

Attention Please!*

Olivier Gossner

CREST, CNRS, École Polytechnique
and London School of Economics

Jakub Steiner

University of Zurich and CERGE-EI

Colin Stewart

University of Toronto

April 15, 2019

Abstract

We study the impact of manipulating the attention of a decision-maker who learns sequentially about a number of items before making a choice. Under natural assumptions on the decision-maker’s strategy, directing attention toward one item increases its likelihood of being chosen regardless of its value. This result applies when the decision-maker can reject all items in favor of an outside option with known value; if no outside option is available, the direction of the effect of manipulation depends on the value of the item.

keywords: sequential sampling, marketing, persuasion, attention

JEL codes: D8, D91

1 Introduction

The struggle for attention is a pervasive phenomenon. Attention-seeking behavior plays an important role in advertising, finance, industrial organization, psychology, and biology.¹

*We have benefited from comments of Ian Krajbich, Arina Nikandrova, Marco Ottaviani, and seminar and conference participants at Bern, Caltech, Cergy-Pontoise, Northwestern, SFU, Vienna, Warwick, Madrid. Pavel Ilinov, Jan Sedek, and Jiaqi Zou provided excellent research assistance. This work was financially supported by the French National Research Agency (ANR), “Investissements d’Avenir” (ANR-11-IDEX0003/LabEx Ecodec/ANR-11-LABX-0047) (Gossner), ERC grant 770652 (Steiner), and by the Social Sciences and Humanities Research Council of Canada (Stewart).

¹See, e.g., Fogg-Meade (1901) for advertising, Lou (2014) for finance, Eliaz and Spiegel (2011) for industrial organization, Orquin and Loose (2013) for psychology, and Dukas (2002) for biology.

The main message is consistent across fields: drawing attention toward an item increases its demand.

The existing literature provides two main explanations for how grabbing attention can increase demand.² One is that they directly affect preferences. While difficult to disprove, a theory of changing preferences offers limited predictive power and makes welfare analysis challenging. The other major explanation is that attention-grabbing behavior itself conveys information—either directly or through signaling—and thereby changes beliefs. But this second channel alone does not suffice to explain the empirical evidence. In fact, there is a sizable body of evidence showing that manipulating attention has a direct influence on demand even when devoid of information.³

We identify a mechanism through which grabbing attention increases demand without influencing preferences or changing the information available to the decision-maker. In our model, a decision-maker learns sequentially about the quality of a number of items by allocating her attention to one of the items in each period. Paying attention to an item generates a noisy signal about its value. Due to cognitive limitations, the decision-maker can focus only on one item at a time; while she pays attention to a given item, her belief about its value evolves stochastically, while her beliefs about the other items remain the same. Once she is sufficiently certain of an optimal choice, she stops learning and either chooses one of the items or an outside option. Starting from a given strategy governing the decision-maker’s attention, we introduce an attention-grabbing manipulation that induces the decision-maker to focus on one “target” item for a fixed duration. We show that, under general conditions, such a manipulation increases demand for the target item.

For the most part, we focus on a setting with binary values, in which each item is either good or bad; we extend our results to allow for more than two values in Section 6.⁴ The decision-maker can choose one of the items or an outside option of known value that she prefers to a bad item but not to a good one. In the context of consumer marketing, we think of the items as substitutable brands of a good, and the outside option as the choice not to purchase any item. We impose a simple stopping rule: the decision-maker ceases to learn about an item once she is sufficiently certain of its value. If she believes an item is sufficiently likely to be good, she stops learning and chooses that item. If she believes an item is sufficiently likely to be bad, she continues to learn about the other items until she finds one that is likely to be good or determines that all items are likely to be bad (in which case she chooses the outside option).

²See Bagwell (2007) for a survey in the context of advertising.

³See, e.g., Chandon et al. (2009), Krajbich and Rangel (2011), or, for a survey, Orquin and Loose (2013).

⁴Binary-value models are common in the literature on sequential sampling; see, e.g., Wald (1945) or, for more recent work, Che and Mierendorff (2017) and Morris and Strack (2017).

