Human revealed preferences for multi-component rewards

Alexandre Pastor-Bernier*, Konstantin Volkmann*, Arkadiusz Stasiak, Fabian Grabenhorst and Wolfram Schultz

* These authors contributed equally to the study.

Department of Physiology
Development and Neuroscience
University of Cambridge
Cambridge CB2 3DY
United Kingdom

Corresponding author:
Wolfram Schultz
Department of Physiology, Development and Neuroscience
University of Cambridge
Cambridge CB2 3DY
United Kingdom
Email: ws234direct@gmail.com

Email addresses of all authors:
Alexandre Pastor-Bernier: pastor.bernier@gmail.com
Konstantin Volkmann: konstantin.volkmann@googlemail.com
Arkadiusz Stasiak: as863@cam.ac.uk
Fabian Grabenhorst: fabian.grabenhorst@googlemail.com
Wolfram Schultz: Wolfram.Schultz@protonmail.com

Running Title:
Human preferences for multi-component choice options

Keywords: Bundle, decision-making, stochastic choice, psychophysics, monkey

Acknowledgements. We thank Charles R. Plott for discussions and conceptual support, Steve Edgley for help and logistic support, Simone Ferrari-Toniolo for comments on experimental economics and Francine Ganter-Restrepo for engineering expertise. This work was supported by the Wellcome Trust (WT 095495, WT 204811).

Abstract

Realistic, everyday rewards contain multiple components. An apple has taste and size. However, we choose in single dimensions, simply preferring some apples to others. How can such single-dimensional preference relationships refer to multi-component choice options? Here, we measured how stochastic choices revealed preferences for two-component milkshakes. The preferences were intuitively graphed as indifference curves that represented the orderly integration of the two components as trade-off: parts of one component were given up for obtaining one additional unit of the other component without a change in preference. Two-component rewards on higher indifference curves were preferred to those on lower curves, despite widely, and sometimes oppositely, varying components. The systematic, non-overlapping curves satisfied leave-one-out tests, followed predictions by machine learning decoders and correlated with single-dimensional Becker-DeGroot-Marschak (BDM) auction-like bids for the two-component rewards. This accuracy suggests a decision process that integrates multiple reward components into single-dimensional estimates in an orderly and systematic fashion. In inter-species comparisons, human performance matched that of highly experienced laboratory monkeys, as measured by parameter accuracy of the critical trade-off between bundle components. These data elucidate the decision mechanisms for multi-component choice options and attest to the validity of the rigorous concepts of Revealed Preference Theory and their convenient graphic schemes for explaining choices of human and non-human primates. The results of this study encourage formal behavioral and neural investigations of normal, irrational and pathological economic choices.

Introduction

We like apples, in particular if they are sweet. We can state our preference in words, but this may not be accurate because of poor introspection, faulty memory or erroneous report. It would be more accurate to observe our actual choice of apples. But how do we make that choice? When choosing apples, we may prefer a sweeter one even if it is a bit smaller; we trade-in some size for more sweetness. Thus, our preferences do not concern any component alone but their combination. Every reward or economic good contains multiple components, attributes or dimensions, and thus constitutes a bundle vector. The bundle components may be integral parts of a good, like the sweetness and size of the apple, or consist of distinct entities, like the steak and vegetable of a meal. Importantly, each component contributes to the choice. Without considering the multi-component nature of choice options, we can only study gross choices or exchanges, like between an apple and a pear (not really a choice for an apple lover), or between a movie ticket and a meal (not good when hungry). Thus, to understand realistic, fine-grained choices, we should consider that choice options have multiple components.

In contrast to the multi-dimensionality of realistic, vectorial choice options, revealed preferences and subjective reward value ('utility') are single-dimensional. When faced with two options, a rational decision maker can only prefer one option, or its alternative, or be indifferent to them (completeness axiom; Von Neumann & Morgenstern 1944; Mas-Colell et al. 1995). With repeated, stochastic choices, preferences are revealed by choice probability (McFadden 2004); the probability of choosing one option over its alternative varies in a graded, scalar manner. Correspondingly, the utility of one choice option can only be higher, lower or equal to that of its alternative. Further, neural signals reflecting choice options vary only along a single dimension at any given moment and thus are also scalar. Hence the question: how can single-dimensional revealed preferences, utility and neural signals concern vectorial, multi-component choice options? In particular, can we test the issue empirically using well worked-out, rigorous theoretical concepts that should reduce possible confounds?

The issue of vectorial-to-scalar transformation can be formally approached by considering that the same scalar measure may arise from oppositely varying vector components. I may trade-in a larger, less sweet apple for a smaller but sweeter one without loss or gain in utility (I do not want to lose anything, nor is it easy to gain anything). Thus, the increase in sweetness is compensated by the decrease in size. This trade-off is captured by the two-dimensional indifference curves (IC) of Revealed Preference Theory that provide a convenient graphic formalism (Fisher 1892; Samuelson 1937; 1938). Equally revealed preferred but differently composed multi-component bundles are plotted on one and the same IC. The continuous nature of these ICs reflects the trade-off that characterises the well-ordered integration of multiple bundle components into a single-dimensional preference relationship and utility. For these reasons, ICs are a mainstay of economic theory and textbooks (Laidler & Estrin 1989; Kreps 1990; Varian 1992; Mas-Colell et al. 1995; Perloff 2009).

Previous research investigated multi-component choice options in several ways. Without referring to the IC graphism, studies tested quality and price of television sets (Simonson 1989), comfort and fuel consumption of cars (Simonson 1989), payoff amount and probability (Tversky 1969; Soltani et al. 2012; Levy et al. 2012), shape and color of cards (Fellows 2016), artificial objects and appendages (Pelletier & Fellows 2019), various food components (Suzuki et al. 2017; DiFeliceantonio et al. 2018; Harris et al. 2018), hats and shoes (Thurstone 1931), pastries and payoffs (MacCrimmon & Toda 1969), and monetary token from two accounts (Choi et al. 2007; Kurtz-David at al. 2019). Studies using ICs were restricted to hypothetical outcomes, such as hats and shoes (Thurstone 1931), pencils and payoff (MacCrimmon & Toda 1969), and different monetary tokens (Choi et al. 2007; Kurtz-David et al. 2019). Without actual experimentation, conceptual studies used the IC scheme for explaining choice inconsistencies, such as violations of the Independence of Irrelevant Alternatives (IIA) (Rieskamp et al. 2006; Chung et al. 2017; Gluth et al. 2017). Despite the importance of these economic concepts, what is lacking are more empirical tests of the IC formalisms in controlled laboratory situations using actual, tangible outcomes. Such experiments should help to elucidate decision mechanisms involving multi-component choice options, empirically scrutinize the theoretical concepts and encourage future behavioral and neuroeconomic studies.

The purpose of this study was to advance our understanding of decision processes for multi-component choice options and to empirically validate the theoretical concepts for further behavioral and neuroeconomic studies on multi-dimensional choice options. Using the rigorous concepts of Revealed Preference Theory, we estimated revealed preferences in normal humans and compared them with those in monkeys. The perspective of serving for further human and animal studies imposed substantial constraints, including specific event timing, repeated trials commensurate with stochastic preferences (McFadden 2004; Stott 2006), and operant choice between bundles of immediately consumable payouts comparable to tangible liquid and food rewards for animals (Kagel et al. 1975; Pastor-Bernier et al. 2017). Sitting in front of a computer monitor, human participants chose between two visually presented bundles that contained the same two fatty and sugary milkshakes in independently varying quantity. The choices conformed with the ICs of economic theory, satisfied leave-out statistics, corresponded well to decoder predictions, and showed mechanism-independence with Becker-DeGroot-Marschak (BDM) bidding. The trade-off accuracy of humans compared well with that of laboratory monkeys. These results encourage future studies of fine-grained 'irrational' IIA anomalies in normal and brain-damaged humans who might fail to properly consider all components of choice options, for studying underlying normal and pathological brain processes. The results should also inform neuroeconomic studies in monkeys whose superb and understandable cognitive and behavioral abilities allow high-resolution neuronal investigations and generalization across primate species.

Method

Implementation of basic concepts. Revealed Preference Theory conceptualizes observable preference relationships between multi-component choice options. Decision makers reveal their preference at the moment of choice. We respected general notions of discrete choice models: the employed choice sets were exhaustive and had finite numbers of options (two) that were mutually exclusive; both options were fully known to each participant and presented simultaneously. The statistical analysis of neural responses in humans and animals requires the use of multiple trials. Therefore, we used a version of Revealed Preference Theory in which preferences are stochastic (McFadden 2004). Thus, we assessed revealed preference from the probability of multiple choices, rather than by traditional single-shot economic tests. The probability of choosing one option over its alternative varies in a single-dimensional, graded manner and thus is scalar. A more complete and traditional description of Revealed Preference Theory includes budget constraints (Mas-Colell et al. 1995); however, in this initial experimental study, we aimed to obtain as straightforward and easily interpretable empirical data as possible and therefore did not test aspects of budget constraint. In particular, the acquisition of any chosen bundle always required exactly one single finger movement (button press on a computer keyboard). We implemented these concepts in human participants as follows:

(1) Participants chose between two bundles that contained the same two distinct milkshakes in independently varying scalar amounts (component A, component B, each measured in ml; Figure 1A). Thus, each choice option consisted of a two-dimensional bundle vector. The milkshakes constituted rewards, as evidenced by the participants' voluntary consumption. We kept effort cost for obtaining any bundle and budget constant for all choice options (single computer keyboard button press).

(2) Stochastic preference was revealed by the measurable probability of repeated choice. The probability of choosing a given bundle depended on all components of all bundles present on that trial (as opposed to 'lexicographic' preferences for only one component) (Figure 1B, c).

(3) Equal choice probability for two bundles indicated equal revealed stochastic preference ($P = 0.5$ each bundle). The point at which both bundles were chosen with equal probability (point of subjective equivalence or choice indifference point, IP) was estimated from an S-shaped psychophysical function fitted to the probabilities of repeated choices while varying one component of one bundle and keeping all other components constant (Figure 1C).

(4) Every bundle had a subjective value ('utility') for the decision maker that depended only on the amounts of both milkshakes. A bundle was chosen with a higher probability than any other bundle in the same option set if and only if its utility was higher than in any other bundle in that

choice set. In other words, the preference relationship between two bundles was monotone if the higher utility of one bundle implied that it was preferred to its alternative (Mas-Colell et al. 1995). A bundle was chosen with equal probability of $P = 0.5$ against another bundle if the two bundles had the same utility. Notions (2) - (4) are compatible with basic assumptions of stochastic choice theories (Luce 1959; McFadden & Richter 1990; McFadden 2004; Stott 2006) (although we analysed the choices with the probit model rather than McFadden's logit model because of its less restrictive assumptions; see below).

(5) Each bundle was graphically displayed at the intersection of the x-coordinate (component A) and y-coordinate (component B) of a two-dimensional graph (Figure 1D). A bundle that was as revealed preferred as its alternative was graphically represented as a two-dimensional IP. Multiple IPs aligned on a single, continuous indifference curve (IC) (Figure 1E). Thus, an IC reflected scalar utility and choice probability for the vectorial bundle in an orderly and systematic manner.

(6) Bundles could be equally revealed preferred despite variation in bundle composition, such that some amount of one component was given up in order to gain one unit of the other component without change in preference (marginal rate of substitution, MRS). This trade-off was graphically characterised by two parameters: (i) the IC slope, which reflected the relative utility (currency) of the two bundle components and could be variable and asymmetric between x-axis and y-axis; (ii) the curvature, which captured any slope change between IC center and IC periphery (Figure 1D).

(7) Bundles with larger amounts of both components were revealed preferred to bundles with smaller amounts (as long as the 'value function' for each component did not decrease with amount, indicating that 'more was better') (Figure 1E). Any bundle above an IC (farther away from the origin) was revealed preferred to any bundle on that IC, any bundle on an IC was revealed preferred to any bundle below that IC (closer to the origin); thus any bundle on a higher IC was revealed preferred to any bundle on a lower IC.

(8) The preference relationship between two bundles may hold even when one component of the revealed preferred bundle had a smaller amount than the alternative bundle (physical non-dominance, requiring overcompensation by the other bundle component) (Figure 1E, arrows). This aspect, together with the equal-preference trade-off, constituted a crucial characteristic for the integration of the physical values of both bundle components into single-dimensional preference relationships and utility.

Human participants. A total of 24 human participants (11 males, 13 females; mean age 25.4 years, range 19-36 years) completed a binary choice task for measuring revealed preferences and performed a Becker-DeGroot-Marschak (BDM) control task. None of the participants had diabetes or lactose intolerance, nor did they require specific diets, to avoid medical and cultural interference. All participants had a known appetite for milkshakes and provided written informed consent based on a detailed information sheet. The Local Research Ethics Committee of the Cambridgeshire Health Authority approved the study.

Stimuli, rewards and delivery apparatus. The human participants viewed quantitative, colored visual stimuli on a computer monitor, which represented the two milkshakes and their amounts in each of the two bundles (Figure 1A). Each bundle stimulus consisted of two vertically aligned rectangles. The color of each rectangle indicated the bundle component. The vertical position of a bar in each rectangle indicated the physical amount of each component (higher was more).

After extensive piloting with various liquids and liquidized foods, we found milkshakes with a controlled mix of sugar and fat to give the most reliable behavioral performance. As the milkshakes were delivered separately, with a 500 ms interval (see below), drinks containing either only sugar or only fat were deemed to be too unnatural. Thus, in our bundles, component A (top, blue) consisted of a low-sugar high-fat milkshake (25% whole milk, 75% double cream, no sugar), and component B (bottom, red) consisted of a high-sugar low-fat milkshake (10% sugar in skimmed milk). The Psychtoolbox in Matlab (Version R2015b) running on a Dell Windows computer served for stimulus display and recording of behavioral choices.

