



ARTICLE

Received 6 Jun 2016 | Accepted 25 Oct 2016 | Published 15 Nov 2016

DOI: 10.1057/palcomms.2016.82

OPEN

Tracking the decoy: maximizing the decoy effect through sequential experimentation

Maurits C Kaptein¹, Robin Van Emden² and Davide Iannuzzi³

ABSTRACT The decoy effect is one of the best known human biases violating rational choice theory. According to a large body of literature, people may be persuaded to switch from one offer to another by the presence of a third option (the decoy) that, rationally, should have no influence on the decision-making process. For example, when asked to choose between a laptop with a good battery but a poor memory and a laptop with a poor battery but a good memory, customers may be induced to shift their preference if the offer is accompanied by a third laptop that has a battery as good as the latter but even worse memory—an effect that has clear applications in marketing practice. Surprisingly, renowned decoy studies have resisted replication, inducing scholars to challenge the scientific validity of the phenomenon and question its practical relevance. Using a treatment allocation scheme that takes inspiration from the lock-in amplification schemes used in experimental physics, we were able to explore the entire range of decoy attribute values and demonstrate that some of the reproducibility issues reported in the literature result from a suboptimal initial conditions. Furthermore, we demonstrate that our approach is able to sequentially identify the features of the decoy that maximize choice reversal. We thus reinstate the scientific validity and practical relevance of the decoy effect and demonstrate the use of lock-in amplification to optimize treatments.

¹ Statistics and Research Methods, Tilburg University, Tilburg, The Netherlands ² Pavlov Behavioral Science, The Hague, The Netherlands ³ VU University Amsterdam, Department of Physics and Astronomy and LaserLab, Amsterdam, The Netherlands Correspondence: (e-mail: m.c.kaptein@uvt.nl)

Introduction

The decoy effect may well be one of the most famous of human biases (Frederick *et al.*, 2014) that violates neoclassical economics' rational choice theory (Pettibone, 2012). According to a large body of literature (see, for example, Huber *et al.*, 1982; Huber and Puto, 1983; Ratneshwar *et al.*, 1987; Simonson, 1989; Tversky and Kahneman, 1991; Tversky and Simonson, 1992; Ariely and Wallsten, 1995; Heath and Chatterjee, 1995; Wedell and Pettibone, 1996; Sen, 1998; Kivetz *et al.*, 2004; Ariely, 2008; Frederick *et al.*, 2014; Lynn and Yang, 2014), people may be induced to switch from one offer to another by the presence of a third option, the so-called decoy, that rationally should not be of influence in the decision-making process (Huber *et al.*, 1982). The decoy has to be chosen such that it looks alike, yet is slightly inferior, to the target option to impart an effect. Note that the inferiority of the decoy ensures that it is not chosen itself. Under these circumstances, the introduction of a decoy has been shown to increase the choice share of the target option. This decoy (or "asymmetric dominance") effect clearly violates the "axiom of regularity" of rational choice theory: when given a choice set A , where $A \subset B$, the probability of choosing an element x from B cannot be greater than the probability of choosing x from A (Luce, 1959; Yang, 2013).

Although a large number of studies have demonstrated the decoy effect in various settings over the last few decades (see Huber *et al.*, 2014, for an overview), some of these studies have recently resisted replication. While it is not disputed that the decoy effect exists under certain well-described circumstances, recent failures to replicate earlier studies has induced some scholars to challenge the broader scientific importance of the phenomenon and to question its practical relevance (Frederick *et al.*, 2014; Lynn and Yang, 2014).

Tantalized by this controversy, and inspired by the emphasis of both Huber *et al.* (2014) and Simonson (2014) on the importance of correct choice set configuration, we have designed a large-scale study that allowed us to investigate how subtle changes in the decoy offer influence its effect, providing an insight into boundary conditions for successful replications. Note that since it is well-known that strong prior preferences for either the target or the competitor might inhibit the effect of introducing a decoy (Huber *et al.*, 2014), our study focusses explicitly on scenarios where it is likely that the choices are constructed.

Specifically, we applied our novel experimental method, which gravitates around an allocation scheme that takes inspiration from the lock-in amplification schemes used in physics (Meade, 1983; Scofield, 1994). This allowed us to explore the entire range of decoy attribute values, enabling us to conclude that the reproducibility issues reported in the literature are likely the result of a suboptimal choice of the initial conditions. Furthermore, and most relevant for the practical applications of a decoy, we show that our approach can sequentially and automatically identify the features of the decoy option that maximize choice reversal, making our method readily applicable in the digital marketplace.

Possible explanations for the decoy effect. Although the literature on the decoy effect has a strong history, there is as of now little agreement on the mechanisms that cause the observed decision reversal. Numerous explanations have been offered, ranging from highly cognitive and effortful to more automatic (Yang, 2013).

Two general categories of plausible mechanisms have been suggested (Wedell and Pettibone, 1996). *Value shift*-based mechanisms assert that the decoy makes the target seem more attractive, independent of the objective values of either target or competitor (Pettibone and Wedell, 2000). *Value added*

mechanisms on the other hand assert that the decoy makes the target appear the safer choice by raising its value relative to the competitor. An exemplar of the latter is work by Simonson (1989), who demonstrated that a decoy makes it easier for people to justify the choice for the target. However, neither of the mechanisms seems able to fully explain the decoy effect across different moderating conditions. Thus, it is up to future research to provide a more comprehensive explanatory framework (see also the discussion in Yang (2013), or Pettibone and Wedell (2007)). In the current article, however, we do not seek additional explanatory mechanisms, but rather explore the contextual effects of the location of the decoy and practical methods to optimize this location, irrespective of its underlying cause.