The decision-maker’s learning is governed by an *attention strategy* that maps beliefs in any given time period to a (possibly random) item of focus. An attention strategy generates, for each profile of values of the items, a stochastic process over beliefs and items of attention, and a probability that each item is chosen. We refer to these probabilities as *interim demands* for the items. The impact of attention manipulation is captured by the difference in interim demands under the *baseline* and *manipulated* attention strategies, where the baseline strategy is the one the decision-maker employs in the absence of manipulation, while the manipulated strategy induces the decision-maker to focus on a *target* item for a fixed duration, after which she returns to her baseline strategy.

We show that manipulation of attention increases demand and decreases the time to decision in favor of the target item. Also, when the manipulated strategy leads to the choice of an item other than the target, it takes longer to do so than does the baseline strategy. These results hold for any realization of the items’ values. In particular, manipulating attention increases demand even if the target item is bad.

The key to understanding the effect of manipulation is to consider the path of learning for each possible realization of the sequence of signals for each item. Given such a realization, we can view an attention strategy as selecting, in each period, an item for which to uncover one more step along the sequence. The choice of item at the end of the process can be thought of as resulting from a kind of approval contest: the decision-maker continues to learn until she approves of one of the items, or until she finds all of them to be unworthy of approval. For a given realization of signals, there may be multiple items the decision-maker would approve of were she to pay enough attention to them. The choice then comes down to which of these items she approves of first. Directing attention toward one of these items accelerates the process of approval for this item while slowing it down for the other items. Consequently, the likelihood that the target item is chosen increases.

This simple intuition ignores the significant complication that manipulating attention generally affects future attention choice. It could happen, then, that manipulation toward a target item leads to a path along which the decision-maker pays much less attention to the target item afterwards, more than compensating for the direct effect of increased attention. If there are only two items (not including the outside option), then this cannot happen: our results hold regardless of the baseline attention strategy. With more than two items, we require two additional assumptions. First, the attention strategy should be stationary: focus in each period must depend only on the current belief, not on the current time. The second assumption is a form of *independence of irrelevant alternatives* (IIA): conditional on not focusing on an item i , the probability of focusing on each other item is independent of the belief about the value of item i (though it may depend on the beliefs about items other

than i). Together, these two assumptions allow us to consider learning about the target item separately from learning among the remaining items, effectively reducing the problem to one with two items.

To formalize these intuitions, we rely on a technique known in probability theory as *coupling*. In short, we construct a joint probability space in which we fix, for each item, the outcome of the learning process that would arise if the decision-maker focused only on that item. We refer to a profile of realizations of these learning processes across items as a *draw*. We show that, for *every* draw, manipulating attention toward an item both increases demand for that item and decreases decision time if that item is chosen.

Both stationarity and IIA are needed for our results in the sense that their conclusions do not hold if we dispense with either assumption; we provide counterexamples in Section 4. Both assumptions, though restrictive, are automatically satisfied if the attention strategy is optimized given the stopping rules in our model: we prove that, for a general class of learning costs, optimal strategies have a (stationary) Gittins index structure as in the theory of multi-armed bandits. It follows that these strategies are stationary and satisfy IIA.

The presence of an outside option plays an important role in our analysis. Such an outside option appears naturally in consumer choice as the option not to purchase any item, and in finance as the option to invest in a risk-free asset. If there were no outside option, the decision-maker could choose by a process of elimination rather than approval; that is, she could seek to eliminate items that she believes to be bad and ultimately choose an item—the last one remaining—with little knowledge of its value. In this case, manipulating attention toward an item may increase the chance that it is eliminated before the other items, thereby decreasing the demand for it. In the limit as learning becomes arbitrarily precise, we find that, for any stationary attention strategy that satisfies IIA, manipulation increases demand for the target item if its value is high, and decreases demand if its value is low.⁵

When each item can take on more than two values, the presence of an outside option is no longer sufficient to generate our results. For example, with two items, it could be that the decision-maker is confident both are better than the outside option, making the problem effectively the same as one with no outside option. However, our results go through as long as the decision-maker stops and chooses an item only when her belief about that item falls within a given set. We think of this set as consisting of those beliefs at which the decision-maker is sufficiently certain about the item’s value. Requiring a degree of certainty about the chosen item is natural if its value affects subsequent decisions. In the binary values case, then, the key role played by the outside option is that it makes the decision-maker choose

⁵This is reminiscent of the finding in Armel et al. (2008) that manipulating attention tends to increase the choice probability for “appetitive” items and decrease the probability for “aversive” ones.

an item only when she is confident that it is good.