The two milkshakes were delivered to the participants using a custom-made syringe pump system. The two milkshakes were contained in two 50-ml syringes, respectively, attached to two piston pumps (NE-500, New Era Pump Systems Inc; www.syringepump.com). Each pump

delivered programmable quantities of milkshake with milliliter precision via a silicone tube approved for delivery of food stuffs (VWR International Ltd). Each pump was controlled by a Dell Windows computer using a National Instruments card (NI-USB-6009) via the Matlab Data Acquisition Toolbox.

Behavioral task. Each trial started with an initial fixation cross in the center of a computer monitor in front of the participant. After a pseudorandomly varying interval (mean 0.5 s), the two bundle stimuli appeared simultaneously at pseudorandomly alternating fixed left and right positions on the monitor; each bundle stimulus contained the same two milkshakes with independently set amounts (Figure 1A). The participant chose one of the two bundle stimuli by pressing a single button once (left or right computer keyboard arrow for corresponding choice of left or right bundle). The two milkshakes from the chosen bundle were delivered together with a probability of $P = 0.2$; i. e. every fifth chosen bundle was paid out on average using a Poisson distribution, and there was no payout of any milkshake on the remaining trials. Component B was delivered after a constant interval of 500 ms after component A. This constant delay, rather than simultaneous delivery or pseudo-randomly alternating sequential liquid delivery, prevented uncontrolled liquid interactions, maintained discriminability of the liquids and amounted to constant temporal discounting. Thus, the utility for component B reflected the subjective value of the milkshake itself and a temporal discount due to longer delay.

Although participants were instructed to not eat or drink up to four hours prior to the testing, satiety was a concern due to the high sugar and fat content of the milkshakes. We addressed the issue by the $P = 0.2$ payout schedule, by limiting each payout to maximally 10.0 ml, and by delivering not more than 200 ml of total liquid to each participant on each session.

Psychophysical assessments. We used a standard psychophysical staircase procedure (Green & Swets 1966) to estimate the IP at which each of the two bundles were chosen with equal probability ($P = 0.5$ each option), revealing equal preference for each option. The procedure required repeated testing, which was also required for the subsequent neuroimaging experiment with the same participants (to be reported elsewhere).

We started the psychophysical procedure by setting, in the Reference Bundle, component A to 0.0 ml and component B to 2.0 ml, 5.0 ml or 8.0 ml (Figure 1B). We then varied the alternative Variable Bundle in a systematic fashion; first, we set its component A one unit higher (mostly 0.5 ml, 1.0 ml or 2.0 ml higher), thus specifying the amount of component A gained by the participant from the choice; then we selected randomly (without replacement) one amount from a mean of 7 fixed amounts of component B (multiples of 0.5 ml), spanning the whole, constant range of tested amounts; we repeated the selection until all 7 amounts had been tested once (Figure 1C). Then we estimated each IP from 6 repetitions using sigmoid fitting (see Eqs. 1, 2 below), requiring 42 choices per IP. At the IP, the amount of component B was usually lower in the Variable Bundle compared to the Reference Bundle. In this way, we assessed at each IP how much of component B the participant was willing to give up in order to gain one unit of component A, always relative to the constant Reference Bundle.

We obtained further IPs in choices between the constant Reference Bundle and the Variable Bundle whose amounts of component A increased step-wise, thus advancing from top left to bottom right on the two-dimensional x-y indifference map. We are aware that the unidirectional progression of testing may lead to somewhat different IP estimates than testing in the opposite direction or in random sequence (Knetsch 1989). However, in this initial study, we were primarily interested in the systematic assessment of consistent IPs rather than exploring potential pitfalls.

To obtain three levels of revealed preference, we used three starting amounts of component B in the Reference Bundle (2.0 ml, 5.0 ml, or 8.0 ml). In total, we estimated 4 IPs (from 5 test bundles) at each of 3 levels of revealed preference, resulting in a total of 12 IPs, derived from 504 choices among 84 different option sets in each participant (7 psychophysically tested amounts for each of 12 IPs; 6 repetitions).

Statistical analysis. In order to estimate the choice IPs numerically, we obtained a sigmoid fit to the empirically assessed choice frequencies via a general linear regression. To do so, we used the

Matlab function *glmfit* (Matlab R2015b; Mathworks) on a binomial distribution with a probit link function, which is the inverse of the normal cumulative distribution function (G). Specifically, the generalised linear regression $y = \beta_0 + \beta_1 B_{var} + \varepsilon$ can be rewritten after applying the link function as:

$$G(y) = \beta_0 + \beta_1 B_{var} + \varepsilon \qquad \text{Eq. 1}$$

with y as number of times the subject chose the Variable Bundle in the current block from a series of six repetitions, β_0 as offset constant, β_1 as regression slope coefficient, B_{var} as reward amount (ml) of component B in the Variable Bundle, and ε as residual error. We choose the probit model because it assumes that random errors have a multivariate normal distribution, which makes it attractive as the normal distribution provides a good approximation to many other distributions. The model does not rely on the assumption of error independence and is used frequently by econometricians (Razzaghi 2013). By comparison, the logit model is computationally simpler but has more restrictive assumptions of error independence. As preliminary data had revealed a similar fit for the logit as for the probit model we used the probit model due to its less restrictive assumptions. Thus, we estimated the IPs from the sigmoid fit provided by the probit model, using the following equation:

$$\text{Indifference Point} = -(\beta_0 / \beta_1) \qquad \text{Eq. 2}$$

with β_0 and β_1 as coefficients of the general linear regression (Eq. 1). By using a probit instead of a logit function, we obtained the coefficients from the output of the probit analysis (Amemiya 1981).

We obtained single ICs, separately for each individual participant, from a set of individual IPs by weighted least-mean-square, non-linear regression (as opposed to the probit regression for estimating each IP). We applied a weight in order to account for within-participant choice variability; the weight was the inverse of the standard deviation of the titrated amount of the B-component at the corresponding IP (the IP having been estimated by the probit regression analysis). We estimated the best-fitting β coefficients from least-mean-square fitting to obtain the equal-revealed-preference IC (which is a utility level) and wrote the basic hyperbolic equation in our notation as:

$$IC = \beta_0 + \beta_1 B + \beta_2 A + \beta_3 BA + \varepsilon \qquad \text{Eq. 3}$$

with A and B as amounts of components A and B (ml) referring to the x-axis and y-axis, respectively, β_1 and β_2 as slope coefficients of the least-mean-square, non-linear regression, and β_3 as curvature coefficients. As IC is a constant, we merged the other constants offset (β_0) and error (ε) into a common final constant k. To draw the ICs, we computed the amount of component B as a function of component A from the derived equation:

$$B = k - (\beta_2/\beta_1)A + \beta_3 A \qquad \text{Eq. 3a}$$

To graphically display a fitted IC (Figure 2A, B), we plotted the preset amount of component A on the x-axis, and the computed fitted amount of component B (Eq. 3a) on the y-axis. The error on the hyperbolic curve was measured as 95% confidence interval. The higher the error around an IP the less weight was given to this point when the IC was calculated. This model resulted in good fits in earlier work (Pastor-Bernier et al. 2017). In this way, the IPs of 5 equally revealed preferred but differently composed bundles aligned as a single fitted IC. The three ICs representing increasing revealed preference levels (low, medium, high) were located increasingly farther away from the origin (Figures 1E; 2A, B). The indifference map of 3 x 5 IPs was unique for each participant (Figure 3).

The IC shape is derived from hyperbolic fit and is quantified by two coefficients: slope and curvature. The IC slope coefficient, derived from the ratio of regression slope coefficients (β_2 / β_1), reflects the currency relationship between the components and describes the participant's preference for component A relative to component B. For example, an IC slope of - 60^0 indicates that

component A is valued twice as much as the same ml amount of component B. The curvature coefficient (β_3) quantifies the constancy in the trade-off between bundle components. A linear curve (curvature coefficient = 0) indicates a constant rate of exchange for the bundle components, indicating that the components are perfect substitutes. A more convex IC (curvature coefficient > 0) indicates a varying rate of exchange, suggesting that the participant is giving up lesser amounts of component B to obtain one unit of component A when advancing on the IC from top left to bottom right. For a more intuitive measure, we quantified the curvature by measuring the largest perpendicular distance between the IC and the line between the x-axis and y-axis intercepts (Figure 2E):

$$d = \max (B_{IC} - B_{linearIC}) \qquad \text{Eq. 4}$$

with d as maximal perpendicular distance (ml) (whereas β_3 is a best-fitted, estimated parameter, and thus less conservative), B_{IC} as amount of component B on the IC (ml), and $B_{linearIC}$ as amount of component B at the line connecting the x- and y-intercepts (constant amount of component A, x-axis; ml). This simplified curvature measure reflects the change in trade-off between the two components across the tested range of reward amounts, in ml of component B.

We used logistic regression on trial-by-trial choices to confirm that the measured choices were explained by bundle components rather than other factors. In a random-effects analysis, we fitted a logistic regression to the data from each individual participant and then averaged the obtained β coefficients and p-values across all participants. We used the following regression:

$$y = \beta_0 + \beta_1 RefB + \beta_2 VarA + \beta_3 VarB + \beta_4 RT + \beta_5 VarPos + \beta_6 PChoice + \varepsilon \qquad \text{Eq. 5}$$

with y as either 0 or 1 (not-choice or choice of Variable Bundle), A and B as amounts of bundle components A and B (ml), RefB as amount of component B in the Reference Bundle (ml), VarA and VarB as amount of components A and B in the Variable Bundle (ml), RT as reaction time (ms; interval between appearance of the two bundle stimuli and participant's key press), VarPos as left or right position of the Variable Bundle on the computer monitor relative to the Reference Bundle (0 or 1), and PChoice as choice in the previous trial (0 or 1) (Figure 2G). Each β coefficient was standardised by multiplication with standard deviation of the respective independent variable and division by standard deviation of y (dependent variable). A subsequent one-sample t-test against 0 served to assess the significance of the beta (β) coefficients in the population of the 24 participants.

Satiety may have occurred and could have affected the preferences for the two bundle components in an uncontrolled manner, even though the bundle rewards were only paid out on every fifth trial on average and were limited to a total of 200 ml. A prime suspected effect might have been a differential devaluation between the two bundle components that would result in changed currency relationship between the two components. Such change between the two components should be manifested as gradual change in instantaneous choice probability near the IPs over repeated test steps of 42 trials each (Figure S1). We calculated the instantaneous choice probability as:

$$y = \sum (n=1 \text{ to } 6) (CV / TS) \qquad \text{Eq. 6}$$

with y as instantaneous probability (P = 0.0 to 1.0), CV as not-choice or choice of Variable Bundle (0 or 1), and TS as test step (1-6).

We used a leave-one-out analysis to assess the meaningful representation of revealed preferences by the fitted ICs and to test the accuracy of the hyperbolic IC fit to the IPs. In this analysis, we removed one IP per IC (but not the initial Reference Bundle set at x = 0), fitted the curve again with the hyperbolic model and assessed the deviation between the original IC and the IC with the left-out IP. We defined the deviation as the difference of component B between the original, left-out IP and the refitted IC (Figure 4B):

$$d = B_{IP} - B_{refit} \qquad \text{Eq. 7}$$

with d as difference (in ml; y-axis), B_{IP} as amount of component B of the left-out IP (ml), and B_{refit} as amount of component B on the refitted IC (ml). Thus, a difference of 0 ml suggested that removal of one IP did not affect the shape of the IC at all, whereas any difference unequal to 0 ml quantified the violation of this assumption.

Decoding of preference levels. To confirm the contribution of each IP to the two-dimensional representation of revealed preferences, we determined the accuracy (in percent correct) with which a randomly selected bundle, defined by the amounts of the two components A and B (in ml), could be assigned to its original revealed preference level as opposed to any one other level (binary distinction). By definition, each bundle that was psychophysically estimated to be as much revealed preferred as the Reference bundle constituted an IP; all bundles to which participants were choice indifferent against the same Reference Bundle constituted a series of IPs. In our experiment, three different Reference Bundles defined three preference levels (low, medium, high: component B: 2.0 ml, 5.0 ml or 8.0 ml, respectively; component A was always 0.0 ml; Figure 5A). The decoder used as inputs only bundles at the psychophysically estimated IPs (to which an IC was fitted using Eqs. 3, 3a), rather than bundles positioned on the fitted ICs.

Our main test employed a binary support vector machine (SVM) decoder separately on each individual participant. We used similar methods as previously described for predicting choice from neuronal activity (Tsutsui et al. 2016). The SVM algorithm considered 5 IP bundles from each of 2 revealed preference levels (total of 10 IPs that had been assessed 12 times at each position in each participant) (Figure 5A). Each of the 2 preference levels was associated with a matrix of 2 columns (containing the x- and y-coordinates of bundle components A and B, respectively) and 5 rows (containing the 5 bundles). The 5 bundles were randomly selected (with replacement) from 60 bundles on each level (due to the random procedure with replacement, some bundles may have entered the algorithm multiple times, and not all five bundles may have been used for a given analysis). We left out 1 randomly selected bundle from the 10 bundles, trained the SVM algorithm with the remaining 9 bundles, and assessed whether the SVM decoder assigned the left-out bundle to its original revealed preference level or to another level. Thus we used 90% of the data for training the decoder and 10% for testing its classification performance. We repeated this procedure 10 times with the same selected 2 x 5 bundles but with a new randomly selected left-out bundle and calculated decoder accuracy as percent correct classification in these 10 trials. We repeated the random selection of the 2 x 5 bundles and the 10-trial accuracy assessment 150 times. For final decoding accuracy, we averaged the percentages from these 150 iterations (Table 1). We applied this procedure separately to all three possible combinations of two revealed preference levels (i. e. low and medium, medium and high, low and high). The SVM was implemented with custom written software in Matlab R2015b (Mathworks) using the functions *svmtrain* and *svmclassify* with linear kernel (our previous work had shown that use of nonlinear kernels did not improve decoder performance; Tsutsui et al. 2016).