The effect of the position of the decoy. The current article focusses on the often discussed yet relatively little-researched explanation for the decoy effect that the probability of the choice reversal is related to the relative positioning of the decoy option on the respective attributes. To illustrate, consider a computer company that wants to increase its market share in the laptop segment against one of its competitors. The company offers a product (the *target*) that comes with a better battery but less memory than the competing one (the *competitor*). As neither product is strictly dominant over the other, personal preference will determine the (population average) probability of choosing the target. Let us now imagine that the company that sells the target product introduces a *decoy* laptop that has the same battery of the target but a less memory. In this situation, it is to be expected that if (for example,) the amount of memory of the decoy is too close to that of the target, the two products may appear as basically identical and the decoy will hardly have an effect on the choice reversal. If on the other hand the amount of memory of the decoy offer is too far from that of the target, potential buyers will not perceive the decoy as a plausible alternative, thereby also rendering the decoy ineffective. Huber *et al.* (2014) identify such effects of the position of the decoy in attribute space as one of the possible explanations for the recent failures to replicate existing decoy studies.

In this article, we limit ourselves to the most researched type of decoy: asymmetrically dominated decoys. An asymmetrically dominated decoy is defined as an option, which is completely dominated by the target on at least one attribute, and where the decoy itself does not possess an attribute that is superior to the target (Lynn and Yang, 2014). That is, we limit ourselves to a region of the attribute space indicated by the shaded rectangle in Fig. 1, left panel.

It is of interest to hypothesize the possible effects of different positions of the decoy in the attribute space depicted in Fig. 1. The decoy literature has, until recently (Trueblood *et al.*, 2014), tended to stay close to the sparse representation as presented in Huber *et al.* (1982) and the right-hand side of Fig. 1. In this literature, different types of decoys are identified: the range decoy (R) is defined as an option, that is slightly weaker than the target on the target's weakest attribute: a range decoy increases the range of an attribute dimension on which the competitor is weakest. The frequency decoy (F) refers to an option that increases the frequency of the attribute dimension on which the target is superior. The range-frequency decoy (denoted RF) is a combination of range and frequency manipulations, and the extreme range (R^*) decoy represents a more extreme version of the range decoy. Importantly, it is demonstrated in (Huber *et al.*, 1982) that different decoy positions result in different sizes of the effect.

Figure 2 represents our hypothesized target choice probabilities as a function of the decoy location for the laptop scenario.

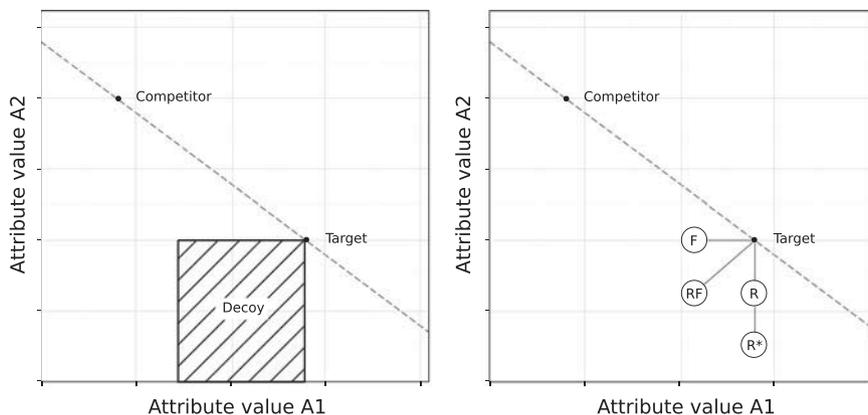


Figure 1 | Illustration of possible sizes of the decoy effect for different decoy positions. Left panel: placement of asymmetrically dominated decoys. Right panel: Positions of range (R), extreme range (R*), frequency (F), and range-frequency decoys (RF).

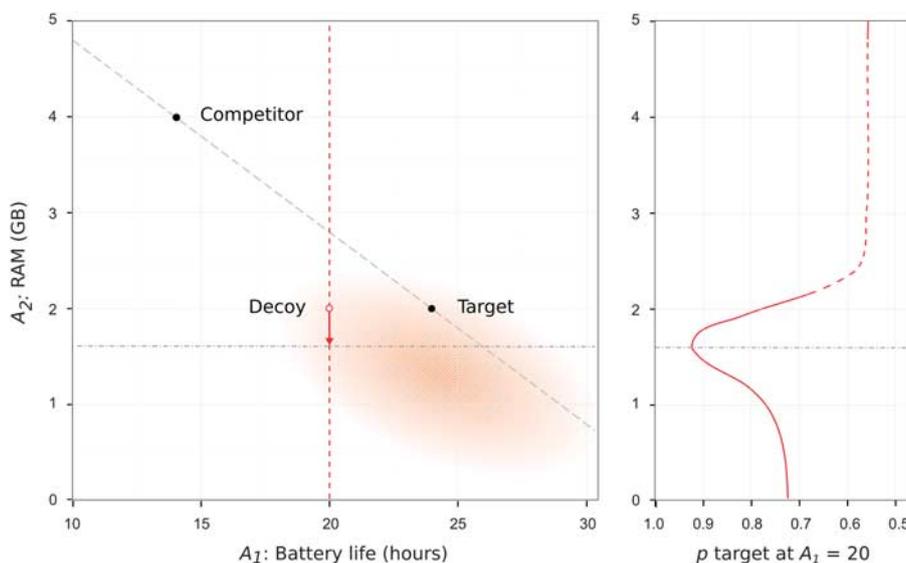


Figure 2 | Density map (left panel) of the possible effect of decoy location in two dimensional attribute space for the laptop scenario. The right panel shows the expected probability of choosing the target as a function of the second attribute.

The figure shows that the competitor *C* scores higher on attribute 2 (the RAM), while the target *T* has a higher score on attribute 1 (the hours of battery life). Introducing a decoy *D*, which is asymmetrically dominated by the target, can lead to an increased preference for *T* compared with *D*. However, the exact position of the *optimal* decoy (that is, the position that maximizes the number of choices for *T* in this space) is unknown. Following Wedell (1991), we hypothesize the effect of the decoy to be first increasing, and then decreasing as the distance between target and decoy passes a threshold value. By superimposing an underlying continuous distribution on the multitude of reported discrete decoy measurements (based on, for example, Huber *et al.*, 1982; Ratneshwar *et al.*, 1987; Simonson, 1989; Tversky and Simonson, 1992; Ariely and Wallsten, 1995; Kivetz *et al.*, 2004; Trueblood *et al.*, 2013) and cross-referencing these with a probability density matrix derived from the Multiattribute Linear Ballistic Accumulator (Trueblood *et al.*, 2014), we constructed Fig. 2. The red dashed vertical line shows different possible placements of the decoy *D* as a function of attribute 2 for a fixed choice of attribute 1; the left panel of the figure shows the hypothesized probability of choosing the target as a function of

attribute 2. The shaded region in the left panel indicates our expectations regarding the effect of the placement of *D* on the probability of choosing *T* (dense colour regions = higher probability).

