Possible explanation, analogous to explanation of some kinds of perceptual biases:
 - decisions are based on a **noisy mental representation** of the decision situation, rather than a precisely correct representation
 - the representation is imprecise because of the **finite processing capacity** of the circuits devoted to this task
 - "biases" may actually represent **Bayes-optimal** judgments, conditional on the imperfect representation that is available as a basis for the judgments
 - moreover, the particular nature of the imprecision in the mental representation may be explained as an efficient use of finite processing capacity ["efficient coding hypothesis"]

• How can biased judgments be optimal?

- How can biased judgments be optimal?
- If judgment has to be based on a mental representation of the situation that is **stochastically** related to the actual situation
 - conditional probabilities p(s|x) of different subjective representations s in the case of any actual state x—
 - then it is not possible, even in principle, to perfectly recover the actual state from the subjective representation s

- How can biased judgments be optimal?
- If judgment has to be based on a mental representation of the situation that is **stochastically** related to the actual situation
 - conditional probabilities p(s|x) of different subjective representations s in the case of any actual state x—
 - then it is not possible, even in principle, to perfectly recover the actual state from the subjective representation s
- One can at best suppose that judgment is based on an optimal estimate of the state, $\hat{x}(s)$

— e.g., **posterior mean** of x, where posterior is conditional on s (using Bayes' Rule)

• There will then be a **probability distribution** of subjective estimates \hat{x} associated with any actual state x (resulting from randomness of the subjective representation s),

— and the estimates need not equal the actual value even **on average**:

$$\mathrm{E}[\hat{x}|x] \neq x$$

• There will then be a **probability distribution** of subjective estimates \hat{x} associated with any actual state x (resulting from randomness of the subjective representation s),

— and the estimates need not equal the actual value even **on average**:

$$\mathrm{E}[\hat{x}|x] \neq x$$

 Note that this would not be true if one assumed that the subjective representation is simply a nonlinear transformation of the state, such as the K-T value function

— if s = v(x), optimal estimate would be $\hat{x} = v^{-1}(s) = x$ for all x

• Is there a reason to treat the subjective representation of a situation as random?

- Is there a reason to treat the subjective representation of a situation as random?
- In lab experiments, subjects' choices often involve a random element: same subject need not choose the same way, if same options are repeated (sometimes only minutes later)

— yet probability of choice often varies systematically with characteristics of options presented

Mosteller and Nogee (1951)



Woodford (Columbia)

Efficient Coding

June 2015 11 / 44

- Is there a reason to treat the subjective representation of a situation as random?
- In lab experiments, subjects' choices often involve a random element: same subject need not choose the same way, if same options are repeated (sometimes only minutes later)

— yet probability of choice often varies systematically with characteristics of options presented

- Theories like prospect theory make a **deterministic** prediction about such choices
 - goal is to correctly predict modal behavior

— but a more complete theory would explain the random variation in choices as well

Woodford (Columbia)

 Randomness of responses a common feature of perceptual judgments

— Mosteller-Nogee figure analogous to a "psychometric function"

 Randomness of responses a common feature of perceptual judgments

— Mosteller-Nogee figure analogous to a "psychometric function"

• Standard explanation (dating back to Fechner): judgment based on a random "percept" *s*

— a distinct probability distribution $p(\boldsymbol{s}|\boldsymbol{x})$ for each true stimulus magnitude \boldsymbol{x}

 Randomness of responses a common feature of perceptual judgments

— Mosteller-Nogee figure analogous to a "psychometric function"

• Standard explanation (dating back to Fechner): judgment based on a random "percept" *s*

— a distinct probability distribution $p(\boldsymbol{s}|\boldsymbol{x})$ for each true stimulus magnitude \boldsymbol{x}

 now understood in terms of stochastic responses of neurons in cortical regions involved in sensory processing (described by "tuning curves")

• When should **randomness of coding** of a magnitude result in **bias** in the average estimate of that magnitude, if optimal use is made of the information that has been coded?