We supplemented the SVM procedure with binary linear discriminant analysis (LDA) that provided visualization of the different levels of revealed preference (Figure 5). We used the same IPs and the same data matrices as for the SVM analysis (and the same IPs as used for the hyperbolic fitting of the three indifference curves, ICs). We obtained two variances; the discriminant 1 eigenvector captured the best separation between the three revealed preference levels as 'across-level variance' (colors in Figure 5); the discriminant 2 eigenvector captured the best within-level separation between five bundles on each of the three preference levels as 'within-level variance' (symbols in Figure 5). The results indicate visually the discrimination accuracy on the two axes of the two-dimensional plots. We also assessed the numeric accuracy of decoding as percent of correctly assigning a randomly selected bundle to its original revealed preference level. The decoder used the Matlab functions *fitcdiscr* and *predic* on z-normalised data from individual participants. For the LDA, our limited data required pooling from multiple participants. As revealed preferences are private and subjective, and therefore difficult to compare between individual participants, the LDA results should be considered as merely supportive and not as stand-alone data.

Mechanism-independent validation. To relate the estimated revealed preferences to inferred utility, we implemented a Becker-DeGroot-Marschak (BDM) mechanism akin to a second price

auction (Becker et al. 1964). The BDM value bids were then compared with the three levels of revealed preference within each participant.

In the BDM, the participant received a fresh monetary endowment (20 UK pence) on each trial. The participant bid for a bundle against a pseudorandomly set computer bid (retrieved from a normal distribution with replacement). If the participant's bid was higher than or equal to the computer bid then she received both component rewards of the bundle and paid an amount equal to the computer bid. If her bid was lower than the computer bid, she lost the auction, paid nothing and did not receive any bundle reward. The participant was informed about a win or a loss immediately after placing the bid; when winning the bid, the participant received the bundle rewards in the same sequence and frequency (every fifth trial on average) as in the choice task assessing revealed preferences.

We showed each participant single bundles that were randomly selected (without replacement) from a set of 15 bundles (5 equally revealed preferred bundles at each of 3 levels; the 15 bundles had been used to fit the 3 ICs shown in Figures 2A, B and 3). A given bundle was set to the participant's psychophysically estimated IP (Figure 6A). We presented each of the 15 bundles 12 times, resulting in 180 trials in total, and considered the mean of these bids as the BDM-estimated utility. The participant indicated her bid by moving a cursor horizontally on the computer monitor with left and right keyboard arrows (Figure 6A inset). The BDM bid was registered from the cursor position at 5.0 s after onset of presentation of the horizontal bidding scale.

We first assessed the basic question whether the monetary bids increased for higher valued bundles but were similar for equally valued bundles (which constituted IPs), using Spearman rank correlation analysis confirmed by one-way Anova with post-hoc Tukey-Kramer test. Then we performed a random-effects analysis with a general linear regression with a normal (Gaussian) link function on separate data from each participant and averaged the obtained β coefficients and their p-values across participants. We used the following regression:

$$y = \beta_0 + \beta_1 \text{PrefLev} + \beta_2 \text{AmBundle} + \beta_3 \text{TrialN} + \beta_4 \text{PrevBid} + \beta_5 \text{Consum} + \varepsilon \quad \text{Eq. 8}$$

with y as monetary bid, PrefLev as revealed preference level (low, medium, high), AmBundle as summed ml amount of components A and B in the currency of component A ($A + (k - (\beta_2/\beta_1)A + \beta_3 A)$ as in Eq. 3a), TrialN as trial number, PrevBid as BDM bid in previous trial (UK pence), and Consum as accumulated drinks consumption for component A and component B up to this point in the experiment (ml) (Figure 6C). Each β coefficient was standardised by multiplication with standard deviation of the respective independent variable and division by standard deviation of y (dependent variable). A subsequent one-sample t-test against 0 assessed the significance of the beta (β) coefficients in all 24 participants.

To assess the internal consistency of BDM bids, we used binary SVM analysis on bids from individual participants in analogy to SVM decoding of preference levels. We tested the same IPs as used for hyperbolic IC fitting (Figure 6A). Each of the 2 preference levels was associated with a matrix of 1 column (containing the bids to each bundle) and 5 rows (containing the 5 bundles). The remainder of the bundle selection, leave-out and repetition procedure was identical to that used for the SVM decoding of preference levels (see above). Thus, the SVM decoder for BDM bids assessed the accuracy with which the left-out bundle belonged to its original revealed preference level. We supplemented the SVM analysis of the BDM bids with analogous LDA for supportive visualization.

Finally, we compared hyperbolically fitted BDM isolines directly with hyperbolically fitted revealed preference ICs (rather than with revealed preference levels just described), separately for each individual participant. This procedure required to present BDM bids on the same scale as revealed preference ICs. To this end, we fitted isolines of same BDM-bids in analogy to fitting same-preference ICs. We fitted a hyperbolic function to the measured mean BDM bids in analogy to Eq. 3:

$$\text{BDMBid} = \beta_0 + \beta_1 B + \beta_2 A + \beta_3 BA + \varepsilon \quad \text{Eq. 9}$$

with β_1 and β_2 as regression slopes, and β_3 as curvature coefficients, and A and B as amounts of components A and B (ml), respectively. Coefficients β_1 and β_2 were standardised by multiplication

with standard deviation of components B and A, respectively (independent variables), and division by standard deviation of BDMBid (dependent variable). We obtained separate β coefficients from all participants and averaged them and their p-values across participants. A subsequent one-sample t-test used the individual beta (β) coefficients from all 24 participants to test overall significance against 0.

To compare BDM bids with ICs, we graphically displayed BDM isolines along which all mean BDM bids were equal. As a BDM isoline is a constant, we merged the constants offset (β_0) and error (ε) into a common final constant k. To draw the BDM isolines, we computed the amount of component B as a function of component A from the derived equation:

$$B = k - (\beta_2/\beta_1)A + \beta_3 A \qquad \text{Eq. 9a}$$

To display a three-dimensional map, we graphed colored BDM isoline zones on the z-axis as a function of the amounts of components A (x-axis) and B (y-axis) (Figure 7A). For a two-dimensional map of BDM isolines, we plotted the preset amount of component A on the x-axis and the amount of component B computed from the isolines (Eq. 9a) on the y-axis (Figure 7B). For comparison, we plotted the revealed-preference ICs on the same two-dimensional map using the same scale. We also compared numerically, separately for each participant, CIs and slope and curvature coefficients between hyperbolically fitted BDMBids (Eq. 9a) and hyperbolically fitted revealed preference ICs (Eq. 3a), using the paired Wilcoxon test (Figure 7C-E).

Comparison with monkeys. In order to compare the results of this human study across a closely related species, we re-analyzed existing data from a previous experiment on rhesus monkeys; all methods of this study have been described (Pastor-Bernier et al. 2017). That study tested in an analogous manner the same notions of Revealed Preference Theory as stated above. Each monkey chose with a hand movement onto a touch monitor between two visually presented bundles composed of the same two liquids (fruit juices or water) in varying amounts. We used analogous psychophysical procedures and statistics to estimate IPs and ICs. For the current comparisons, we assessed the accuracy of integration of the two option components by the monkeys with two measures (Figure 8): (1) the 95% CIs of the psychophysical fits to the choice probabilities used for estimating each IP, which indicated how well the animals had estimated the IPs (Figure 2A inset); (2) the CIs of the hyperbolic fits of ICs to all equally preferred IPs, which indicated the accuracy of the trade-off between the two bundle components that characterizes the value integration from both bundle components (Figures 2A, B; 3).

Results

Performance assessment with graphic indifference curves (IC). The human participants repeatedly chose between two composite visual stimuli that represented two-component bundles (see Methods notion 1; Figure 1A). Each bundle stimulus consisted of two colored vertical rectangles; blue indicated component A, a low-sugar high-fat milkshake; red indicated component B, a high-sugar low-fat milkshake; a vertically positioned bar inside each rectangle indicated the amount of each milkshake (higher was more). We assessed choices between a preset Reference bundle and a Variable bundle whose component A was set to a specific test amount and whose component B varied pseudorandomly across a wide range (Figure 1B). The choice probabilities in all 24 participants followed the change of component B in an orderly fashion, thus revealing monotonic preference relationships (notions 2 - 4; Figure 1C). We estimated the indifference point (IP; $P = 0.5$ for each bundle) from six repetitions using the probit choice function (Eqs. 1, 2); each IP required 42 choices. Thus, a two-dimensional IP indicated the amounts of the two components of the Variable Bundle that was as much revealed preferred, and had the same utility, as the constant Reference Bundle. To obtain a series of IPs, we repeated this procedure by setting increasing amounts of component A in the Variable Bundle while keeping the Reference Bundle constant. All IPs of such a series were as much revealed preferred as the constant Reference Bundle and, by transitivity, equally revealed preferred to each other.

We obtained single ICs from such a series of five equally revealed preferred IPs by hyperbolic fitting (Eqs. 3, 3a) (notions 5 and 6; Figure 1D, E). Such an IC defined the trade-off between the two components of an equally preferred bundle: it indicated how much of component B a participant gave up for obtaining one more unit of component A without a change in utility. Thus, the IC characterised the orderly integration of both bundle components into a single-dimensional estimate. The continuous ICs were asymmetric between the x-axis and y-axis, indicating different subjective weighting of the two milkshakes; the convex IC curvature suggests that lower amounts of both milkshakes together were as much preferred as higher amounts of singular milkshakes (possibly reflecting gradually flattening, concave utility functions and/or complementarity between high-sugar and high-fat components).

We estimated three sets of 5 IPs with their respective, hyperbolically fitted ICs by presetting the Reference Bundle to three different amounts of component B (2.0 ml, 5.0 ml or 8.0 ml; component A being always 0.0 ml). Figure 2A (inset) shows the psychophysical assessment of an example IP. Although the ICs varied in slope and curvature between participants (for two example participants, see Figure 2A, B; for all participants, see Figure 3), the ICs of bundles with larger reward amounts were located farther away from the origin, thus representing revealed preference for larger rewards (notion 7). The three ICs were well ordered and failed to overlap. Only the 95% confidence intervals (CI) overlapped partly in four of the 24 participants (17%). Closer inspection of Figures 2A, B and 3 shows that some bundles on higher ICs had a smaller amount of component B than bundles on lower ICs, indicating revealed preference despite partial physical non-dominance (notion 8; involving overcompensation by higher amount of component A). As with the trade-off, such preferences for bundles with one physically lower component indicate the integration of the values of both bundle components into single-dimensional utility.

Taken together, the maps of systematic and continuous ICs reflect the decision process that results in single-dimensional preferences for multi-component, vectorial choice options. These results very much correspond to, and thus validate empirically, the intuitive schemes of Revealed Preference Theory. The following tests will address these validations more quantitatively.

IC coefficients. The shape of ICs reflects the trade-off between the components and can be quantified by slope and curvature coefficients of hyperbolic fits (Eq. 3a). The global IC slope, between y-axis and x-axis intercepts, was measured as ratio of the two regression coefficients β_2 / β_1 in Eq. 3; it indicated how much the participant was globally willing to give up in order to obtain one unit of the other component; the measure reflected the relative utility (currency) of the two bundle components. The IC slopes steeper than -45° indicate that a participant gave up a higher amount of component B (high-sugar, y-axis) for a smaller amount of component A (high-fat, x-axis), thus indicating higher subjective value of fatty than sugary milkshake. The IC slope was -71° ± 6.5° (mean ± standard error of the mean, SEM; range -45° to -76°; $N = 24$ participants; Figure 2C). The higher valuation of high-fat component A over high-sugar component B amounted to a factor of 3:1 in 18 of the 24 participants (75%). The predominantly asymmetric trade-off between the two milkshake components documents that each component contributed to bundle preference in its own distinct way.

The IC curvature showed substantial convexity in 18 of the 24 participants (75%), as indicated by β_3 coefficients from the hyperbolic fit (Eq. 3) that were significantly larger than 0.0 (8.89; mean ± 5.9 SEM; $p<0.05$, t-test; Figure 2D). For graphically assessing IC curvature, we measured the distance between the IC center and a straight line connecting equally revealed preferred bundles at the x-and y-intercepts, in units of ml on the y-axis (Figure 2E). This distance ranged from 0.09 ml (quasi-linear IC) to 3.76 ml (most convex IC) (mean of 1.28 ml ± 0.19 SEM; Figure 2F). The distribution of the IC distance was overall similar to that of the β_3 curvature coefficient from Eq. 3. The two highest histogram bars in Figure 2F show data from six participants with rather similar, considerably convex IC, and from six other participants with rather similar but quasi-linear IC. Thus, the coefficients confirmed numerically the well-ordered nature of the representation of revealed preferences by the ICs.

Control for other choice variables. To test whether the choices reflected the components of the bundles rather than other, unrelated factors, we performed a logistic regression analysis separately

for each of the 24 participants, using the following regressors: amount of each bundle component, reaction time, Reference Bundle position on participant's monitor, and previous trial choice (Eq. 5). The standardised beta (β) coefficients and p-values were assessed for each participant and then averaged across all 24 participants; they demonstrated that the choice of the Variable Bundle was negatively correlated with the amount of component B in the Reference Bundle (RefB: β = -0.43 ± 0.16, mean ± SEM; P = 0.020 ± 0.005) (component A was constant 0.0 ml, see Methods) and positively correlated with both components A and B in the Variable Bundle (VarA: β = 0.67 ± 0.16; P = 0.009 ± 0.004; VarB: β = 0.94 ± 0.33; P = 0.012 ± 0.009) (Figure 2G). The beta (β) coefficients for these three variables differed significantly from 0 (P = 0.012, P = 0.00088 and P = 0.00028, respectively; one-sample t-test), confirming the validity of the β's. Thus, more frequent choice of the Variable Bundle correlated with lower amounts in the Reference Bundle and with higher amounts of either component in the Variable Bundle. The result suggests that both bundle components, rather than a single component alone, were important for the observed preference relationships. The remaining variables, including reaction time, position of Reference Bundle on the monitor and previous trial choice, failed to significantly account for current choice of the Variable Bundle (P = 0.754 – 0.988 ± 0.003 – 0.290). Thus, the revealed preference relationships concerned the bundles with their two components rather than other task factors.