The current study. On the basis of the above analysis of the existing work, we hypothesize that the replication issues reported in recent literature may be because of a suboptimal positioning of the decoy in attribute space. In the current experiment we replicate a number of decoy studies to investigate this hypothesis, but contrary to the original studies we sample from a broad range of attribute values along one of the dimensions. Hence, we can explore the effect of the decoy *as a function of its position*. Second, to restore the practical applicability of decoy strategies, it is necessary to design a method that can efficiently explore the attribute space and properly position the decoy offer in practical situations. For this purpose, we examine the performance of an algorithm that is routinely used in experimental physics to sequentially position the decoy such that the choice reversal is optimized.

Table 1 | Overview of the five decoy tasks used in this study and their attribute values

Scenario	C		T		D		Consistent?
	A ₁	A ₂	A ₁	A ₂	A ₁	A ₂ [*]	
Laptop	14	4.0	24	2.0	20	2.0	N
Hotel	179.00	5	129.00	15	119.00 [†]	35 [†]	N
Economist	Online	56.95	Both	125.00	Print	125.00 [†]	N
Beer	8.6	70	7.8	50	7.8 [†]	30	Y
Juice	70	2.00	50	1.20	30	1.10 [†]	Y

Method

To settle the debate on the reproducibility of the decoy effect, and to establish a method for the optimal positioning of a decoy offer, we ran an experiment in which a total of $N=7125$ respondents were subjected to five different decoy tasks that have been examined at least twice before in the existing literature. Next to the laptop scenario described above, we included a hotel room scenario with attributes *walking distance* to destination and *price* of the room (Simonson, 1989; Lynn and Yang, 2014), the well-known Economist subscription scenario with attributes *online, print, or combined subscription* and *price* of the subscription (Ariely, 2008; Lynn and Yang, 2014), a 6-pack of beer scenario with attributes *price* and *quality rating* (Huber *et al.*, 1982; Lynn and Yang, 2014), and a frozen orange juice scenario with again a *quality rating* and *price* (Ratneshwar *et al.*, 1987; Frederick *et al.*, 2014). Table 1 highlights whether previous replications have been consistent or not. The table further presents the positions of the competitor C, target T, and decoy D options in attribute space (A₁ and A₂) in the original studies. The * indicates the attribute manipulated in the current study. The † indicates that higher values of this attribute correspond to a lower probability of choosing the target.

A number of respondents ($n=506$) were confronted with a binary choice task without a decoy option to enable estimation of the probability of choosing the target, denoted $p^{(T)}$. For $n=1467$ respondents the value of one of the decoy attributes, $A_2^{(D)}$, was chosen uniformly at random within a reasonable range to enable estimation of the probability $p^{(TD)}$ of choosing the target as a function of the position of the decoy. Finally, for two sets of $n=2556$ and $n=2596$ respondents the value of $A_2^{(D)}$ was optimized sequentially starting from two different initial values. The sequential allocation was done using an experimentation scheme called lock-in amplification (Kaptein and Iannuzzi, 2015) (see details below), which allowed us to approximate the value of the derivative at a specific point of the (unknown) function relating the attribute value to the probability of selecting the target. Using a gradient ascent method (see, for example, Poggio *et al.*, 2011), we could subsequently update the attribute value to sequentially find the position of the decoy that maximizes target selection. This approach, which is similar to that used in physics to stabilize experimental setups (Meade, 1983; Scofield, 1994), is particularly suitable for decoy positioning in applied settings: compared to random exploration, lock-in amplification is robust to high noise, does not require the specification of a functional form, and has the ability to follow drifting variables (which, in our case, translates to possible drifts in the optimal decoy position over time) (Gaber *et al.*, 2005; Kaptein and Iannuzzi, 2015). Thus, the study was set up such that we could replicate five previously replicated decoy studies, extend them by analyzing the effect of the positioning of the decoy along the full range of one of the attributes, and test a sequential allocation scheme that can be used to position decoy offers in practice.

Participants and procedure. Our participants consisted of 7125 crowd-sourced workers from Amazon Mechanical Turk, each completing a Human Intelligence Task (HIT). In total 7134 participants started the survey, but 9 did not fill out any of the questions and were thus excluded from the analysis. Amazon Mechanical Turk is a popular web service that enables so-called “requesters” to pay individuals known as “workers” a fee on completion of an online task. It has proven a reliable tool for enlisting participants in online social research (Eriksson and Simpson, 2010; Horton and Chilton, 2010; Ipeirotis *et al.*, 2010; Paolacci *et al.*, 2010; Buhrmester *et al.*, 2011; Mason and Suri, 2012). MTurks built-in system of qualifications was used to ensure that mainly US-based workers with a HIT approval rate of at least 90% and at least 100+ submitted prior HITS could participate, guaranteeing the quality of the collected data.

We implemented our study by embedding a custom survey platform within the MTurk worker interface. The attribute values of the decoy options in the sequential allocation conditions were dynamically updated using StreamingBandit server running the lock-in feedback algorithm (see below for details of the algorithm) (Kaptein and Iannuzzi, 2015; Kaptein and Kruijswijk, 2016). Hence, in the two sequential optimization conditions—which differed only in their initial starting point—the positioning of the decoy was updated sequentially such as to converge to a position in attribute space that would maximize the probability of choosing the target. Note that the order of the options (target, competition and decoy) within each scenario was randomized for each participant to control for possible effects of

the ordering of the options. In each condition the HIT consisted of an initial page requesting informed consent followed by the decoy scenarios and a short closing survey querying the country of residence, gender, age, and education of respondents. As a reimbursement, the participants received a payment of US \$0.40 through their Amazon Mechanical Turk account. All resulting datasets have been made publicly available in the Dataverse repository (van Emde *et al.* 2016).