- When should **randomness of coding** of a magnitude result in **bias** in the average estimate of that magnitude, if optimal use is made of the information that has been coded?
- If estimate $\hat{x}(s)$ is the posterior mean, then

$$\mathbf{E}[\hat{x}|x] = \int \tilde{x} \, \pi(\tilde{x}) \, L(\tilde{x}|x) \, d\tilde{x}$$

where $\pi(x)$ is the **prior** over possible values of the magnitude, and

$$L(\tilde{x}|x) \equiv \int \frac{p(s|\tilde{x})p(s|x)}{p(s)} \, ds$$

is the average relative likelihood of state \tilde{x} , averaging over the subjective representations *s* produced by state *x*

Woodford (Columbia)

$$\mathrm{E}[\hat{x}|x] = \int \tilde{x} \, \pi(\tilde{x}) \, L(\tilde{x}|x) \, d\tilde{x}$$

• So $E[\hat{x}|x]$ can differ from x either because

- prior $\pi(\tilde{x})$ is not symmetrical around x, or
- average relative likelihood $L(\tilde{x}|x)$ is not symmetrical around x

$$\mathbf{E}[\hat{x}|\mathbf{x}] = \int \tilde{x} \, \pi(\tilde{x}) \, L(\tilde{x}|\mathbf{x}) \, d\tilde{x}$$

• So $E[\hat{x}|x]$ can differ from x either because

- prior $\pi(\tilde{x})$ is not symmetrical around x, or
- average relative likelihood $L(\tilde{x}|x)$ is not symmetrical around x
- First effect: bias toward the **prior mean**, if average relative likelihood is not very sharply peaked around x (i.e., coding is imprecise)

- effect is stronger, the less the precision of the coding

$$\mathrm{E}[\hat{x}|x] = \int \tilde{x} \, \pi(\tilde{x}) \, L(\tilde{x}|x) \, d\tilde{x}$$

• So $E[\hat{x}|x]$ can differ from x either because

- prior $\pi(\tilde{x})$ is not symmetrical around x, or
- average relative likelihood $L(\tilde{x}|x)$ is not symmetrical around x
- First effect: bias toward the **prior mean**, if average relative likelihood is not very sharply peaked around x (i.e., coding is imprecise)

- effect is stronger, the less the precision of the coding

— so a **non-linear bias** can result, if **degree of precision of coding** varies with *x*

Woodford (Columbia)

$$\mathbf{E}[\hat{x}|\mathbf{x}] = \int \tilde{x} \, \pi(\tilde{x}) \, L(\tilde{x}|\mathbf{x}) \, d\tilde{x}$$

Second effect: even with uniform prior (or symmetric around x), likelihood L(x̃|x) can be asymmetric if degree of precision of coding varies with x

— less precise coding as x increases $\Rightarrow L(\tilde{x}|x)$ larger for $\tilde{x} > x$ than for $\tilde{x} < x$ to same extent

$$\mathbf{E}[\hat{x}|\mathbf{x}] = \int \tilde{x} \, \pi(\tilde{x}) \, L(\tilde{x}|\mathbf{x}) \, d\tilde{x}$$

Second effect: even with uniform prior (or symmetric around x), likelihood L(x̃|x) can be asymmetric if degree of precision of coding varies with x

— less precise coding as x increases $\Rightarrow L(\tilde{x}|x)$ larger for $\tilde{x} > x$ than for $\tilde{x} < x$ to same extent

- results in bias toward the direction of less precise coding

$$\mathbf{E}[\hat{x}|\mathbf{x}] = \int \tilde{x} \, \pi(\tilde{x}) \, L(\tilde{x}|\mathbf{x}) \, d\tilde{x}$$

Second effect: even with uniform prior (or symmetric around x), likelihood L(x̃|x) can be asymmetric if degree of precision of coding varies with x