To assess potential consumption effects, we searched for signs of satiety. We followed choice probability across the total test duration in each of the 24 participants. We selected two bundles that were situated above and below the IP, respectively. These two bundles contained high-sugar low-fat (above IP) and low-sugar high-fat (below IP) milkshakes. We plotted choice probability over six repeated test steps (Figure S1). Choice probabilities fluctuated only insignificantly, without conspicuous upward or downward trend (above IP: F (5, 41) = 0.28, $P > 0.05$; below IP: F (5, 41) = 1.53, $P > 0.05$; 1-way repeated measures Anova with post-hoc Tukey Test). Even at the final, sixth step, choice probability differed only insignificantly from any other step. Thus, the revealed references did not seem to be importantly confounded by satiety for neither sugar nor fat within the amounts and concentrations used in our experiment.

Internal validation. We performed a leave-one-out analysis to assess the contribution of individual bundles to the ICs obtained from hyperbolic fits to the empirically estimated IPs in humans. We removed one IP at a time (except the initial Reference Bundle at x = 0.0) from the set of five IPs to which a given IC had been fitted, and then refitted the IC with the remaining four IPs using the same hyperbolic model (see Methods, Eqs. 3, 3a). This was done separately for each of the three ICs, separately in all 24 participants (total of 288 IPs removed from 72 ICs). We found consistency in the refitted ICs in four measures (Figure 4A). First, none of the 72 refitted ICs overlapped with the refitted ICs at different levels in the same participant, thus demonstrating maintained IC separation despite one left-out IP. Second, none of the 72 refitted ICs overlapped with the 95% CIs of original ICs at different levels, confirming IC separation despite one left-out IP. Third, most refitted ICs (66 of 72 ICs, 92%) fell inside the 95% CIs of the original ICs, and the remainder curves (6 of 72 ICs, 8%) showed only some portions outside the 95% CIs of the original ICs, thus refuting possible overweighted influence of individual IPs on ICs. Fourth, numeric comparisons showed only insignificant deviations between refitted ICs and the IPs that had been left out when refitting the curves (vertical distance of 0.05 ± 0.13 ml in all 24 participants; mean ± standard error of the mean, SEM; N = 336; P = 0.98 against normal distribution; t-test) (Figure 4B, C), confirming absence of overweighted IP influence on ICs. These four results suggest that the hyperbolically fitted ICs captured the IPs consistently and provided valid representations of the revealed preferences.

We used a linear support vector machine (SVM) decoder as different statistical procedure to confirm the contribution of each IP to the two-dimensional representation of revealed preferences by ICs. In each participant, we set a given test bundle to one of the psychophysically determined IPs (Figure 5A) and assessed the accuracy with which the decoder assigned that bundle to its original preference level (each preference level was defined by a series of empirically estimated IPs but was not a fitted IC). SVM decoding accuracy ranged largely from 70% to 100% ($P = 2.055 \times 10^{-101}$), although a few lower values were observed (Table 1 left).

We supplemented the SVM analysis by visualization of decoding using linear discriminant analysis (LDA). The considerable amount of data necessary for reasonable LDA required us to pool data from several participants, which violates a basic tenet of economic theory that prohibits pooling of subjective preferences across individual participants. To somewhat contain expected inaccuracies, we normalised IPs across participants (z-score normalization for reward B along the y-axis; reward A had been set experimentally to identical values on the x-axis) and restricted the analysis to specific subsets of participants. The LDA confirmed the SVM results in all subsets (Figure 5); the first linear discriminant assigned bundles to the three revealed preference levels, as shown by spatial separation of the three colored groups, with a numeric accuracy of 80-100% ($P = 1.148 \times 10^{-97}$). By contrast, the second discriminant failed to accurately assign bundles to different positions on same preference levels, as shown by the mix of the five shapes representing bundle position. These characteristics were seen in six participants whose fitted ICs showed the highest similarity in convexity (Figure 5B, C), in six participants with linear ICs (Figure 5D, E) and, for comparison, in all 24 participants (Figure 5F, G) (for distinction of participants based on IC curvature, see two highest bars in Figure 2F). Thus, LDA decoding followed the fundamentals of ICs: preference for bundles on higher ICs but indifference along ICs.

Taken together, the two decoders confirmed three distinguishable levels of IPs, and LDA in addition confirmed indifference between IPs on same levels. As these IPs constituted the basis for hyperbolic fitting of the three ICs, the decoder results validated also the fitting procedure and confirmed that the empirically estimated ICs represented well the revealed preferences.

Mechanism-independent validation. The ordered representation of revealed preferences by our ICs can be further validated by comparison of the key features of ICs with subjective values inferred from a different estimation mechanism. To this end, we used a monetary Becker-DeGroot-Marschak (BDM) bidding task that is equivalent to a second-price auction and estimates participants' truthful, 'incentive compatible' bids on a trial-by-trial basis. The property of truthful revelation makes BDM an indispensable tool of experimental economics and explains its increasing popularity in neuroscientific studies of human decision making (Plassmann, O'Doherty & Rangel 2007; Medic et al. 2014).

For these comparisons, the same 24 participants made monetary BDM bids (UK pence) for 15 bundles that had been set, individually in each participant, to five psychophysically estimated IPs on each of the three preference levels (in ml of each of the two component milkshakes). Each preference level was defined by IP bundles that were equally revealed preferred to the initial Reference Bundle with component A set to 0.0 ml (Figure 6A; the same 15 bundles were used for hyperbolic IC fitting, but the preference levels used for this analysis were defined by the IPs and not the fitted ICs). Although BDM is typically performed in single shots, we obtained average BDM bids for each bundle from 12 repetitions, thus aiming to approach the nature of data collection during stochastic choice with psychophysical estimation of each IP (42 trials).

The BDM bids followed the ordered preference levels, as shown by significant positive correlation between the bids for same-level bundles (means from all five bundles) and the three levels (Spearman rank correlation: rho = 0.60 ± 0.05; mean ± SEM; $N = 24$ participants; $P < 0.01$; confirmed by one-way Anova with post-hoc Tukey-Kramer test between all bundles across the three levels: $P < 0.01$) (Figure 6B; red, green, blue). A random-effects analysis separately for each of the 24 participants (Eq. 8) demonstrated a relationship of BDM bids to preference level (PrefLev: $\beta = 0.47 \pm 0.09$, mean across all 24 participants ± SEM; $P = 0.016 \pm 0.015$; β-coefficient difference from 0: $P = 0.000026$, one-sample t-test) and bundle amount (AmBundle: $\beta = 0.15 \pm 0.13$; $P = 0.020 \pm 0.017$; $P = 0.0278$), rather than trial number (TrialN: $\beta = -0.10 \pm 0.25$; $P = 0.726 \pm 0.354$), previous trial bid (PrevBid: $\beta = 0.12 \pm 0.11$; $P = 0.676 \pm 0.427$) or consumption history (Consum: $\beta = 0.12 \pm 0.11$; $P = 0.224 \pm 0.185$) (all β's were standardised) (Figure 6C). A specific analysis (Eq. 9) demonstrated that the BDM bids reflected both bundle components (component A: $\beta = 0.6534 \pm 0.0866$, mean ± SEM; $P = 0.0324 \pm 0.0150$; β difference from 0: $P = 1.1613\ 0884 \times 10^{-7}$, one-sample t-test; component B: $\beta = 0.6425 \pm 0.0585$, $P = 0.0289 \pm 0.0202$; β difference from 0: $P = 1.2770 \times 10^{-10}$). Thus, the BDM bids followed well the revealed preference levels and took both bundle components into account.

The BDM bids were similar for equally revealed preferred bundles. To comply with the notion of trade-off between different bundle components, equally revealed preferred bundles should elicit similar BDM bids despite varying bundle composition. We conducted a one-way Anova on all BDM bids for bundles positioned on the same preference level, separately for each level and each participant. The bids for bundles located on the same preference level failed to differ statistically (low level: $P = 0.54 \pm 0.06$; medium level: $P = 0.69 \pm 0.06$; high level: $P = 0.78 \pm 0.05$; mean ± SEM, $N = 24$ participants). Only three participants showed significantly different bids for bundles on the medium preference level, and one participant showed different bids for bundles on the high preference level ($P < 0.05$). Thus, most participants showed similar BDM bids for equally revealed preferred bundles on same revealed preference levels (while being sensitive to bundle amounts, as shown by increasing bids across levels). The BDM bids confirmed the equal-preference trade-off between bundle components independent of the binary choice mechanism used for assessing revealed preferences.

The BDM bids were internally consistent. We used decoders to test the distinction of bundle position across but not along preference levels. Using BDM bids, a binary SVM decoder showed good accuracy of assigning a test bundle to its original preference level in individual participants (mostly 50-70%; $P = 3.789 \times 10^{-9}$) but failed to distinguish bundles along same levels (43-51%; $P = 0.1433$) (Table 1 right). We used an LDA, which is able to decode several levels together, as supplementary test and found good visual assignment of the test bundles (Figure 6A) to the three revealed preference levels in our population of 24 participants (first discriminant; numeric accuracy of 88-100%; $P = 9.46 \times 10^{-12}$; Figure 6D; three colored symbol groups) but not to different positions on same preference levels (second discriminant; numeric accuracy of 43-51%; mix of shapes) (note the reservations above when combining data from multiple participants). Thus, the BDM bids followed well the two-dimensional scheme of revealed preference represented by the ICs.

The BDM bids matched well the revealed preference ICs in direct comparison. A stronger mechanism-independent validation of the IC scheme may be achieved by direct graphic comparison between BDM bids and hyperbolically fitted ICs. To this end, we estimated isolines that were fit to BDM bids using Eqs. 9, 9a and compared them with ICs that had been hyperbolically fitted to IPs (Eqs. 3, 3a). The BDM isolines represented BDM bids from 12 averages, and the revealed preference ICs derived from fits to IPs estimated from 42 stochastic choices. We found that the BDM isolines increased for IPs on increasing preference levels (farther away from the origin), but were similar for IPs on the same revealed preference level (same distance from origin) (Figure 7A). The BDM isolines matched the revealed preference ICs within their 95% CIs in every one of the 24 participants (Figure 7B). Statistical comparisons showed significantly higher CIs of BDM isolines compared to revealed preference ICs (Figure 7C; $P < 0.0884 \times 10^{-8}$, Wilcoxon paired test; $N = 24$ participants). Despite their larger variability, both BDM isoline slope and curvature coefficients, derived from the respective β_2 / β_1-ratio and β_3 regression coefficient in Eq. 9, failed to differ significantly from the respective slope and curvature coefficients of revealed preference ICs (Eq. 3; Figure 7D, E; both $P > 0.05$, Wilcoxon test on BDM vs. IC coefficients paired from each participant; $N = 24$). Thus, despite larger variability, the BDM bids matched well the revealed preference ICs when assessed in a comparable way.

Taken together, the BDM bids provided good validation of the graphic representation of revealed preferences by the ICs and thus strengthen the validity of these conceptual schemes.

Performance comparison with non-human primates. The systematic IPs reflecting the trade-off between the two bundle components resulted in well-ordered ICs of our human participants and indicated underlying decision processes that properly integrated the utilities of all option components. The question arises to what extent these basic measures of multi-component integration in humans compared with these measures in non-human primates whose lack of verbal interaction makes their behavior primarily dependent on experienced reinforcement. Using the same experimental design and economic concepts as in our human participants, we had tested two rhesus monkeys during several months and estimated > 600 IPs that conformed to three to six convex or linear, negatively sloped ICs with each of five different types of two-liquid bundles (Pastor-Bernier et al. 2017). The ICs were consistent, as shown by out-of-sample prediction, transitivity and

independence of number of choice options, thus reflecting the animals' extended laboratory experience and allowing comparison with human performance.

We used two crucial measures for the decision process that is apparent in revealed preferences for multi-component choice options. First, to assess the accuracy of the estimated IPs, we computed the 95% CIs of the psychophysical fits to the choice probabilities when assessing individual IPs (which were used for constructing ICs). The CIs for IPs in humans were in a similar range but slightly smaller in comparison with all but one of the five bundle types in monkeys (Figure 8: compare A with B-F); the higher human accuracy was overall significant ($P = 0.000461$; Kolmogorov-Smirnov test; Figure 8G). Second, we tested specifically the accuracy of the fit of ICs to the IPs that were derived from the critical trade-off between the two bundle components; this trade-off reflected the integration of the two bundle components into single-dimensional measures. To this end, we computed the CIs of the hyperbolic IC fits to the estimated IPs. A narrow CI would indicate a good match of ICs to IPs and thus a good trade-off and utility integration between the two bundle components. The human participants showed slightly smaller CIs, and thus higher accuracy, than the monkeys ($P = 0.0018$; Figure 8H).

Thus, despite differences in number of participants ($N = 24$ for humans, $N = 2$ for monkeys) and task experience (single days for humans, several months for monkeys), the CI comparisons for IPs and ICs both suggested comparable, although slightly better, choice accuracy in humans compared to the highly experienced laboratory monkeys. These results suggest similarly good correspondence between the empirical ICs and the schematic conceptual graphs in both species.