Of our participants, 46.4% were female. The majority of our participants were between 25 and 44 years of age (65.3%), and resided in the United States (99.2%). Over 90% completed a high school (or higher) education.

Materials

Figure 3 displays a screenshot, including the exact text, of the laptop scenario as used in this study. The exact descriptions used for the other four scenarios are as follows:

- (Hotel) “You’re going to have a job interview out of town and will need to book a hotel room for the trip. you will not be reimbursed for the hotel room. You find the following deals online, which one would you choose?”
 D: Hotel W: 40 min away from the interview site. \$119/nt.
 C: Hotel J: 5 min away from the interview site. \$179/nt.
 T: Hotel N: 15 min away from the interview site. \$129/nt.
- (Economist) “Welcome to the Economist Subscription Centre. Pick the type of subscription you want to buy or renew”:
 C: Economist.com subscription—\$56.95.
 D: Print subscription—\$143.90.
 T: Print & Web subscription—\$125.00.
- (Beer) “If you had to choose one of the following brand of beer to purchase, which one would it be? Average Quality Rating Scale: (100 = Best, 0 = Worst)”
 D: Price / sixpack: \$7.80—Average Quality Rating: 30
 C: Price / sixpack: \$8.60—Average Quality Rating: 70
 T: Price / sixpack: \$7.80—Average Quality Rating: 50
- (Juice) “Below you will find some brands of frozen concentrated orange juice. You know only the price and the quality ratings made by consumer reports. Given that you had to buy one brand based on this information alone, which one would it be? In case of the quality rating, 100 = ideal.”
 D: Brand W: Price per can \$1.15. Quality rating 30
 T: Brand N: Price per can \$1.20. Quality rating 50
 C: Brand J: Price per can \$2.00. Quality rating 70

The lock-in amplification algorithm. Here we detail the idea of lock-in amplification in continuous time. Let’s assume that y , the probability of choosing the decoy, is a continuous function of x : $y=f(x)$ (in our case, $p^{(TD)}=f(A_2)$). Let’s further assume that x oscillates with time according to:

$$x(t) = x_0 + A\cos(\omega t) \tag{1}$$

where ω is the angular frequency of the oscillation, x_0 its central value, and A its amplitude. For relatively small values of A , Taylor

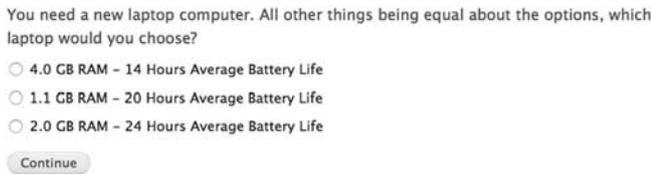


Figure 3 | Screenshot of the laptop decoy scenario as used in this study.

expanding $f(x)$ around x_0 to the second order, one obtains:

$$y(x(t)) = f(x_0) + (x_0 - x_0 - A\cos(\omega t)) \left(\frac{\partial f}{\partial x} \Big|_{x=x_0} \right) + \frac{1}{2} \left(x_0 - x_0 - A\cos(\omega t) \right)^2 \left(\frac{\partial^2 f}{\partial x^2} \Big|_{x=x_0} \right) \tag{2}$$

which can be simplified to:

$$y(x(t)) = k - A\cos(\omega t) \left(\frac{\partial f}{\partial x} \Big|_{x=x_0} \right) + \frac{1}{4} A^2 \cos(2\omega t) \left(\frac{\partial^2 f}{\partial x^2} \Big|_{x=x_0} \right) \tag{3}$$

where $k = f(x_0) + 1/4(\partial^2 f/\partial x^2|_{x=x_0})$. Thus, for small oscillations y becomes the sum of three terms: a constant term, a term oscillating at angular frequency ω , and a term oscillating at angular frequency 2ω .

Suppose we can actively *manipulate* x and *measure* y , and that f is continuous and only has one maximum and no minimum. Further suppose we are interesting in finding $x_{max} = \text{argmax}_x f(x)$ and that our measurements contain noise:

$$y(t) = f(x(t)) + \epsilon \tag{4}$$

where ϵ denotes the noise and $\epsilon \sim \pi()$, where π is some probability density function and $E[\epsilon|x] = 0$. In physical lock-in amplifiers (Scofield, 1994), one multiplies the observed y variable by $\cos(\omega t)$. Using equation (3) and equation (4), we obtain:

$$y_\omega(t) = \cos(\omega t) \left[k - A\cos(\omega t) \left(\frac{\partial f}{\partial x} \Big|_{x=x_0} \right) + \frac{1}{4} A^2 \cos(2\omega t) \left(\frac{\partial^2 f}{\partial x^2} \Big|_{x=x_0} \right) + \epsilon \right] \tag{5}$$

where y_ω is the value of y after it has been multiplied by $\cos(\omega t)$. This can be written more compactly as:

$$y_\omega = -\frac{A}{2} \left(\frac{\partial f}{\partial x} \Big|_{x=x_0} \right) + k_{2\omega} \cos(\omega t) + k_{2\omega} \cos(2\omega t) + k_{3\omega} \cos(3\omega t) + \epsilon \cos(\omega t) \tag{6}$$

where

$$k_\omega = k + A^2/8 \left(\frac{\partial^2 f}{\partial x^2} \Big|_{x=x_0} \right), \tag{7}$$

$$k_{2\omega} = -A/2 \left(\frac{\partial^2 f}{\partial x^2} \Big|_{x=x_0} \right), \tag{8}$$

$$k_{3\omega} = A^2/8 \left(\frac{\partial^2 f}{\partial x^2} \Big|_{x=x_0} \right). \tag{9}$$

Integrating y_ω over a time $T = (2\pi N)/(\omega)$, where N is a positive integer and T denotes the time needed to integrate N full oscillations, one obtains:

$$y_\omega^* = -\frac{TA}{2} \left(\frac{\partial f}{\partial x} \Big|_{x=x_0} \right) + \int_0^T \epsilon \cos(\omega t) \tag{10}$$

Depending on the noise level, one can tailor the integration time, T , in such a way to reduce the second addendum of the right hand of equation (10) to negligible levels, effectively averaging out the noise in the measurements. Under these circumstances, y_ω^* provides a direct measurement of the value of the first derivative of f at $x = x_0$. This provides a logical update strategy for x_0 : if $y_\omega^* < 0$, then x_0 is larger than the value of x that maximizes f ; likewise, if $y_\omega^* > 0$, x_0 is smaller than the value of x that maximizes f . Thus, based on the oscillation observed in y_ω we are now able to move x_0 closer to $x = \text{argmax}_x f(x)$ using an update rule $x_0 := x_0 + \gamma y_\omega^*$, where γ quantifies the learn rate of the procedure (Kaptein and Iannuzzi, 2015). Note that the lock-in amplification algorithm depends on a starting value, x_0 , an amplitude A , a learn-rate γ , and an integration time T . The exact values as used in the experiment are presented in Table 2. Also, the actual discrete algorithm used in the studies is described in the Supplementary Materials.