Our result is robust to many aspects of the learning process. The information structure for each item is general, allowing for any number of signal realizations and dependence on the current belief about the item. The decision-maker need not be Bayesian; we can, for instance, reinterpret her beliefs as intensities of accumulated neural stimuli in favor of each item (as, e.g., in Platt and Glimcher, 1999), which can evolve according to an arbitrary stationary Markov process (not necessarily corresponding to Bayesian updating). We also allow for attention strategies that are not optimal, making our results robust with respect to the structure of the attention costs if these strategies were chosen optimally.

Related literature Evidence that increased attention boosts demand comes from several fields. In marketing, Chandon et al. (2009) show that drawing attention to products—for instance, with large displays or placement at eye level—increases demand. In finance, Seasholes and Wu (2007) show that attention-grabbing events about individual stocks increase demand for them. In biology, Yorzinski et al. (2013) study the display strategies through which peacocks grab and retain the attention of peahens during courtship. In each of these contexts, the decision-maker has an outside option not to choose any available item; our results suggest that this is an important feature driving the effect of attention-grabbing behavior on demand.

Our assumptions on attention allocation are rooted in psychology. Though humans are able to pay attention to multiple stimuli simultaneously, such division of attention is difficult, especially when the stimuli are similar to each other (e.g., Spelke, Hirst, and Neisser, 1976). Psychologists distinguish between exogenous and endogenous attention, where the first is beyond the decision-maker’s control and is triggered by sudden movements, bright colors and such, while endogenous attention shifts are controlled by the decision-maker (Mayer et al., 2004). We can interpret attention allocation during our manipulation window as being exogenous, whereas our baseline attention strategy fits the endogenous attention interpretation.

Our model builds on a long tradition in statistics and economic theory originating in Wald (1945), who proposed a theory of optimal sequential learning about a single binary state. A growing literature studies optimal sequential learning about several options when attention must focus on one item at a time (Mandelbaum, Shepp, and Vanderbei, 1990; Ke, Shen, and Villas-Boas, 2016; Ke and Villas-Boas, 2019; Nikandrova and Pancs, 2018; Austen-Smith and Martinelli, 2018). The structure of the optimal learning strategy varies depending on the costs and information structure. Our results on the impact of attention manipulation are independent of these considerations; however, relative to this literature,

we make simplifying assumptions on the rules that govern termination of learning. In a different vein, Che and Mierendorff (2017) study sequential allocation of attention between two Poisson signals about a binary state. In contrast, in our model, the decision-maker chooses among signals about multiple independent states.

In the drift-diffusion model of Ratcliff (1978), a decision-maker tracks the difference in the strength of supporting evidence between two actions, making a choice when this difference becomes sufficiently large.⁶ Krajbich, Armel, and Rangel (2010) explicitly incorporate attention choice in this model, and introduce an exogenous bias in the accumulated signal toward the item on which the decision-maker is currently focusing. This extended drift-diffusion model accommodates empirical findings showing that exogenous shifts in attention tend to bias choice (see, e.g., Armel, Beaumel, and Rangel, 2008; Milosavljevic et al., 2012). The closest DDM models to ours are the so-called “race models” in which evidence in support of distinct alternatives is integrated in separate accumulators, with the choice driven by whichever accumulator reaches its stopping boundary first; see Bogacz et al. (2006) for a review. Relative to this literature, whose primary modeling goal is to fit choice data, we focus on foundations for the mechanism by which attention affects demand.

Optimal sequential learning about several items is related to the theory of multi-armed bandits (Gittins and Jones, 1974). We exploit this connection to show that optimal attention strategies satisfy IIA by using the Gittins index characterization.