Discussion

This study used basic formalisms of Revealed Preference Theory in order to empirically investigate the nature of human economic decision processes for multi-component rewards. We measured behavioral choices of small, tangible and immediately consumable rewards that were delivered in repeated trials in a controlled, non-verbal, laboratory setting. The estimated ICs were systematic and corresponded well to the graphic schemes of economic textbooks and decision research that are used for interpreting human choices (Laidler & Estrin 1989; Kreps 1990; Varian 1992; Mas-Colell et al. 1995; Perloff 2009). The experimental feasibility opens the possibility for concept-driven, empirical neuroimaging studies of rational and irrational behavior in normal and brain-damaged individuals.

Experimental controls and comparisons attested to the validity of these measurements; the preferences were graphically well represented by the asymmetric, nonlinear and non-overlapping ICs, as shown by leave-one-out statistics and accuracy of SVM and LDA decoding; the correspondence to BDM bids demonstrated validity beyond the specific binary choice eliciting mechanism; the similarity in performance accuracy with laboratory monkeys suggested sharing of this decision mechanism with an evolutionary close species. At the basis of this decision mechanism lies the orderly vectorial-to-scalar transformation; bundles are vectors, but utility and preference are scalars. The transformation was particularly well seen in the continuous, graded trade-off between bundle components. In contrast to 'lexicographic' preferences that concern only a single component, scalar revealed preferences concern both bundle components: preference was unchanged when reduction in one bundle component was compensated by increase in another bundle component; bundles with one smaller component could even be preferred to bundles with one larger component when the other component varied enough in the opposite direction (partial physical non-dominance), suggesting proportionate contributions of individual bundle components. Thus, the revealed preferences took both bundle components into account, and in a very systematic manner.

The representation of vectorial, multi-component choice options by single-dimensional neural signals is an open issue in experimental neuroeconomics. Our behavioral tests provide a formal, concept-driven foundation for investigating such signals in humans. These tests require repeated trials rather than one-shot tests typical for experimental economics; the multi-trial nature is well captured by stochastic choice theories that facilitate data interpretation due to decades of decision research. The immediately consumed, small payouts facilitate comparisons with animal studies and control for temporal discounting. The tangible payouts after small sets of trials do not rely on language and should reduce mental ruminations about future rewards and assure reliable

cooperation by the participants, thus reducing interfering and confounding brain activity. We were particularly interested in estimating whole maps of well-ordered, fitted ICs derived from multiple IPs that conform to predictive mathematical functions, rather than testing preferences for a few bundles with limited general validity. As a result, neuroeconomic work using the demonstrated empirical tests could investigate decision mechanisms not resolvable by purely behavioral examination. For example, it is unknown whether the bundle components are combined into a common neural utility estimate or are coded separately and integrated at a later stage. The ensuing neural decision process might occur between whole bundles with all components integrated, or involve hierarchical or parallel component-by-component comparisons. The systematic ICs may also allow to investigate neural underpinnings of specific theories, such as the switching of attentional processes between components conceptualised by multialternative decision field theory (Roe et al. 2001). Finally, our detailed ICs would be helpful for investigating neural mechanisms underlying inconsistent decision making when choice options are added (independence of irrelevant alternatives, IIA), beyond the decoy and attraction effects already being addressed (Chung et al. 2017; Gluth et al. 2017).

Our quantitative, psychophysically controlled preference assessment required repeated choices (Green & Swets 1966). Such a multi-trial approach is dictated by statistical requirements of neural studies, corresponds to the standard elicitation of choice functions in reinforcement learning (Sutton & Barto 1998) and allows comparison with animal studies requiring similar schedules (Kagel et al. 1975; Pastor-Bernier et al. 2017). In humans and monkeys, these methods deliver systematically varying choice probabilities rather than single, all-or-none selection. The visible trial-by-trial variations are assumed to reflect underlying random processes that make the choice process stochastic, as captured by choice probability and probabilistic choice functions in random-choice theories (Luce 1959; McFadden & Richter 1990; McFadden 2004; Stott 2006). We appreciate that our multi-trial schedule is at odds with the frequently employed, standard, single-shot, deterministic assessment of ICs in experimental economics (Thurstone 1931; MacCrimmon & Toda 1969), but we believe that the obtained consistent and robust ICs validate the approach.

Economic choice often involves substantial but imaginary sums of money and consumer items, or random, singular payouts (Simonson 1989; Simonson, Tversky 1993; Rieskamp et al. 2006). By contrast, our payout schedule was tailored to requirements of neural studies and allowed the participants to actually consume the chosen rewards over tens and hundreds of trials (while controlling for satiety effects), as is typically required for animal experiments. These behavioral choices resembled small daily activities, such as consuming snacks and drinks, and were met by the good motivation of our participants. In this way, we obtained well-ordered ICs that provided valid representations of revealed preferences, without requiring imagined items or unreasonably large sums of money.

Previous investigations of economic choice of multi-component options have revealed several inconsistencies, including reference biases, differences between willingness-to-pay and willingness-to-sell, and violations of independence of irrelevant alternatives such as the similarity, compromise, asymmetric dominance and attraction effects (Tversky 1972; Knetsch & Sinden 1984; Simonson 1989; Tversky & Simonson 1993; Bateman et al. 1997; Rieskamp et al. 2006). These phenomena may well be due to the intrinsically adaptive nature of biological processes (Soltani et al. 2012; Li et al. 2018), rather than invalidating the revealed preference concept itself. We aimed to avoid interference from adaptive processes by designing stable and highly reproducible test conditions in a well-controlled laboratory environment, non-reliance on verbal report, single, uninterrupted test sessions, singular changes of bundle components, constant direction of testing (from top left to bottom right on the revealed preference map), and preventing satiety by limiting total milkshake intake to 200 ml. As a result, our IPs remained stable over successive testing steps. We used the exact same conditions for eliciting BDM bids, which may have facilitated their correspondence to revealed preference ICs. With these testing conditions, we avoided some known compromising factors, which might help to identify more basic factors underlying decision irregularities.

The coefficients of hyperbolic fits to the ICs characterize numerically the representation of revealed preferences. The slope coefficient indicates the relative weighting of the two milkshakes. For example, the amount of equally revealed preferred single milkshakes (graphed at the respective x-axis and y-axis intercepts of the two-dimensional map) was lower for high-fat (component A)

than high-sugar (component B) milkshakes in our participants, which was represented by a IC slope steeper than -45 degrees; thus, participants would have revealed preferred high-fat over high-sugar milkshakes if they came in same amounts. Such asymmetric IC slopes are also seen with various bundles in primates (Pastor-Bernier et al. 2017). A key reason for the asymmetry may be the use of a physical amount scale; in addition, the revealed preference among the milkshake bundles depended also on their fat and sugar content, for which there is no simple common physical scale. Thus, the scaling of milkshakes in units of physical amount would explain the IC asymmetry. Further, the slope for identical bundle compositions varied between our participants, which demonstrates an additional subjective component in revealed preferences. Despite these scaling and subjectivity issues, our estimated ICs had well-ordered slopes and failed to overlap.

Further, the convex IC curvature in most of our participants indicated that unproportionately smaller amounts of combined sugar-fat milkshakes (at IC center) were equally revealed preferred as larger amounts of milkshakes containing primarily sugar or fat (positioned closer to IC boundaries near x- and y-intercepts). The BDM isolines confirmed the convex curvature; unproportionately smaller amounts of combined sugar-fat milkshakes (at isoline center) elicited the same bids as larger amounts of milkshakes containing primarily sugar or fat (positioned closer to isoline x- and y- intercepts). A previous study on food snacks found also higher BDM bids for sugar-fat combinations than for fat or sugar alone, despite similar calories (DiFeliceantonio et al. 2018). However, not all ICs must be convex; bundles with one unattractive component showed concave ICs in primates (Pastor-Bernier et al. 2017). In addition, IC convexity may be ascribed to concave utility of each bundle component (Perloff 2009); more of the same component (closer to IC x- or y-intercepts) has lower marginal gain, and therefore the decision-maker is willing to trade in more of a component of which she has much for one unit of the other component of which she has little. This may be the reason why most participants showed convex ICs. Thus, the systematic estimates of the defining slope and curvature coefficients suggest that the ICs reliably represented well-ordered preferences for the milkshake bundles.

The estimating mechanism for BDM bids differs substantially from that underlying revealed preference ICs. Both the auction-like bidding and the mechanics of cursor movement differed from the choice between two simultaneously presented options. BDM bidding is incentive compatible, such that erroneous or deceiving bids lead to suboptimal outcome, as conceptualised by the expected cost of misbehavior (Lusk et al. 2007); bidders should state exactly their subjective value to avoid paying an exaggerated price or foregoing a desired object. With these properties, BDM bidding constituted a well-controlled, authoritative test for eliciting true subjective values, thus providing a useful validation mechanism for preferences that were revealed by binary choice. Indeed, the obtained SVM- and LDA-consistent BDM bids followed the preference scheme of ICs, namely higher bids for revealed preferred bundles and similar bids with choice indifference despite varying bundle composition. Most strikingly, hyperbolically fitted BDM isolines closely resembled the hyperbolically fitted revealed preference ICs, both in terms of graphics and numeric coefficients. The only notable difference was higher BDM bid variability. Our data align well with, and extend, the previously noted correspondence between binary choices and BDM bids for single-component options and bundle choices on paper or via verbal communication (Roosen et al. 2017). Thus, the BDM estimating mechanism resulted in bids that corresponded well to the graphic scheme of ICs conceptualised by Revealed Preference Theory and thus validated our empirical IC assessments.

The revealed preferences of our human participants indicated a similar level of integration of the two-component choice options as seen in our previous study on rhesus monkeys (Pastor-Bernier et al. 2017) and, in a more general way, also observed in earlier rodent experiments (Kagel et al. 1975; 1995). In two monkeys, we estimated > 600 IPs that conformed to three to six convex or linear, negatively sloped ICs with each of five different types of two-liquid bundles. Preferences for bundle types with other liquid compositions were characterized by concave or positively sloped ICs reflecting satiety or disfavored juices. Their ICs represented revealed stochastic preferences in a consistent and orderly manner, as shown by out-of-sample prediction by homothetic polynomials, various transitivity tests, and axiomatic compliance with independence of number of choice options (two and three) with additional > 900 IPs. In contrast to the unexperienced humans, the monkeys had several months of experience with the specific choice options and their reward distributions in

well-controlled, stable laboratory conditions. Despite these differences, the numeric comparisons of choice accuracy for estimating IPs and ICs demonstrated very similar performance. The IP estimation reflects the reliability of the general stochastic decision-making process, as inferred from choice probability; its similar accuracy in both species is reassuring and a necessary requirement for establishing ICs. By contrast, the IC estimation from fits to the IPs reflects specifically the trade-off that defines the graded and orderly vectorial-to-scalar transformation; its similar accuracy suggests that this crucial expression of the decision process is rather precise in both species. Thus, the robustness of the performance across species, despite their different cognitive capacity and task experience, demonstrates the basic nature of this economic decision process and validates key assumptions of Revealed Preference Theory.

References

Alphonce, R., & Alfnes, F. (2017). Eliciting consumer WTP for food characteristics in a developing context: Application of four valuation methods in an African market. *J Agricult Econ, 68,* 123-142.

Amemiya, T. (1981). Qualitative response models: A survey. *J Econ Lit, 19,* 1483-1536.

Bateman, I., Munro, A., Rhodes, B., Starmer, C., & Sugden, R. (1997). A test of the theory of reference-dependent preferences. *Quart J Econ, 112,* 479-505.

Becker, G. M., DeGroot, M. H., & Marschak, J. (1964). Measuring utility by a single-response sequential method. *Behav Sci, 9,* 226-232.

Choi, S., Fisman, R., Gale, D., & Kariv, S. (2007). Consistency and heterogeneity of individual behavior under uncertainty. *Am Econ Rev, 97,* 1921–1938.

Chung, H.-K., Sjöström, T., Lee, H.-J., Lu, Y.-T., Tsuo, F.-Y., Chen, T.-S., Chang, C.-F., Juan, C.-H., Kuo, W.-J., & Huang, C.-Y. (2017). Why Do Irrelevant Alternatives Matter? An fMRI-TMS study of context-dependent preferences. *J Neurosci, 37,* 11647–11661.

DiFeliceantonio, A. G., Coppin, G., Rigoux, L., Thanarajah, S. E., Dagher, A., Tittgemeyer, M., & Small, D. M. (2018). Supra-additive effects of combining fat and carbohydrate on food reward. *Cell Metab, 28*, 33-44.

Fellows, L. K. (2006). Deciding how to decide: ventromedial frontal lobe damage affects information acquisition in multi-attribute decision making. *Brain, 129,* 944-952.

Fisher, I. (1892). Mathematical Investigations in the theory of value and prices. *Trans Connecticut Acad, 9,* 1–124.

Gluth, S., Hotaling, J. M., & Rieskamp, J. (2017). The attraction effect modulates reward prediction errors and intertemporal choices. *J Neurosci, 37,* 371-382.

Green, D. M., & Swets, J. Signal detection theory and psychophysics: New York: Wiley, 1966.

Harris, A., Clithero, J. A., & Hutcherson, C. A. (2018). Accounting for Taste: A multi-component neurocomputational model explains the neural dynamics of choices for self and others. *J Neurosci, 38,* 7952–7968.

Kagel, J. H., Battalio, R. C., Rachlin, H., Basmann, R. L., Green, L., & Klemm, W. R. (1975). Experimental studies of consumer demand behavior using laboratory animals. *Econ Inquiry, 13*, 22-38.

Kagel, J. H., Battalio, R. C., & Green, L. (1995). *Economic Choice Theory: An Experimental Analysis of Animal* Behavior. Cambridge University Press, Cambridge.

Knetsch, J. L. (1989). The endowment effect and evidence of nonreversible indifference curves. *Am Econ Rev, 79,* 1277-1288.

Knetsch, J. L., & Sinden, J. A. (1984). Willingness to pay and compensation demanded: Experimental evidence of an unexpected disparity in measures of value. *Quart J Econ, 99,* 507-521.