Ethics committee psychology (CEP) approval. The Psychology Research Ethics Committee (in Dutch: “Commissie Ethiek Psychologie”, CEP) considers applications for ethical approval for research conducted within the Institute of Psychology of Leiden University, the Netherlands. The Psychology Ethics Committee has reviewed and approved the current research under the title of “Decoy Amazon Mechanical Turk Study—2016”. CEP code: CEP16-0203/42. Contact ethiekpsychologie@fsw.leidenuniv.nl for more information.

Results

In this section, we describe the results of our experiment. First, we detail how we estimate, based on the data, the quantities of interest (for example, \hat{p}^T , \hat{p}^{T_D} , and $\hat{p}^{T_D}(A_2)$). Subsequently, we present the results and discuss both the effects of the positioning of the decoy on the probability of selecting the target, as well as the performance of the lock-in amplification algorithm to sequentially position the decoy. Note that the main results are presented in Fig. 4, which shows the estimated probability of selecting the target when the decoy is absent, $\hat{p}^{(T)}$, the probability of selecting the target when the decoy is present as a function of its position, $\hat{p}^{(T_D)}$, the values of C and T on the second attribute, and the optimal position of the decoy, $p_{max}^{(T_D)}$, as uncovered by lock-in feedback. Also, the figure shows the convergence path(s) of the lock-in feedback for each scenario using two different starting points in the horizontal lower panels.

Table 2 | Values of the tuning parameters of the lock-in amplification algorithm as used in the experiment

Scenario	A_1^{orig}	A_2^{alt}	A	T	γ
Laptop	2.0	1.0	0.4	150	0.06
Hotel	35	55	2.8	150	0.06
Economist	125	160	18.0	150	0.06
Beer	30	10	7.0	150	0.06
Juice	1.10	1.80	0.22	150	0.06

Estimation procedure. Before discussing the results, we first detail our (Bayesian) estimation procedure for both \hat{p}^T and $\hat{p}^{T_D}(A_2)$ as presented in Fig. 4. To estimate \hat{p}^T we assume the observed number of choices for the target to follow a binomial distribution with parameters θ , $n = s + f$, where the former is the probability of a success and the latter the total number of trials (here we exclude, as is standard in the decoy literature, trials in which the decoy itself was selected (Lynn and Yang, 2014)). We use s to refer to the number of times the target was selected, f the number of times the competitor was selected, and n the total number of available observations in this condition. We thus have a binomial sampling distribution:

$$p(s|\theta, n) = \frac{n!}{s!(n-s)!} \theta^s (1-\theta)^{n-s} = \text{Binom}(\theta, n).$$

Next, we use a fairly uninformative, uniform prior² on θ :

$$p(\theta|\alpha, \beta) = \text{Unif}[0, 1] = \text{Beta}(\alpha = 1, \beta = 1) \\ = \frac{1}{B(\alpha, \beta)} \theta^{\alpha-1} (1-\theta)^{\beta-1}$$

where $B()$ is the Beta function (for example, Wilcox, 1981). We obtain the posterior:

$$p(\theta|s, n, \alpha, \beta) = \frac{p(s|\theta, n)p(\theta|\alpha, \beta)}{p(s|n, \alpha, \beta)} \\ = \text{Beta}(\alpha = s + 1, \beta = f + 1)$$

(see, for example, Gelman *et al.*, 2013, for a derivation)). The estimates of \hat{p}^T presented in Fig. 4 are the 2.5th percentile, the expected value, and the 97.5th percentile of this posterior distribution of θ .

For the estimation of $\hat{p}^{T_D}(A_2)$ we assume that each individual observation $y_i \in \{0, 1\}$ (where $\sum_{i=1}^n y_i = s$) is sampled from a Bernoulli(θ), and θ depends on the value of A_2 via a standard probit link

$$\theta = \Phi(\beta_0 + \beta_1 A_2 + \beta_2 A_2^2 + \dots + \beta_k A_2^k) \tag{11}$$

where $\Phi()$ is the standard normal CDF. We examined different orders k of the polynomial (appropriately centring the data): Fig. 4 shows the results for $k=4$, at which point increasing k does not seem to visually improve model fit. Using the data augmentation scheme as described in Albert and Chib (1993), and implemented in MCMCpack (Martin *et al.*, 2011), we obtain $m = 10000$ draws from the posterior distribution $p(\beta|y, A_2)$ using an improper flat prior for β : $p(\beta) \propto 1$. Given the m draws of β , we obtain m draws from the posterior predictive distribution $p(\theta|\beta, y, A_2)$ for each of 100 values of A_2 spaced equally within the ranges displayed in Fig. 4 for each scenario. The figure presents the 2.5th percentile of these draws, their mean, and the 97.5th percentile.

Explaining the replication failures: the positioning of the decoy. Figure 4 shows the results for each of the five scenarios. Our results readily provide an explanation for the reason why the

laptop scenario has been inconsistent in prior replications: the position of the decoy in the original studies (the vertical dashed line labelled A_2^{orig}) is such that the probability of choosing the target in the absence or in the presence of the decoy is too close to trigger a noticeable decision reversal. Hence, while the effect has been observed in the initial study it is no surprise that it was relatively small and inconsistent. Our data however also shows that, for smaller values of $A_2^{(D)}$, the difference between the offers becomes sufficiently large to generate a clear decoy effect.