2 Main result in a simple setting

In this section, we show in the simplest possible setting that temporarily manipulating attention toward an item increases the probability that the decision-maker (DM) chooses it.

The DM chooses one among two items $i \in \{1, 2\}$ of unknown values $v^i \in \{0, 1\}$ and an outside option with known value $z \in (0, 1)$. The two values v^i are independent ex ante, and each is equal to 1 with prior probability p_0^i . The DM learns sequentially about the value of each item, and can vary the focus of her learning as specified below. When she stops learning, she selects an item or the outside option based on whichever one has the highest posterior expected value.

Let p_t^i denote the DM’s belief about each item i at the beginning of the period t and write \mathbf{p}_t for the pair of beliefs (p_t^1, p_t^2) . At the beginning of each period $t = 0, 1, \dots$, the DM chooses an item ι_t on which to focus in period t . She receives a signal that is informative about the value of item ι_t and independent of the value of the other item. In this section, the signal takes on values 0 and 1, and for each v , the realized signal is equal to v with

⁶Ratcliff’s model is essentially equivalent to a sequential sampling model in the style of Wald (1945).




Figure 1: One possible strategy, together with three possible paths of beliefs. Each dot depicts a belief (p^1, p^2) that can be reached following some sequence of signals. The horizontal or vertical arrows at each point specify the direction of movement of the belief according to whether the strategy focuses on item 1 or 2, respectively. The bold dot identifies the prior belief. At each step of the learning process, the belief moves to one of the two adjacent points in the direction of the arrows. The dashed path of beliefs terminates with the choice of item 1, the solid path with the choice of item 2, and the dotted path with the choice of the outside option.

probability λ , where $\lambda > 1/2$; that is, $\Pr(x_t = 1 \mid v^{t_t} = 1) = \lambda = \Pr(x_t = 0 \mid v^{t_t} = 0)$. Upon observing a signal realization, the DM updates her belief according to Bayes' rule. In particular, the DM's belief about the item she is focusing on changes while her belief about the other item remains fixed. Thus, letting

$$p[+] = \frac{\lambda p}{\lambda p + (1 - \lambda)(1 - p)} \quad (1)$$

and $p[-] = \frac{(1 - \lambda)p}{(1 - \lambda)p + \lambda(1 - p)}$

for each $p \in (0, 1)$, we have that the posterior belief $p_{t+1}^{t_t}$ is either $p_t^{t_t}[+]$ or $p_t^{t_t}[-]$ according to whether x_t is 1 or 0, respectively, and $p_{t+1}^i = p_t^i$ for item $i \neq t_t$. Attention allocation is governed by a (pure) *attention strategy* $\alpha : [0, 1]^2 \rightarrow \{1, 2\}$ that specifies the item of focus $t_t = \alpha(\mathbf{p}_t)$ of a DM with beliefs \mathbf{p}_t .

The DM stops and makes a choice once she is sufficiently sure that either (i) one of

the items is of high value or (ii) both items are of low value. Accordingly, we introduce thresholds \underline{p} and \bar{p} such that $\underline{p} < z, p_0^i < \bar{p}$. We define stopping regions $F^i = \{\mathbf{p} : p^i \geq \bar{p}\}$ for $i = 1, 2$, $F^{oo} = \{\mathbf{p} : p^1, p^2 \leq \underline{p}\}$ and $F = F^1 \cup F^2 \cup F^{oo}$. Learning stops in the period $\tau = \min\{t : \mathbf{p}_t \in F\}$ with the DM choosing item i if $\mathbf{p}_\tau \in F^i$ and the outside option if $\mathbf{p}_\tau \in F^{oo}$. See Figure 1 for an illustration.

Let τ^i denote the period in which the DM chooses item i ; that is, $\tau^i = \tau$ if item i is chosen and $\tau^i = \infty$ otherwise. Note that τ and τ^i depend on the attention strategy; accordingly, we write $\tau(\alpha)$ and $\tau^i(\alpha)$ if the attention strategy α is not otherwise clear from the context.