Kreps, D. M. (1990). A Course in Microeconomic Theory. Harlow: Pearson Education Ltd.

Kurtz-David, V., Persitz, D., Webb, R., & Levy, D. J. (2019). The neural computation of inconsistent choice behavior. *Nat Comm, 10,* 1583.

Laidler, D. E. W., & Estrin, S. (1989). Introduction to Microeconomics. New York: Philip Allan.

Levy, I., Belmaker, L., Manson, K., Tymula, A., & Glimcher, P. W. (2012). Measuring the subjective value of risky and ambiguous options using experimental economics and functional MRI methods. *J Vis Exp*, *67,* e3724.

Li, V., Michael, E., Balaguera, J., Castañón, S. H., & Summerfield, C. (2018). Gain control explains the effect of distraction in human perceptual, cognitive, and economic decision making. *Proc Natl Acad Sci, 115*, E8825-E8834.

Luce, R. D. (1959). Individual Choice Behavior: A Theoretical Analysis. New York: Wiley.

Lusk, J. L., Alexander, C., & Rousu, M. C. (2007). Designing experimental auctions for marketing research: The effect of values, distributions, and mechanisms on incentives for truthful bidding. *Rev Market Sci, 5*, 1-32.

MacCrimmon, K. R., & Toda, M. (1969). The experimental determination of indifference curves. *Rev Econ Stud, 36,* 433-451.

Mas-Colell, A., Whinston, M. D., & Green, J. R. (1995). Microeconomic theory. New York: Oxford Univ Press.

McFadden, D., & Richter, M. K. (1990). Stochastic rationality and revealed stochastic preference. Preferences, Uncertainty, and Optimality, Essays in Honor of Leo Hurwicz, Westview Press: Boulder, CO, 161-186.

McFadden, D. L. (2004). Revealed stochastic preference: A synthesis. *Econ Theory, 26*, 245–264.

Medic, N., Ziauddeen, H., Vestergaard, M. D., Henning, E., Schultz, W., Farooqi, I. S., & Fletcher, P. C. (2014). Dopamine modulates the neural representation of subjective value of food in hungry subjects. *J Neurosci, 34*, 16856-16864.

Pastor-Bernier, A., Plott, C. R, & Schultz W. (2017). Monkeys choose as if maximizing utility compatible with basic principles of revealed preference theory. *Proc Natl Acad Sci, 114*, E1766-E1775.

Pastor-Bernier, A., Stasiak, A., & Schultz, W. (2019). Orbitofrontal signals for two-component choice options follow indifference curves of Revealed Preference Theory. *Nat Comm, 10,* 4885.

Perloff, J. M. (2009). Microeconomics. Boston: Addison-Wesley.

Pelletier, G., & Fellows, L. K. (2019). A critical role for human ventromedial frontal lobe in value comparison of complex objects based on attribute configuration. *J Neurosci, 39*, 4124-4132.

Plassmann, H., O'Doherty, J., & Rangel, A. (2007). Orbitofrontal cortex encodes willingness to pay in everyday economic transactions. *J Neurosci, 27,* 9984-9988.

Razzaghi, M. (2013). The Probit Link Function in Generalized Linear Models for Data Mining Applications. *J Mod Appl Stat Meth, 12,* Issue 1, Article 19.

Rieskamp, J., Busemeyer, J. R., & Mellers, B. A. (2006). Extending the bounds of rationality: Evidence and theories of preferential choice. *J Econ Lit, 44,* 631-661.

Roe, R. M., Busemeyer, J. R., & Townsend, J. T. (2001). Multialternative decision field theory: a dynamic connectionist model of decision making. *Psychol Rev, 108*, 370-392.

Roosen, J., Marette, S., & Blanchemanche, S. (2010). Value elicitation using BDM and a discrete choice mechanism. *Econ Bull, 30,* 1554-1563.

Samuelson, P. A. (1937). A note on measurement of utility. *Rev Econ Stud* 4, 155-161.

Samuelson, P. A. (1938). A note on the pure theory of consumer's behavior. *Economica, 5,* 61-71.

Simonson, I. (1989). Choice based on reasons: The case of attraction and compromise effects. *J Consum Res, 16,* 158-174.

Soltani, A., De Martino, B., & Camerer, C. (2012). A range-normalization model of context-dependent choice: A new model and evidence. *PLOS Comput Biol, 8,* e1002607.

Stott, H. P. (2006). Cumulative prospect theory's functional menagerie. *J Risk Uncert, 32*, 101-130.

Sutton, R. S., & Barto, A. G. (1998). Reinforcement Learning: An Introduction. Cambridge MA: MIT Press.

Suzuki, S., Cross, L., & O'Doherty, J. P. (2017). Elucidating the underlying components of food valuation in the human orbitofrontal cortex. *Nat Neurosci, 20*, 1780–1786.

Thurstone, L.L. (1931). The indifference function. *J Soc Psychol, 2*, 139-167.

Tsutsui, K.-I., Grabenhorst, F., Kobayashi, S., & Schultz, W. (2016). A dynamic code for economic object valuation in prefrontal cortex neurons. *Nat Comm, 7*, 12554.

Tversky, A. (1969). Intransitivity of preferences. *Psychol Rev, 76*, 31-48.

Tversky, A. (1972). Elimination by aspects: A theory of choice. *Psychol Rev, 79*, 281-99.

Tversky, A., & Simonson, I. (1993). Context-dependent preferences. *Management Sci, 39*, 1179-1189.

Varian, H. R. (1992). Microeconomic Analysis. New York: Norton.

von Neumann J, Morgenstern O. (1944). The Theory of Games and Economic Behavior. Princeton: Princeton University Press.

Legends

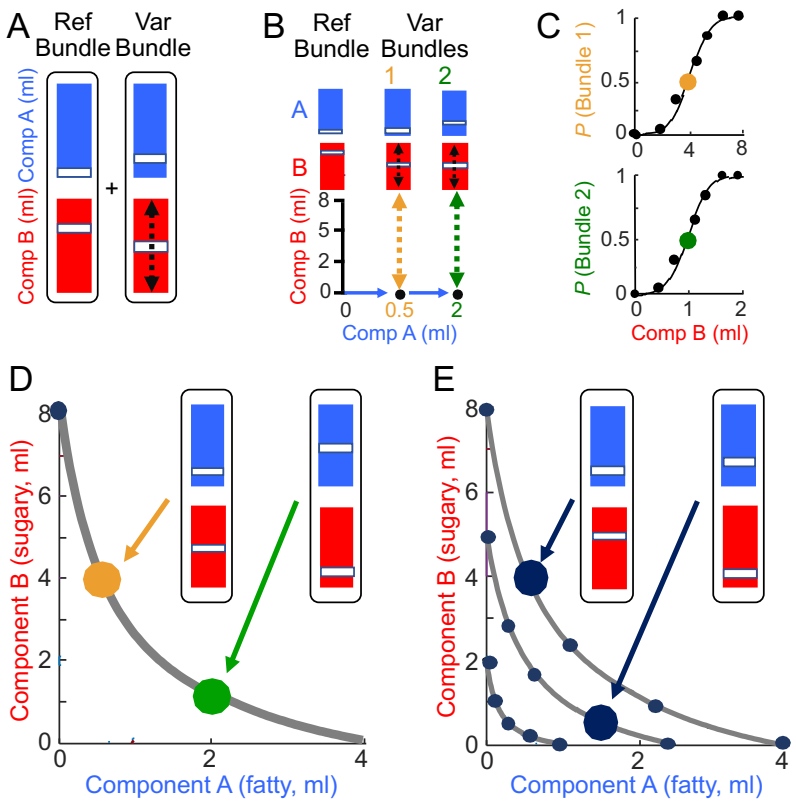

Figure 1. Binary choice task and experimental design.
(A) Bundle stimuli for binary choice. Each bundle contained two components with amounts between 0.0 and 10.0 ml. Component (Comp) A: low-sugar high-fat milkshake, Comp B: high-sugar low-fat milkshake. Participants chose between the Reference Bundle and the Variable Bundle, whose locations on a computer monitor alternated pseudorandomly between fixed left and right positions.
(B) Psychophysical test design. In the Reference Bundle, components A and B were set to 0.0 ml and 8.0 ml, respectively; in the Variable bundle, component A was set to a specific test amount while the amount of component B was psychophysically varied (dashed arrows).
(C) Psychophysical assessment of two example choice indifference points (IP; choice probability P = 0.5 each option; yellow and green) in a typical participant. During repeated trials, the participant chose between the preset Reference Bundle and the Variable Bundle. Component A of the Variable Bundle was set to specific test amounts (here 0.5 ml and 2.0 ml, respectively). Component B of the Variable Bundle varied between seven randomly selected amounts of component B. IPs were estimated from choice probabilities, p (Bundle 1) and p (Bundle 2), respectively, using the probit model (Eqs. 1, 2).
(D) Schematic indifference curve (IC), fitted by a hyperbola (Eqs. 3, 3a) to all equally revealed preferred but differently composed bundles. Yellow and green dots show IPs relative to the bundle marked by blue dot, as estimated in B, C.
(E) Schematic map of three ICs. Increasing distance from origin represents higher utility; all bundles on higher ICs are revealed preferred to all bundles on lower ICs. Heavy dots denote a preference relationship between two bundles with oppositely varying physical amounts of component A (top-IC bundle is revealed preferred to mid-IC bundle despite lower physical amount of component A; partial physical non-dominance).

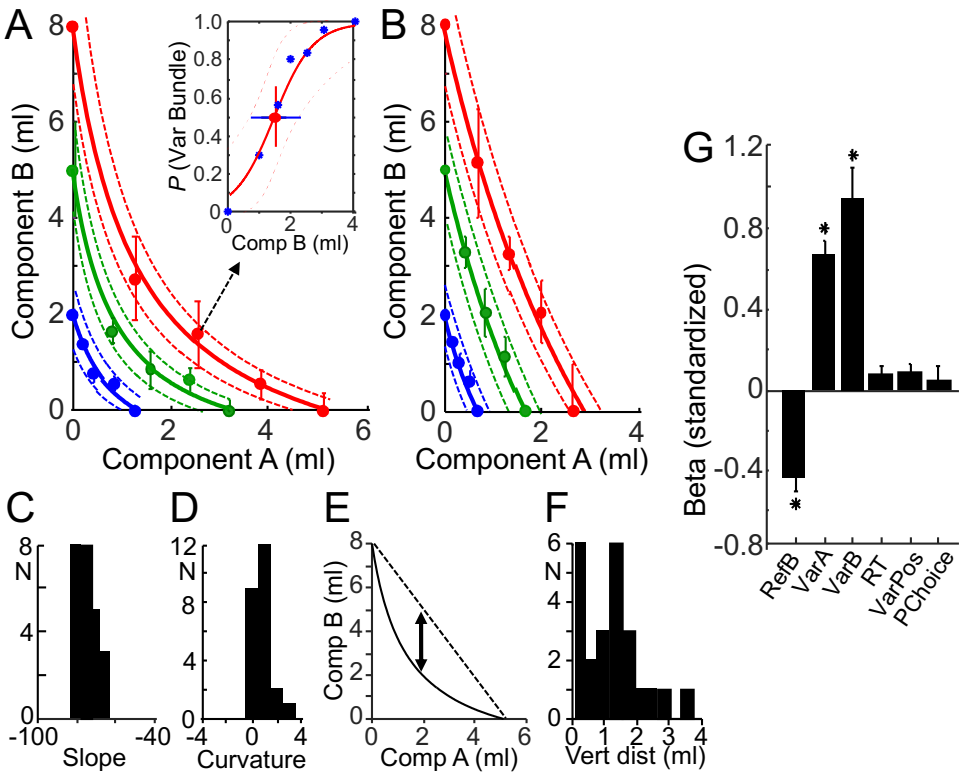

Figure 2. Empirical indifference curves (IC) representing revealed preferences.
(A) Typical convex ICs from an example participant, as seen in 18 of the 24 participants. Component A was a low-sugar high-fat milkshake; component B was a high-sugar low-fat milkshake. Solid lines show hyperbolically fitted ICs, dotted lines show 95% confidence intervals of fits. Dots show bundles that are equally preferred on the same IC (IPs). Inset: psychophysical assessment of indifference point (IP) marked on highest IC (test points in blue, IP estimated by probit regression in red).
(B) Typical linear ICs from another example participant, as seen in 6 of the 24 participants.
(C, D) Distributions of slope and curvature, respectively, of hyperbolically fitted ICs from all 24 participants (coefficients β_2 / β_1 and β_3 in Eq. 3, respectively). N = number of participants.
(E) Scheme of intuitive numeric assessment of IC curvature: maximal vertical distance (ml of component B on y-axis) between fitted IC (curve) and a straight line connecting the x-axis and y-axis intercepts. A distance of > 0.0 ml indicates convexity, whereas a 0.0 ml distance indicates perfect linearity.
(F) Distribution of convex curvature, as measured using the scheme shown in E. The two peaks indicate six participants each with similarly near-linear ICs and similarly convex ICs, respectively.
(G) Specificity of bundle choice, as opposed to unrelated parameters. Bar graph shows standardised beta (β) regression coefficients for choice of Variable Bundle over Reference Bundle (Eq. 5), as assessed for each individual participant and then averaged across all 24 participants. RefB, component B in Reference Bundle; VarA and VarB, components A and B in Variable Bundle; RT, reaction time; VarPos, left-right position of Variable Bundle stimulus; PChoice, choice in previous trial. Error bars show standard error of the mean (SEM). $^*P \leq 0.02$.