For the hotel and the Economist scenarios—which also previously failed to replicate—a similar pattern emerges: The original decoy was positioned such that the effects were small. In both cases, however, it is clear that larger effects can be obtained by moving the decoy to higher values of the attribute. Especially the Economist scenario seems to be ill-positioned in the original studies and, while this scenario potentially provides a convincing demonstration of the effect, its replication failure under the initial context is not surprising. The hotel scenario was also presumably not replicated in the previous literature because of a poorly chosen decoy position, as it seems that only much higher values on $A_2^{(D)}$ may potentially demonstrate a decoy effect. For the beer and the juice scenarios, which have been replicated consistently, the original positioning of the decoy was such that the effects were relatively large, and hence were easily replicated. Although in neither case the decoy seems to be positioned at its absolute maximum, the effects are large enough that it is no surprise that these scenarios provided consistent results in previous studies.

Finding optimal decoy values in practice: lock-in feedback.

Next to highlighting the contextual effects that caused previous replication failures, it is clear from the top panel in Fig. 4 that the lock-in amplification scheme is able to sequentially bring the decoy offer towards the position where the effect is maximized—an extremely important feature from both a practical and a methodological perspective. Hence, next to demonstrating why a number of replications might have failed, the current study demonstrates a readily usable method to introduce and position decoy offers in practice. Clearly, in all of the scenarios, the effect of the decoy could have been demonstrated more strongly if the selection of the decoy attributes had been guided by a lock-in feedback.

Discussion

The fact that the decoy effect has occasionally resisted replication has led some to refrain from using decoys in practice, and even question the scientific importance of the phenomenon (Frederick *et al.*, 2014). Our study, however, confirms that subtle changes in decoy position are of great importance for the strength of the effect, supporting the emphasis of both Huber *et al.* (2014) and Simonson (2014) on the context-dependence of choice behaviour and the significance of correct choice set configuration. This result also supports the initial identification of Huber *et al.* (1982) of a limited area in attribute space for which the decoy effect occurs.

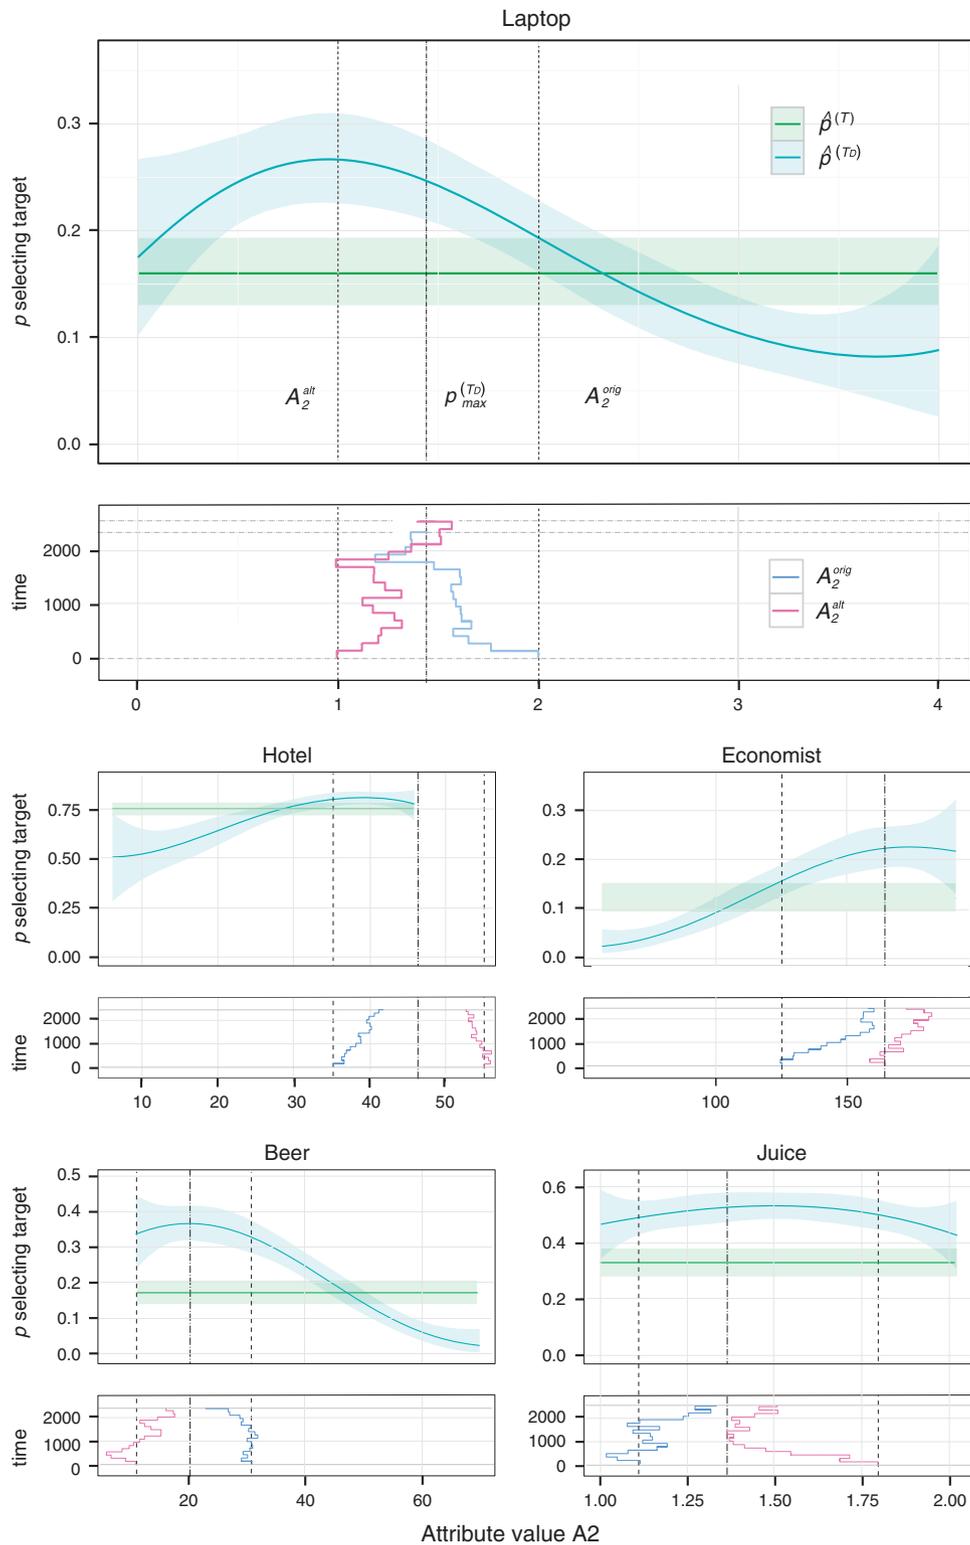


Figure 4 | Results of our decoy experiment. Displayed are the estimated probabilities of selecting the target as a function of attribute 2 for different conditions and their credible intervals.