— less precise coding as x increases $\Rightarrow L(\tilde{x}|x)$ larger for $\tilde{x} > x$ than for $\tilde{x} < x$ to same extent

- results in bias toward the direction of less precise coding
- So more complex biases result from **non-uniformity of the precision of coding** over the range of *x*

Efficient Coding

• But why should there be non-uniformity in the precision of coding?
Efficient Coding

- But why should there be non-uniformity in the precision of coding?
- This is required (quite generally) by the **efficient coding hypothesis**:

— a **finite capacity** for making discriminations among alternative situations is allocated in a way that is **most useful** to the organism, for reducing **average uncertainty** about the situation that exists on a given occasion

Example: Discrimination of Orientation

 Well-established that humans (and animals) can make sharper discriminations between differing orientations that are near-vertical or near-horizontal, than between oblique orientations ("oblique effect": Appelle, 1972)

Orientation Discrimination



Example: Discrimination of Orientation

- Well-established that humans (and animals) can make sharper discriminations between differing orientations that are near-vertical or near-horizontal, than between oblique orientations ("oblique effect": Appelle, 1972)
- Animal neurophysiology studies (e.g., of macaque V1) show this explained by allocation of greater processing resources to the former types of discriminations:
 - larger number of neurons with "preferred orientation" near vertical or horizontal than near oblique angles
 - narrower "tuning widths" for neurons with preferred orientations near vertical or horizontal

(Mansfield, 1974; Li et al., 2003; Wang et al., 2003)

Orientation Discrimination



Woodford (Columbia)

Efficient Coding

June 2015 21 / 44

Example: Discrimination of Orientation

• This can be viewed as **efficient** given the fact that in both natural and man-made environments, horizontally and vertically oriented edges **occur more frequently** than oblique orientations

Orientation Discrimination



23 / 44

Orientation Discrimination



Woodford (Columbia)

Efficient Coding

June 2015 24 / 44

Example: Discrimination of Orientation

- This can be viewed as **efficient** given the fact that in both natural and man-made environments, horizontally and vertically oriented edges **occur more frequently** than oblique orientations
- Specifically, a neural coding scheme in which
 - cell density \sim environmental frequency
 - width of tuning curve $\sim 1/{\rm frequency}$

and hence discrimination thresholds vary inversely with environmental frequency, is efficient in the sense of maximizing the **mutual information** between the stimulus and the neural activation state (Ganguli and Simoncelli, 2012)

Example: Discrimination of Orientation

 This is equivalent to saying that the neural coding scheme minimizes average posterior uncertainty about the stimulus orientation, where uncertainty is measured using Shannon's entropy measure

• This non-uniform allocation of processing resources has additional implications, beyond the non-uniformity of discrimination thresholds

— it also produces $\ensuremath{\textbf{biases}}$ in perceived orientation of visual stimuli

• This non-uniform allocation of processing resources has additional implications, beyond the non-uniformity of discrimination thresholds

— it also produces $\ensuremath{\textbf{biases}}$ in perceived orientation of visual stimuli

- Using a mathematical description of the non-uniformity of
 - the prior probability distribution for stimulus orientations
 - the distribution of preferred orientations for neurons
 - Sthe tuning widths of neurons with different preferred orientations

just mentioned, it is possible to derive the predicted **average** Bayesian estimate of orientation $E[\hat{\theta}|\theta]$

 Result: estimates are biased away from "cardinal" orientations (Wei and Stocker, 2012)

Predicted Bias in Average Perceived Orientation



Woodford (Columbia)

June 2015 29 / 44

- Result: estimates are biased away from "cardinal" orientations (Wei and Stocker, 2012)
- Tomassini *et al.* (2010) document this bias, in an experiment that asks subjects to align dots with the **perceived orientation** of oriented "Gabor patches"

Tomassini et al. (2010)



Woodford (Columbia)