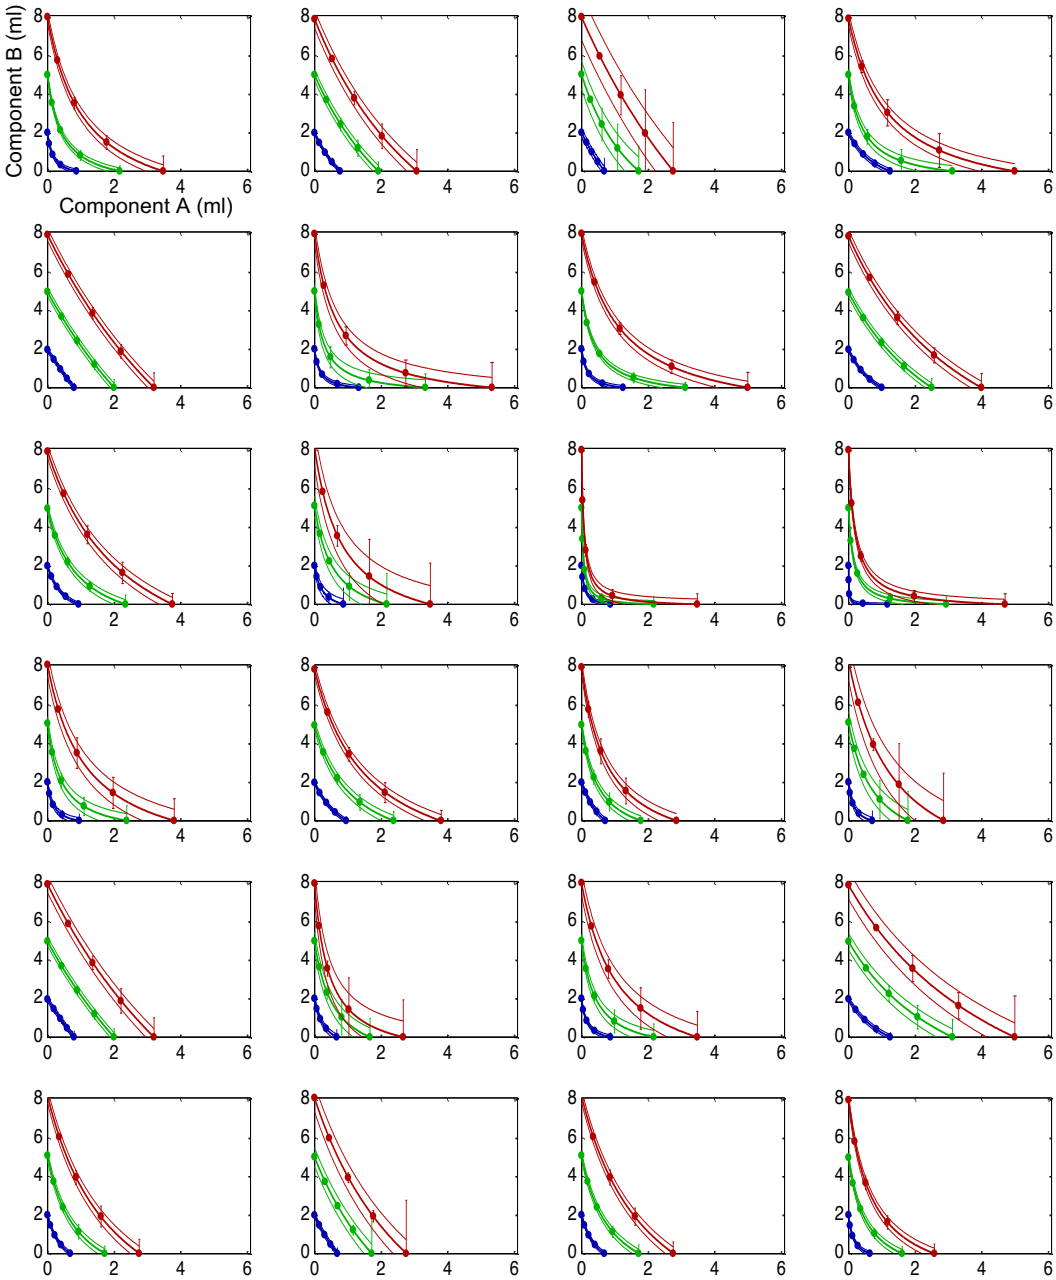

Figure 3. Empirical indifference curves (IC) from all 24 participants. Note the fanning out of the confidence intervals towards the bottom right in each graph (thin lines), which likely reflects the progression of choice testing: the Reference Bundle was kept constant at the y-axis intercept (x = 0), whereas testing with the Variable Bundle progressed from top left to bottom right. Same conventions as for Figure 2.

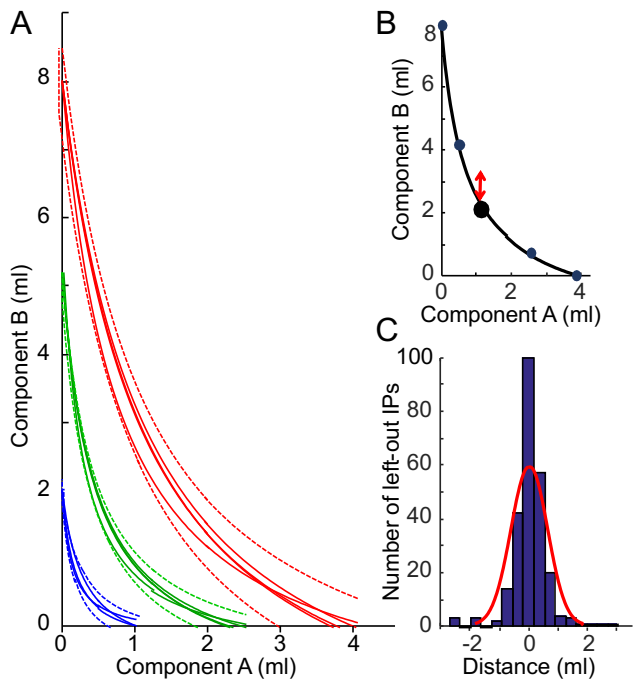

Figure 4. Leave-one-out validation of estimated indifference curves (ICs).
(A) Graphic assessment with an example participant: hyperbolically refitted ICs with one left-out indifference point (IP) (solid lines), plotted together with 95% confidence intervals of the original hyperbolically fitted ICs (dotted lines). The refitting resulted in 4 new ICs (partly overlapping) at each of three levels. None of the refitted IPs fell outside the original confidence intervals.
(B) Scheme of numeric assessment: distance in ml (in ml on y-axis, red) between the left-out IP on the original IC (heavy black dot on black curve) and the refitted IC (green).
(C) Histogram of distance between refitted ICs and left-out IPs across all subjects. Skewness was -0.12, suggesting rather symmetric distribution around the mean (modeled in red).

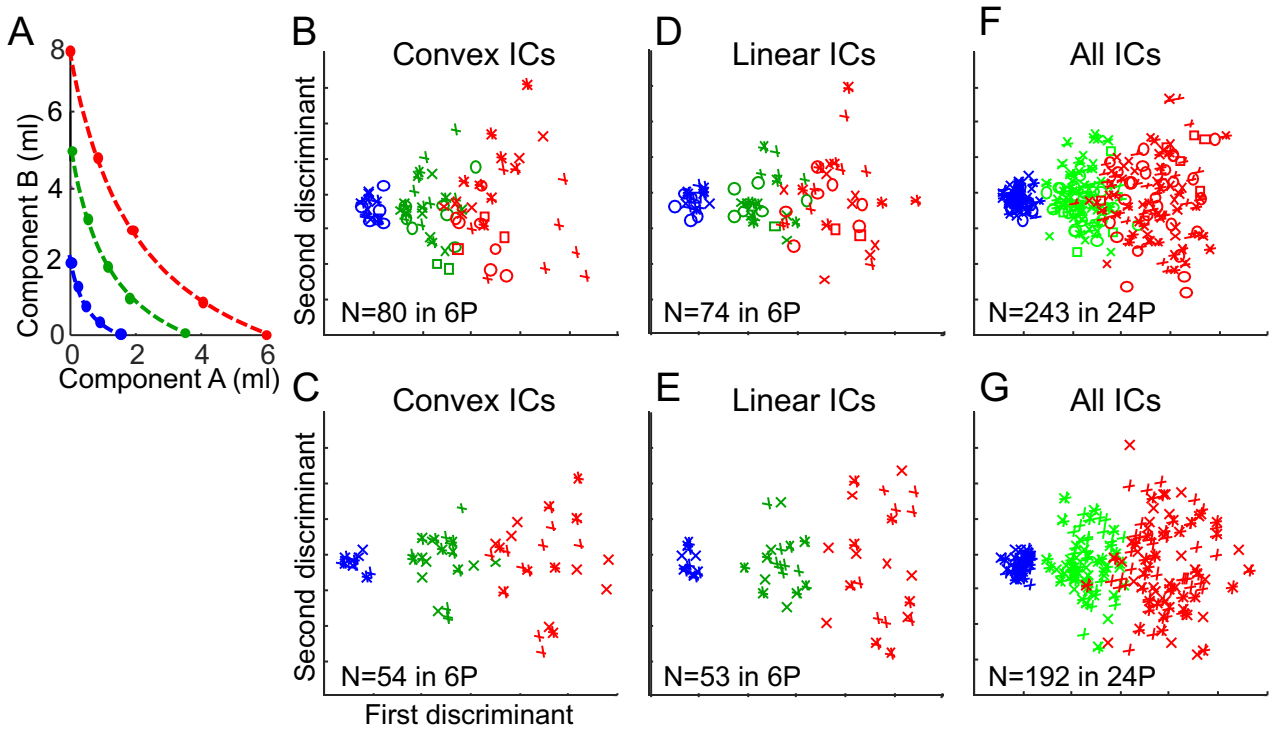

Figure 5. Visualization of bundle decoding using Linear Discriminant Analysis (LDA).
(A) Schematics of bundle decoding at psychophysically estimated points of equal revealed preference (indifference points, IPs, plotted along the dotted lines). Following the notions of revealed preference, LDA should show accurate decoding across the three preference levels (green, blue, red) but not along each level.
(B) LDA bundle distinction between three levels of revealed preference (81-100% binary numeric decoding accuracy between two levels; first discriminant) but not along same preference levels (second discriminant) in N = 80 bundles from six participants (6P) with similar convex ICs (same convention applies to all panels). Bundles on the three preference levels are colored blue, green and red according to distance from origin, red being highest. The five bundles on each preference level are marked from top left to bottom right with 'o', '*', '+', 'x' and '□' symbols. Due to the arbitrariness of the scale, numbers are not indicated.
(C) As B but for partly physically non-dominating bundles (one lower component in preferred bundle than in alternative bundle (97-100% accuracy; first discriminant).
(D) As B but for six participants with linear ICs (90-100% accuracy).
(E) As D but for partly physically non-dominating bundles (96-100% accuracy).
(F) As B but for all 24 participants (18 with convex IC, six with linear ICs) (83-100% accuracy).
(G) As F but for partly physically non-dominating bundles (97-100% accuracy).

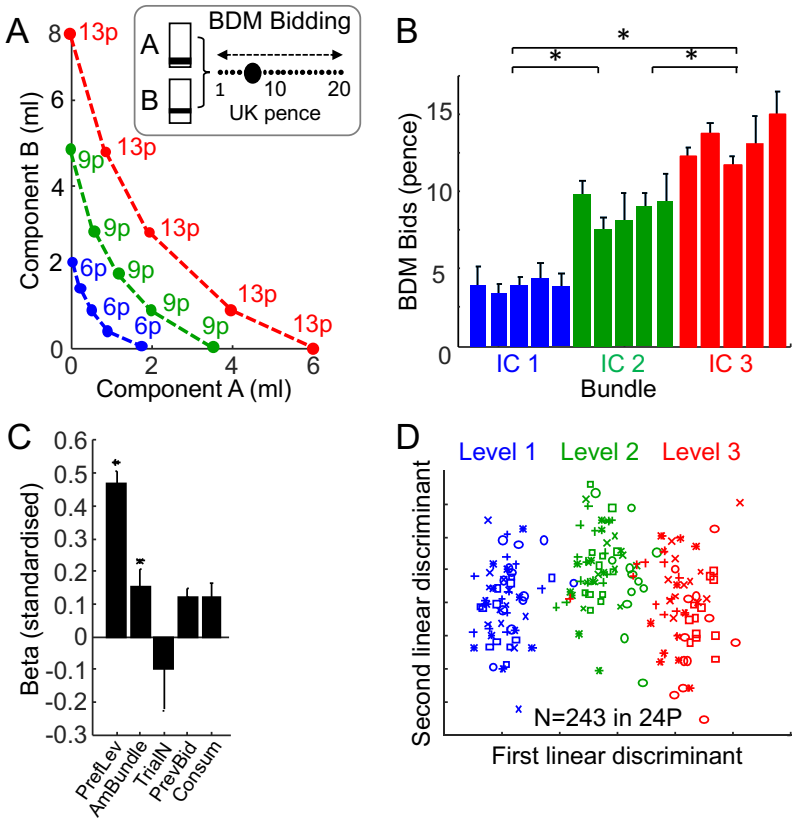

Figure 6. Characteristics of BDM bids for bundles at different revealed preference levels.
(A) Schematics of positions of bundles used for eliciting BDM bids at psychophysically estimated points of equal revealed preference (indifference points, IPs, connected by dotted lines). Following the notions of trade-off and increasing revealed preference, BDM bids should be similar for equally valued bundles (along the dotted lines) but higher for revealed preferred bundles (farther away from origin). We tested 5 bundles per level, 3 levels, 12 repetitions, total of 180 bids. Inset: BDM task. Each participant bid for the visually presented two-component (A, B) bundle by moving the black dot cursor using the leftward and rightward horizontal arrows on a computer keyboard. Numbers indicate example bids (in UK pence).
(B) Mean BDM bids from a typical participant. The bids were rank ordered (Spearman Rho = 0.83, $P < 0.001$) across bundles positioned on increasing indifference curves (blue, green, red) and failed to overlap between lowest, medium and highest levels. There was a significant difference for BDM bids between levels but not within levels (*$P < 0.001$; two-way Anova followed by Tukey-Kramer test). Data are shown as mean ± SEM (standard error of the mean), N = 12 bids per bar.
(C) Specificity of monetary BDM bids, as opposed to unrelated parameters. Bar graph showing the standardised beta (β) regression coefficients for BDM bids (Eq. 8), as assessed for each individual participant and then averaged across all 24 participants. Abbreviations: PrevLev, revealed preference level (low, medium, high); AmBundle, summed currency-adjusted amount of both bundle components; TrialN, trial number; PrevBid, BDM bid in previous trial; Consum, accumulated drinks consumption. Error bars show SEMs. * $P \leq 0.020$.
(D) Visual decoding by Linear Discriminant Analysis (LDA) of BDM bids for bundles on different preference levels (first discriminant; 88-100% numeric decoding accuracy; $P = 9.46 \times 10^{-12}$) and along same preference levels (second discriminant; 43-51% accuracy) (N = 243 bundles; all 24 participants). The LDA used scalar BDM bids for bundles positioned at IPs shown in A. Same conventions as for Figure 5A.