The above rules specify, for any given pair of values $\mathbf{v} = (v^1, v^2)$, the joint stochastic process of beliefs and focus of attention $(\mathbf{p}_t, \iota_t)_t$, with the joint law denoted by $P_\alpha^\mathbf{v}$. For any strategy α and pair of values \mathbf{v} , we let the *interim demand* for item i ,

$$D^i(\mathbf{v}; \alpha) = P_\alpha^\mathbf{v}(\mathbf{p}_\tau \in F^i),$$

be the probability that the DM stops with the choice of i when the true values are \mathbf{v} . (Stopping in F^1 and F^2 are mutually exclusive.)

We are interested in how manipulation of the DM's attention strategy affects her choice.⁷ To this end, given a baseline strategy α , we introduce a *manipulated strategy* β constructed from α by making item 1 the item of focus in the initial period and then returning to α in all subsequent periods. That is, the item of focus $\beta(\mathbf{p}, t)$ in period t for beliefs \mathbf{p} is given by

$$\beta(\mathbf{p}, t) = \begin{cases} 1 & \text{if } t = 0, \\ \alpha(\mathbf{p}) & \text{if } t > 0. \end{cases}$$

We say that an attention strategy α is *non-wasteful* if $\alpha(\mathbf{p}) \neq i$ for any \mathbf{p} such that $p^i \leq \underline{p}$. Non-wasteful strategies do not focus on an item that the DM deems to have low value. Recall that a random variable τ (weakly) first-order stochastically dominates another random variable τ' if $\Pr(\tau \leq t) \leq \Pr(\tau' \leq t)$ for every t .

Proposition 1. *Suppose that the baseline attention strategy α is non-wasteful. For all pairs of values $\mathbf{v} \in \{0, 1\}^2$, manipulating attention toward item 1 in the first period*

⁷We do not explicitly model the mechanism of manipulation. Depending on the context, manipulation could result from changes in visual salience, in relative inspection costs, or in the position of the items on a list (as in online search results).

1. (weakly) increases the demand for item 1 and decreases the demand for item 2; that is,

$$D^1(\mathbf{v}; \beta) \geq D^1(\mathbf{v}; \alpha)$$

and $D^2(\mathbf{v}; \beta) \leq D^2(\mathbf{v}; \alpha),$

and

2. accelerates the choice of item 1 and decelerates the choice of item 2; that is, $\tau^1(\alpha)$ first-order stochastically dominates $\tau^1(\beta)$ and $\tau^2(\beta)$ first-order stochastically dominates $\tau^2(\alpha)$.

When an item has low value, the DM's belief about it tends to drift downward whenever she focuses on it. Yet, perhaps surprisingly, the proposition indicates that manipulating attention toward an item boosts its demand even in this case.

The first-order stochastic dominance results in the second part of the proposition combine two conceptually distinct effects. For example, for item 1, the probability of its being chosen under α is at most as high as under β , and hence the probability that $\tau^1(\beta) = \infty$ is at most as large as the corresponding probability for $\tau^1(\alpha)$. In addition, conditional on strategy β leading to the choice of item 1, there is a sense (outlined below) in which item 1 is chosen no later under strategy α .

Proposition 1 is a special case of Theorem 1 in the next section, and thus we provide only an informal proof here. Imagine that there is a large (countably infinite) deck of cards for each item, with each card showing a signal realization of 0 or 1. In each period t , the attention strategy chooses a deck $\alpha(\mathbf{p}_t)$ from which to draw the next card, and then the DM updates the relevant belief based on the signal shown on that card. Now consider the effect of manipulation on choice for a given sequence of cards in each deck, where manipulation induces the first card to come from the deck for item 1. To avoid trivialities, focus on the case in which, absent manipulation, the DM first draws from deck 2. The following argument shows that the statements in the proposition hold for each sequence of cards in each deck. It follows that they also hold when averaging across sequences conditional on the items' values (or, for that matter, conditional on any other event).