By not only demonstrating the effect of the positioning of a decoy offer, but also providing a method to “lock into” an optimal decoy position, we present future researchers the option to pin down this, and potentially other, quantifiable aspects of online choice behaviour. It seems to us to be of great practical and scientific use to be able to determine optimal stimuli positions

sequentially, with the potential to isolate other contextual influences as the optimal stimuli positions might drift over time. Our proposed method is thus of relevance for future studies into the decoy effect and beyond.

Furthermore, although it has been suggested that decoys are hard to use in practice (Huber *et al.*, 2014), our sequential

optimization method reaffirms its potential use in digital marketing practice. In contrast to offline stores, online shops do not need to produce decoy goods to make use of the effect. They can offer clear numeric, easily optimizable representations of any two product dimensions, such as a price and a quality rating. As such, we expect self-optimizing decoys to be a valuable addition to online dynamic pricing technology (for example, Kannan and Kopalle, 2001; Jiang *et al.*, 2011) and recommender systems (Teppan and Felfernig, 2012). As can be deduced from experiments such as our beer scenario, a correctly positioned online decoy can change a target's choice share by multiple percentage points—an impressive feat. We also see ample room for further research, for example by extending our work to broader human decision making related paradigms such as Multi-Attribute Utility Theory-based modelling (Von Winterfeldt and Edwards, 1986; Felfernig *et al.*, 2008).

To conclude, we have shown that failures to replicate the decoy effect in certain scenarios are likely due to a suboptimal position of the decoy offer in attribute space. Furthermore, we demonstrated that sequential optimization methods borrowed from physics can ensure a correct decoy position. It is reasonable to assume that similar algorithms will soon be used to address other open issues in the social sciences in both theoretical analysis and practical applications.

Notes

- 1 We also analyzed three additional scenarios that had not been replicated before. The results of these can be found in the Supplementary Materials. The analysis of these results shows that, in two cases, no decoy effect should be expected under any circumstances, as the change of the decoy attribute values does not affect the decision process. For the third of these cases, we find a pattern very similar to that reported upon here.
- 2 Note that a prior for which $\alpha, \beta \rightarrow 0$ would be less informative, however, given the sample size $n, n \approx 1500$ the influence of the prior is neglectable.