Efficient Coding

June 2015 31 / 44

Efficient Coding and Value-Based Choice

- Further conjecture: in value-based decisionmaking as well, decisions are based on **subjective representations** of the available options that are random, but in a way that is **efficient**
 - subject to an upper bound on processing capacity (degree of differentiation of posteriors)
 - and for a particular frequency distribution of possible choice situations (environment to which coding scheme is adapted)

Efficient Coding and Value-Based Choice

- Further conjecture: in value-based decisionmaking as well, decisions are based on **subjective representations** of the available options that are random, but in a way that is **efficient**
 - subject to an upper bound on processing capacity (degree of differentiation of posteriors)
 - and for a particular frequency distribution of possible choice situations (environment to which coding scheme is adapted)
- Efficiency must however be defined in terms of the payoffs in a particular class of decision problems: average reduction of payoff-relevant uncertainty (need not mean entropy reduction, as in infomax theory)

Modeling the choice situation in the KT experiment:

• DM must make a judgment about the **relative value** of two lotteries (*a*, *b*), after first receiving an **initial amount** *w*

Modeling the choice situation in the KT experiment:

- DM must make a judgment about the **relative value** of two lotteries (*a*, *b*), after first receiving an **initial amount** *w*
- Each lottery characterized by a payment x_i in each of two equi-probable states i = 1, 2

— hence DM's situation in the event of a given choice characterized by 3 numbers: w, x_1, x_2

Application to the KT Experiment

 Assume DM cares only about expected final wealth: values options (if fully informed) at

$$v = w + \frac{1}{2}x_1 + \frac{1}{2}x_2$$

Application to the KT Experiment

 Assume DM cares only about expected final wealth: values options (if fully informed) at

$$v = w + \frac{1}{2}x_1 + \frac{1}{2}x_2$$

• Suppose task is to estimate the **relative value** of option *a*, $v^a - v^b$

- more specifically, to produce an estimate e that minimizes

$$\mathrm{E}[(e - (v^a - v^b))^2]$$

• Assume DM must estimate relative value on the basis of a **subjective representation** of the choice situation,

$$r = (r_0; r_1^a, r_2^a; r_1^b, r_2^b)$$

where the 5 elements of *r* correspond to observations of the 5 **attributes** $x = (w; x_1^a, x_2^a; x_1^b, x_2^b)$ of the choice situation

• Assume DM must estimate relative value on the basis of a **subjective representation** of the choice situation,

$$r = (r_0; r_1^a, r_2^a; r_1^b, r_2^b)$$

where the 5 elements of *r* correspond to observations of the 5 **attributes** $x = (w; x_1^a, x_2^a; x_1^b, x_2^b)$ of the choice situation

• Attributes must be coded **independently**: coding described by conditional probabilities

$$p_0(r_0|w)$$
, $p_i^m(r_i^m|x_i^m)$ for $m = a, b; i = 1, 2$

with conditional independence of each component of r

• Constraint on precision of coding:

• let the information-processing resources required by a proposed coding scheme be measured by the (Shannon) channel capacity C(p) of the communication channel that produces output signal r with probabilities p(r|x) when supplied with input x

— essentially, a measure of the [log of the] effective number of categories of inputs that can be distinguished, but also defined for "fuzzy" (probabilistic) categories

• Constraint on precision of coding:

let the information-processing resources required by a proposed coding scheme be measured by the (Shannon) channel capacity C(p) of the communication channel that produces output signal r with probabilities p(r|x) when supplied with input x

— essentially, a measure of the [log of the] effective number of categories of inputs that can be distinguished, but also defined for "fuzzy" (probabilistic) categories

• assume a finite upper bound

$$C(p_0) + \sum_i \sum_m C(p_i^m) \leq \bar{C}$$

so that finer distinctions are possible in the case of any attribute only at the cost of making coarser distinctions with regard to some other attributes

Woodford (Columbia)

- Finally, suppose that coding is optimized for a particular **prior** over possible choice situations **x**:
 - the separate attributes $w, x_1^a, x_2^a, x_1^b, x_2^b$ are **independently** drawn from their respective prior distributions
 - prior distribution for each state-contingent payoff x_i^m is same
 [numerical example: normal, mean zero]