The DM chooses item 1 if her belief p_t^1 reaches (at least) \bar{p} before p_t^2 does. Intuitively, inducing the DM to draw first from deck 1 should only cause p_t^1 to reach \bar{p} sooner. There is a complication, however, insofar as manipulation can cause the order of subsequent draws to change since the DM may reach pairs of beliefs that she would not have reached otherwise. The key observation is that, once we have fixed the sequence of the cards, we only need to keep track of how many cards the baseline and manipulated strategies have drawn from

each deck. At the end of the first period, compared to the baseline strategy, the manipulated strategy is further ahead with deck 1 in the sense that more cards have been drawn from deck 1. Correspondingly, the baseline strategy is further ahead with deck 2. In each subsequent period, either the manipulated strategy remains ahead with deck 1 (perhaps pulling even further ahead) and the baseline strategy remains ahead with deck 2, or the numbers of draws from both decks under the baseline strategy “meet” the numbers under the manipulated strategy. In the latter case, the beliefs under the two processes coincide after the period in which they meet (since beliefs are independent of the order in which signals are received). Therefore, the manipulation has no effect on choice if the two processes meet. Accordingly, consider the case in which the two processes do not meet before one of them stops.

Suppose that the baseline strategy leads to the choice of item 1; that is, the belief about item 1 reaches \bar{p} in some period τ before the belief about item 2 reaches \bar{p} . Since the manipulated strategy is further ahead with deck 1 and behind with deck 2 in each period $t < \tau$, it must be that under the manipulated strategy, the belief about item 1 reaches \bar{p} before the belief about item 2 does, and does so no later than period τ . Therefore, the manipulated strategy also leads to the choice of item 1, with this choice occurring no later than under the baseline strategy.

The argument for the statements about the demand and decision time for item 2 is symmetric. Suppose that the manipulated strategy leads to the choice of item 2; that is, under the manipulated strategy, the belief about item 2 reaches \bar{p} in some period $\hat{\tau}$ before the belief about item 1 does. Since the baseline strategy is further ahead with deck 2 and behind with deck 1 in each period $t < \hat{\tau}$, under the baseline strategy, the belief about item 2 reaches \bar{p} before the belief about item 1 does, and does so no later than period $\hat{\tau}$.

3 General result for binary values

We now extend the setting to allow for more than two items, general signal structures, stochastic attention strategies, and arbitrary length of the manipulation window. At the end of the learning process, the DM chooses one item from the set $\mathcal{I} = \{1, \dots, I\}$ or an outside option with known value $z \in (0, 1)$. Each item $i \in \mathcal{I}$ has an uncertain value $v^i \in \{0, 1\}$, and again, $\mathbf{v} = (v^1, \dots, v^I)$.

In each period $t = 0, 1, \dots$, the DM focuses on a single item. Her belief in period t that item i is of high value ($v^i = 1$) is denoted p_t^i , and we write $\mathbf{p}_t = (p_t^i)_i$ for the *belief vector* in period t .⁸ As usual, given any finite set X , we write $\Delta(X)$ for the set of probability distributions over X . A (stochastic) *attention strategy* is a function $\alpha: (\Delta(\{0, 1\}))^I \times \mathbb{N} \rightarrow$

⁸We identify the probability $p_t^i \in [0, 1]$ attached to $v^i = 1$ with the corresponding belief in $\Delta(\{0, 1\})$.

$\Delta(\mathcal{I})$ that specifies a probability distribution over items of focus as a function of the belief vector at the beginning of a period together with the current time. We can think of $\alpha(\mathbf{p}, t)$ as a vector $(\alpha^i(\mathbf{p}, t))_i$, where $\alpha^i(\mathbf{p}, t)$ is the probability with which the DM focuses on item i in period t . An attention strategy is *stationary* if it does not depend on time (i.e., if $\alpha(\mathbf{p}, t) = \alpha(\mathbf{p}, t')$ for all \mathbf{p} , t , and t'), in which case we simply write $\alpha(\mathbf{p})$ for $\alpha(\mathbf{p}, t)$. If, for every \mathbf{p} and t , an attention strategy $\alpha(\mathbf{p}, t)$ assigns probability one to a single item, then we say that it is a *pure strategy*, and abuse notation slightly by writing $\alpha(\mathbf{p}, t)$ for the item of focus.