References

- Albert JH and Chib S (1993) Bayesian analysis of binary and polychotomous response data. *Journal of the American Statistical Association*; **88** (422): 669–679.
- Ariely D (2008) *Predictably Irrational: The Hidden Forces That Shape Our Decisions*. Harper Collins: New York.
- Ariely D and Wallsten TS (1995) Seeking subjective dominance in multi-dimensional space: An explanation of the asymmetric dominance effect. *Organizational Behavior and Human Decision Processes*; **63** (3): 223–232.
- Buhrmester M, Kwang T and Gosling SD (2011) Amazon's mechanical Turk a new source of inexpensive, yet high-quality, data? *Perspectives on Psychological Science*; **6** (1): 3–5.
- Eriksson K and Simpson B (2010) Emotional reactions to losing explain gender differences in entering a risky lottery. *Judgment and Decision Making*; **5** (3): 159.
- Felfernig A, Gula B, Leitner G, Maier M, Melcher R and Teppan E (2008) Persuasion in knowledge-based recommendation. In: *International Conference on Persuasive Technology*. Springer: Berlin, Heidelberg, pp. 71–82.
- Frederick S, Lee L and Baskin E (2014) The Limits of Attraction. *Journal of Marketing Research*; **51** (4): 487–507.
- Gaber MM, Zaslavsky A and Krishnaswamy S (2005) Mining data streams. *ACM SIGMOD Record*; **34** (2): 18.
- Gelman A, Carlin JB, Stern HS, Dunson DB, Vehtari A and Rubin DB (2013) *Bayesian Data Analysis*, Third Edition, CRC Press: New York.
- Heath TB and Chatterjee S (1995) Asymmetric decoy effects on lower-quality versus higher-quality brands: Meta-analytic and experimental evidence. *Journal of Consumer Research*; **22** (3): 268–284.
- Horton JJ and Chilton LB (2010) The labor economics of paid crowdsourcing. In: *Proceedings of the 11th ACM conference on Electronic commerce* ACM: New York, pp. 209–218.
- Huber J, Payne JW and Puto C (1982) Adding asymmetrically dominated alternatives: Violations of regularity and the similarity hypothesis. *Journal of Consumer Research*; **9** (1): 90.
- Huber J, Payne JW and Puto CP (2014) Let's be honest about the attraction effect. *Journal of Marketing Research*; **51** (4): 520.
- Huber J and Puto C (1983) Market boundaries and product choice: Illustrating attraction and substitution effects. *Journal of Consumer Research*; **10** (1): 31–44.
- Ipeirotis PG, Provost F and Wang J (2010) Quality management on Amazon mechanical Turk. In: *Proceedings of the ACM SIGKDD Workshop on Human Computation* ACM: New York, pp. 64–67.
- Jiang Y, Shang J, Kemerer CF and Liu Y (2011) Optimizing E-tailer profits and customer savings: Pricing multistage customized online bundles. *Marketing Science*; **30** (4): 737–752.
- Kannan PK and Kopalle PK (2001) Dynamic pricing, price expectations, pricing in E-Commerce, reference prices. *International Journal of Electronic Commerce*; **3** (5): 63–83.
- Kaptein M and Iannuzzi D (2015) Lock in feedback in sequential experiments. *CoRR*, <http://arxiv.org/abs/1502.00598>, accessed 3 December 2015.
- Kaptein M and Kruijswijk J (2016) Streamingbandit: Developing adaptive persuasive systems. *CoRR*, <http://arxiv.org/abs/1602.06700>, accessed 19 May 2016.
- Kivetz R, Netzer O and Srinivasan V (2004) Alternative Models for Capturing the Compromise Effect. *Journal of Marketing Research*; **41** (3): 237–257.
- Luce RD (1959) *Individual Choice Behavior a Theoretical Analysis*. John Wiley and sons: New York.
- Lynn M and Yang S (2014) More evidence challenging the robustness and usefulness of the attraction effect. *Journal of Marketing Research*; **51** (4): 3–12.
- Martin AD, Quinn KM and Park JH (2011) MCMCpack: Markov chain Monte Carlo in R. *Journal of Statistical Software*; **42** (9): 22.
- Mason W and Suri S (2012) Conducting behavioral research on Amazon's mechanical Turk. *Behavior Research Methods*; **44** (1): 1–23.
- Meade ML (1983) *Lock-in Amplifiers: Principles and Applications*. Mike Meade: Peter Perigrinus, London, UK.
- Paolacci G, Chandler J and Ipeirotis PG (2010) Running experiments on Amazon mechanical Turk. *Judgment and Decision Making*; **5** (5): 411–419.
- Pettibone JC (2012) Testing the effect of time pressure on asymmetric dominance and compromise decoys in choice. *Judgment and Decision Making*; **7** (4): 513.
- Pettibone JC and Wedell DH (2000) Examining models of nondominated decoy effects across judgment and choice. *Organizational Behavior and Human Decision Processes*; **81** (2): 300–328.
- Pettibone JC and Wedell DH (2007) Testing alternative explanations of phantom decoy effects. *Journal of Behavioral Decision Making*; **20** (3): 323–341.
- Poggio T, Voinea S and Rosasco L (2011) Online learning, stability, and stochastic gradient descent. *Artificial Intelligence*; **8** (11): 11.
- Ratneshwar S, Shocker AD and Stewart DW (1987) Toward understanding the attraction effect: The implications of product stimulus meaningfulness and familiarity. *Journal of Consumer Research*; **13** (4): 520–533.
- Scofield JH (1994) Frequency-domain description of a lock-in amplifier. *American Journal of Physics*; **62** (2): 129.
- Sen S (1998) Knowledge, information mode, and the attraction effect. *Journal of Consumer Research*; **25** (1): 64–77.
- Simonson I (1989) Choice based on reasons: The case of attraction and compromise effects. *Journal of Consumer Research*; **16** (2): 158.
- Simonson I (2014) Vices and virtues of misguided replications: The case of asymmetric dominance. *SSRN Electronic Journal*; **1904** (1963): 185–206.
- Teppan E and Felfernig A (2012) Minimization of decoy effects in recommender result sets. *Web Intelligence and Agent Systems: An International Journal*; **10** (4): 385–395.
- Trueblood JS, Brown SD and Heathcote A (2014) The multiattribute linear ballistic accumulator model of context effects in multialternative choice. *Psychological Review*; **121**, (2): 179–205.
- Trueblood JS, Brown SD, Heathcote A and Busemeyer JR (2013) Not just for consumers: Context effects are fundamental to decision making. *Psychological Science*; **24** (6): 901–908.
- Tversky A and Kahneman D (1991) Loss aversion in riskless choice: A reference-dependent model. *The quarterly journal of economics*; **106** (4): 1039–1061.
- Tversky A and Simonson I (1992) Choice in context: Tradeoff contrast and extremeness aversion. *Journal of Marketing Research*; **29** (3): 281–296.
- van Emden R, Kaptein M and Iannuzzi D (2016) Replication data for: Tracking the decoy—maximizing the decoy effect through sequential experimentation. Dataverse, <http://dx.doi.org/10.7910/DVN/FCHU0j>, Accessed 13 October 2016.
- Von Winterfeldt D and Edwards W (1986) *Decision analysis and behavioral research*. Cambridge University Press: Cambridge, NY.
- Wedell DH (1991) Distinguishing among models of contextually induced preference reversals. *Journal of Experimental Psychology: Learning, Memory, and Cognition*; **17** (4): 767.
- Wedell DH and Pettibone JC (1996) Using judgments to understand decoy effects in choice. *Organizational Behavior and Human Decision Processes*; **67** (3): 326–344.
- Wilcox RR (1981) A review of the beta-binomial model and its extensions. *Journal of Educational Statistics*; **6**, 3–32.
- Yang SS (2013) The attraction effect: An overview, its fragility, and a meta-analysis, PhD Thesis, Cornell University: Cornell.

Data availability

The source code, the scripts used to generate the statistics, the estimation scripts and the datasets analyzed in the current study are available in the Dataverse repository: <http://doi:10.7910/DVN/FCHU0J>.

Acknowledgements

D.I. acknowledges the support of the European Research Council (grant agreement n. 615170) and of the Stichting voor Fundamenteel Onderzoek der Materie (FOM), which is financially supported by the Netherlands organization for scientific research (NWO). We would like to thank Dr Dean Eckles, Dr Joris Mulder and Prof. Dr Eric van Dijk for their comments on earlier versions of this manuscript and Dr Giuseppe Ruoso, Dr Grzegorz Gruca for useful discussions.

Additional information

Supplementary Information: accompanies this paper at <http://www.palgrave-journals.com/palcomms>

Competing interests: The Authors declare no competing financial interests.

Reprints and permission information is available at http://www.palgrave-journals.com/pal/authors/rights_and_permissions.html

How to cite this article: Kaptein MC, van Emden R and Iannuzzi D (2016) Tracking the decoy: maximizing the decoy effect through sequential experimentation. *Palgrave Communications*. 2:16082 doi: 10.1057/palcomms.2016.82.



This work is licensed under a Creative Commons Attribution 4.0 International License. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in the credit line; if the material is not included under the Creative Commons license, users will need to obtain permission from the license holder to reproduce the material. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>