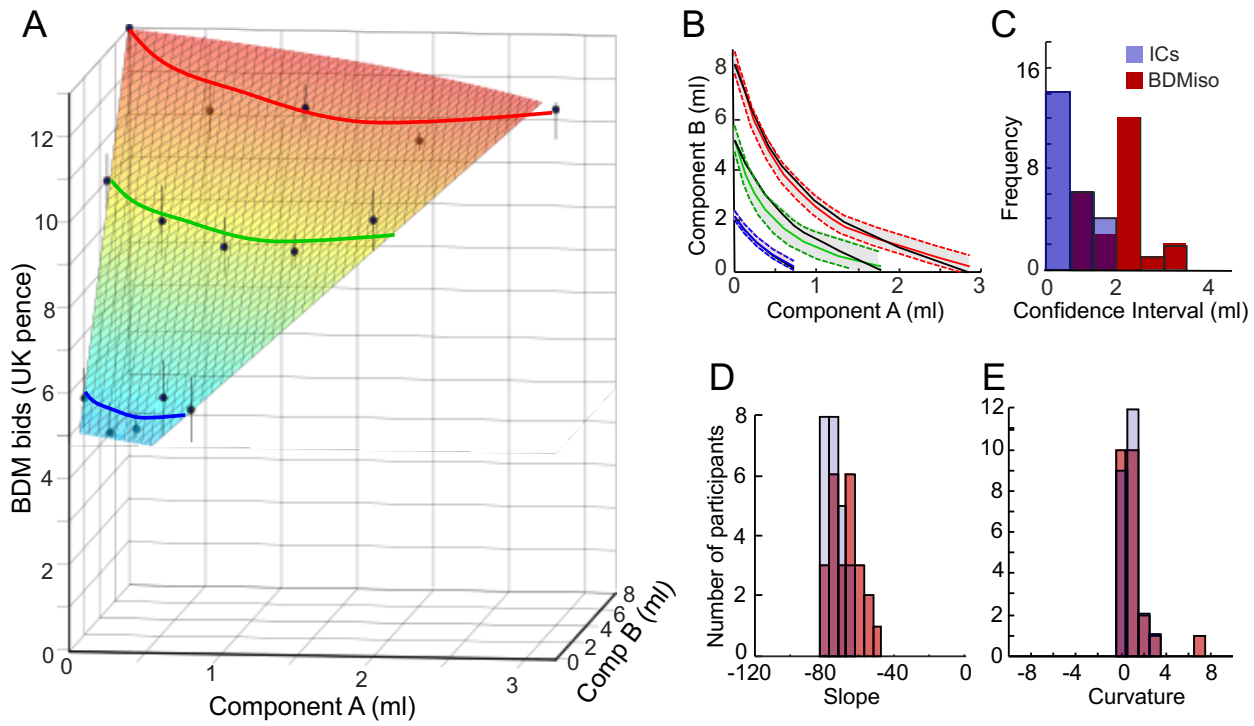

Figure 7. Comparison between BDM bids and indifference curves (IC).
(A) Three-dimensional representation of hyperbolically fitted BDM isolines from a typical participant (mean ± SEM). The BDM bids were made for bundles that had been placed on psychophysically estimated revealed preference indifference points (black dots). The BDM isolines were similar along same levels and increased across revealed preference levels.
(B) Match between hyperbolically fitted isolines of mean BDM bids (black; Eqs. 9, 9a) and hyperbolically fitted revealed preference ICs (blue, green, red; ± 95% confidence intervals, CI, shaded) from a typical participant. Thus, the BDM isolines fell within the respective CIs of the revealed preference ICs.
(C) CIs of fitted BDM isolines (red) and revealed preference ICs (blue) from all 24 participants (averaged along each isoline/IC). The larger 95%CIs of BDM isolines (BDMiso) suggest more variability compared to revealed preference ICs.
(D, E) Comparison of slope and curvature coefficient estimates, respectively, between BDM isolines (red) and revealed preference ICs (blue; same data as shown in Figure 2C, D) from all 24 participants.

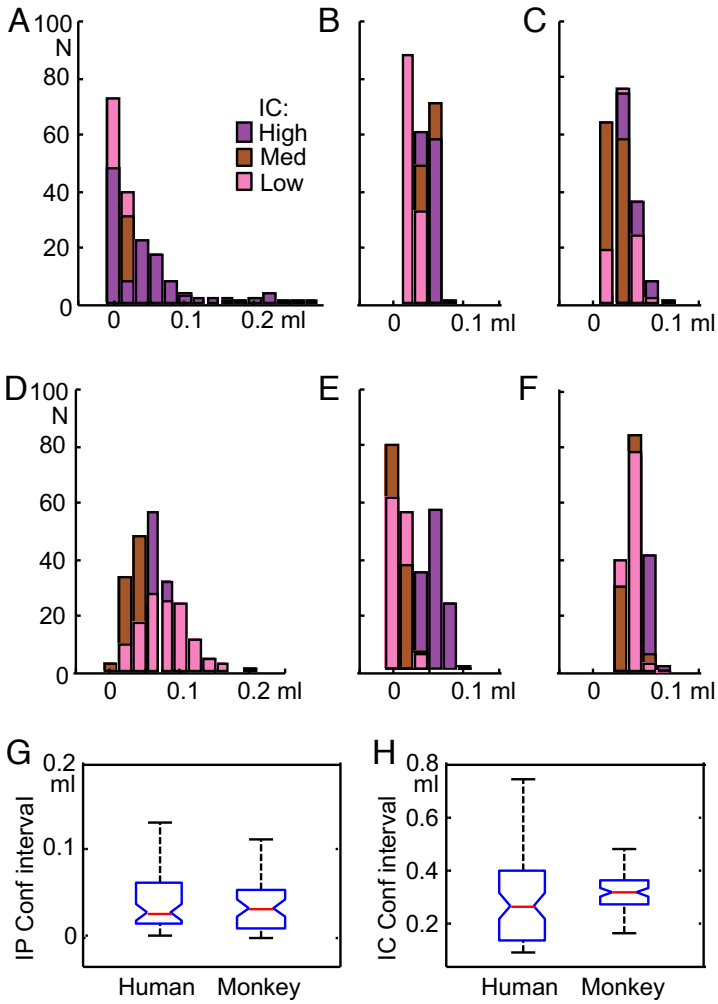

Figure 8. Comparison of choice accuracy with rhesus monkeys.
(A) Confidence intervals (CI, 95%) from psychophysical fits for estimation of 5 choice indifference points (IP) on each of 3 indifference curves (IC) (pink, brown, violet) from each of 24 human participants (0.032 ± 0.005 ml; N = 360 IPs), normalized to reward range for monkeys (0 - 1.0 ml; milkshake). Inset applies to A - F. All columns start at bottom.
(B - F) CIs (95%) from monkeys for five bundle types. N = 40 - 50 IPs / bundle type. To obtain same sample size as in humans, 360 randomly selected IPs were resampled 100 times and averaged in each of 360 bins, separately for each bundle type. All monkey data are from Pastor-Bernier et al. (2017).
(B) CIs (95%) for bundle type (blackcurrant juice, grape juice) in monkey A (0.043 ± 0.008 ml of blackcurrant juice; mean ± SEM).
(C) as B but for bundle (blackcurrant juice, water) (0.040 ± 0.006 ml).
(D) as B but for bundle (blackcurrant juice, strawberry juice) (0.061 ± 0.011 ml).
(E) as B but bundle (blackcurrant juice, water) in monkey B (0.065 ± 0.016 ml).
(F) as B but for bundle (blackcurrant juice, strawberry juice (0.025 ± 0.013 ml).
(G) CIs (95%) from psychophysical fitting for IP estimation. Box plots show medians (0.0310 ml milkshake normalized for humans; 0.0344 ml blackcurrant juice for monkey; P = 0.000461; CIs for N = 360 IPs each species; two-tailed Kolmogorov-Smirnov test; red), 25th and 75th percentiles (blue) and most extreme values (black).
(H) CIs (95%) from hyperbolic fitting for IC estimation. Medians were 0.2658 ml milkshake normalized for humans; 0.3038 ml blackcurrant juice for monkey; P = 0.0018; CIs for N = 72 ICs each species (reduced for monkeys by random selection to correspond to 3 ICs in 24 humans); K-S test.

Table 1. Decoding accuracy, as assessed with support vector machine.

	Bundles			BDM bids		
Participant C – convex IC L – linear IC	Lev1 vs. Lev2	Lev1 vs. Lev3	Lev2 vs. Lev3	Lev1 vs. Lev2	Lev1 vs. Lev3	Lev2 vs. Lev3
C1	83.9	100.0	68.8	70.6	81.0	54.4
C2	86.8	99.8	68.7	44.0	51.0	47.4
C3	72.7	96.0	58.7	56.6	71.3	59.3
C4	88.0	99.5	67.6	67.7	76.3	62.5
C5	99.1	100.0	81.4	64.8	69.9	49.8
C6	93.7	99.5	61.9	53.5	65.9	54.0
C7	66.2	75.5	49.5	70.0	72.2	43.6
C8	72.9	81.8	52.6	49.8	67.2	62.6
C9	80.0	100.0	80.0	68.4	83.8	62.1
C10	99.6	100.0	69.8	47.0	46.6	43.7
C11	90.2	99.5	55.4	68.9	75.7	53.1
C12	96.8	100.0	71.2	47.3	57.9	47.1
C12	86.1	87.7	56.7	58.7	60.9	48.2
C14	80.0	100.0	80.0	63.2	74.3	56.8
C15	99.8	100.0	85.5	67.8	78.6	54.8
C16	100.0	100.0	90.0	57.8	62.5	50.4
C17	100.0	100.0	100.0	44.8	70.9	68.4
C18	99.0	99.9	55.3	68.1	75.3	51.7
L19	100.0	100.0	88.6	43.5	65.8	59.6
L20	100.0	100.0	95.9	69.2	84.9	70.1
L21	100.0	100.0	100.0	62.2	66.0	44.4
L22	100.0	100.0	83.6	54.8	52.5	45.9
L23	100.0	100.0	100.0	63.2	58.4	47.8
L24	100.0	100.0	80.0	61.4	68.5	57.7

Accuracy is shown in % of numbers of bundles and BDM bids correctly assigned to one of the two tested revealed preference levels (averages from 150 iterations of classification of 10 pseudorandomly selected bundles / BDM bids from two tested revealed preference levels). Bundles were tested at psychophysically assessed choice indifference points (IP) on one of three revealed preference levels (corresponding to the three indifference curves, IC, fitted to the IPs). The tested bundles had been pseudorandomly selected from 5 unique estimated IPs on each of two revealed preference levels under study. BDM: Becker-DeGroot-Marschak auction-like bidding mechanism. Lev1, Lev2, Lev3: revealed preference levels, numbered according to distance from origin. The 24 participants were labelled as C or L according to their IPs being fitted best to convex or linear ICs.

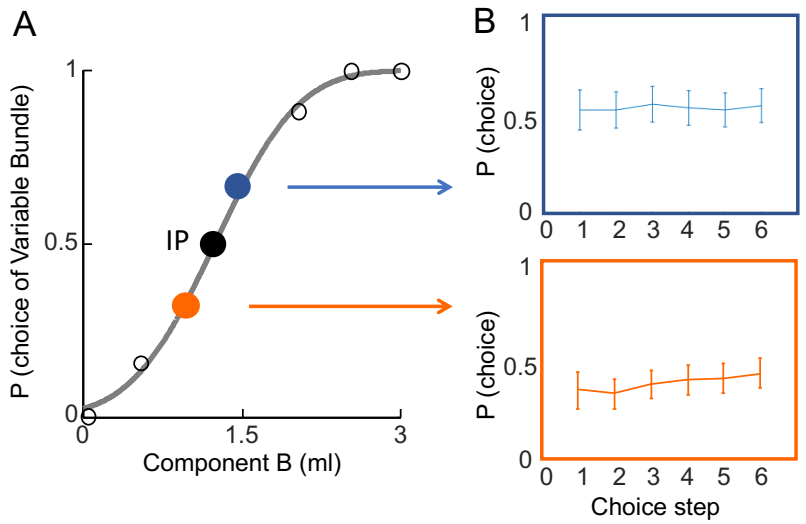

Figure S1. Satiety control.
(A) Choice function obtained by sigmoid fit with the probit model. From seven randomly selected fixed amounts of component B in the Variable Bundle, the two colored dots show two amounts associated with choice probabilities closest to the indifference point (IP).
(B) Lack of satiety across typical test duration. Choice probabilities varied insignificantly above (blue) and below (orange) the IP across six repeated choice steps, using 42 trials at each step. Data were averaged from all 24 participants. The total duration of the six steps was 20 ± 1.38 minutes (mean ± SEM).