- **(**) Relative value of two options independent of $w \Rightarrow C_0 = 0$
 - no capacity used to represent value of w
 - hence choice must be a function only of distributions of gains or losses

() Relative value of two options independent of $w \Rightarrow C_0 = 0$

- no capacity used to represent value of w
- hence choice must be a function only of distributions of gains or losses
- no need for a separate "isolation principle" to deliver this result: follows from same efficient coding consideration as will be used to explain risk attitudes

2 Each state-contingent payoff x_i^m is coded using conditional probabilities p(r|x) that solve

$$\min \mathrm{E}[(\hat{x}(r) - x)^2] \qquad \text{s.t.} \quad C(p) \leq \bar{C}/4$$

Each state-contingent payoff x_i^m is coded using conditional probabilities p(r|x) that solve

$$\min \mathrm{E}[(\hat{x}(r) - x)^2] \qquad \text{s.t.} \quad C(p) \leq \bar{C}/4$$

• Finite capacity \Rightarrow efficient coding makes $\mathrm{E}[\hat{x}|x]$ a nonlinear function of x

Mean Estimated Value vs. True Value



Woodford (Columbia)

June 2015 40 / 44

Explaining the KT Experiment

Prediction of the model proposed here:

• DM should choose lottery a iff

$$\hat{x}(r_1^a) + \hat{x}(r_2^a) > \hat{x}(r_1^b) + \hat{x}(r_2^b)$$

Explaining the KT Experiment

Prediction of the model proposed here:

• DM should choose lottery a iff

$$\hat{x}(r_1^a) + \hat{x}(r_2^a) > \hat{x}(r_1^b) + \hat{x}(r_2^b)$$

• For each attribute, optimal estimate $\hat{x}(r_i^m)$ and conditional probabilities $p(r_i^m | x_i^m)$ are determined as above

Mean Subjective Valuations of Lotteries



(a) has higher MSV when w = 1000, but (b) higher when w = 2000

A Further Implication of the Theory

• Viewing the phenomenon as resulting from finite-precision coding, rather than an arbitrary fact about how different things are valued, not only provides a functional explanation, but also implies that it should be present to a greater or lesser extent depending on degree of scarcity of processing capacity
A Further Implication of the Theory

- Viewing the phenomenon as resulting from finite-precision coding, rather than an arbitrary fact about how different things are valued, not only provides a functional explanation, but also implies that it should be present to a greater or lesser extent depending on degree of scarcity of processing capacity
- Suggestive evidence:
 - DeMartino *et al.* (2006): significant correlation between decreased asymmetry between gain and loss domains and higher activity in rOFC and vmPFC

A Further Implication of the Theory

- Viewing the phenomenon as resulting from finite-precision coding, rather than an arbitrary fact about how different things are valued, not only provides a functional explanation, but also implies that it should be present to a greater or lesser extent depending on degree of scarcity of processing capacity
- Suggestive evidence:
 - DeMartino *et al.* (2006): significant correlation between decreased asymmetry between gain and loss domains and higher activity in rOFC and vmPFC
 - Porcelli and Delgado (2009): acute stress results in increased asymmetry between gain and loss domains

Conclusions

- Important biases in choice behavior can result from decisions based on imprecise subjective coding of features of the choice situation
- It may be possible to understand the form of such representations using similar principles to those that explain aspects of **perceptual coding** in sensory domains

— in particular, efficient allocation of scarce processing capacity

Conclusions

- Important biases in choice behavior can result from decisions based on imprecise subjective coding of features of the choice situation
- It may be possible to understand the form of such representations using similar principles to those that explain aspects of **perceptual coding** in sensory domains

— in particular, efficient allocation of scarce processing capacity

• This is one of the more obvious areas in which findings from computational neuroscience can guide theory development in economics

Woodford (Columbia)