The evolution of beliefs as a function of the attention strategy is as follows. The DM begins with a prior belief vector $\mathbf{p}_0 = (p_0^1, \dots, p_0^I)$. If the beliefs at the beginning of period t are \mathbf{p}_t and the DM focuses on item i in period t , then the belief about item i follows a stochastic transition $\phi^i : \Delta(\{0, 1\}) \times \{0, 1\} \rightarrow \Delta(\Delta(\{0, 1\}))$ that depends on p_t^i and on v^i , while the beliefs about all other items remain unchanged; that is, conditional on \mathbf{p}_t , \mathbf{p}_{t+1} is a random variable such that p_{t+1}^i is distributed according to $\phi^i(p_t^i, v^i)$ and $p_{t+1}^j = p_t^j$ for all $j \neq i$. By fixing the DM's beliefs about items she is not currently focusing on, we are implicitly assuming that she treats the items' values as independent. One special case of this setting is when the DM receives a (serially conditionally independent) signal about the item she focuses on and updates her belief according to Bayes' rule. However, we also allow for non-Bayesian processes; for example, the "belief" p_t^i could alternatively be a neural firing rate, or the DM could under- or over-react to new information relative to Bayes' law.

An attention strategy α naturally induces a process of items of focus $(\iota_t)_t$ and a process of belief vectors $(\mathbf{p}_t)_t$. The law of ι_t conditional on $(\mathbf{p}_0, \iota_0, \dots, \mathbf{p}_{t-1}, \iota_{t-1}, \mathbf{p}_t)$ is $\alpha(\mathbf{p}_t, t)$, and the belief vector \mathbf{p}_{t+1} is drawn as described above, conditional on $(\mathbf{p}_0, \iota_0, \dots, \mathbf{p}_{t-1}, \iota_{t-1}, \mathbf{p}_t, \iota_t)$. Given any vector of values \mathbf{v} , we let $P_\alpha^\mathbf{v}$ denote the joint law of the (Markov) process $(\mathbf{p}_t, \iota_t)_t$.

The time at which the DM stops learning is governed by thresholds \underline{p}^i and \bar{p}^i for each i satisfying $0 \leq \underline{p}^i < p_0^i, z < \bar{p}^i \leq 1$. If the DM's belief p_t^i satisfies $p_t^i \leq \underline{p}^i$ then we say that she is *sufficiently certain* that $v^i = 0$, and similarly, if $p_t^i \geq \bar{p}^i$ then she is sufficiently certain that $v^i = 1$. The DM learns until she is sufficiently certain of an optimal choice. We thus define the stopping region F according to $F = (\bigcup_i F^i) \cup F^{\infty}$, where $F^i = \{\mathbf{p} : p^i \geq \bar{p}^i\}$ and $F^{\infty} = \{\mathbf{p} : p^i \leq \underline{p}^i \text{ for all } i\}$. The DM makes her choice at the stopping time $\tau = \min\{t : \mathbf{p}_t \in F\}$: either $\mathbf{p}_\tau \in F^i$ for some i , in which case this i is chosen, or $\mathbf{p}_\tau \in F^{\infty}$, in which case the outside option is chosen (these cases are mutually exclusive). For any item i , let the *stopping time* τ^i for i be equal to τ if i is chosen and ∞ otherwise. (We allow for the possibility that learning does not stop, in which case $\tau = \tau^i = \infty$ for all i .)

An attention strategy α is *non-wasteful* if $\alpha^i(\mathbf{p}, t) = 0$ for all \mathbf{p} such that $p^i \leq \underline{p}^i$. A non-wasteful strategy never focuses on an item that the DM is sufficiently certain is of low

value.

A special case of particular interest is when the thresholds are given by $\underline{p}^i = 0$ and $\bar{p}^i = 1$ for all i . These thresholds ensure that the DM learns until she is certain that she can make an optimal choice (at which point she stops immediately). One type of learning that eventually leads to certainty is so-called ‘‘Poisson learning’’: the DM is a Bayesian who receives a signal about the item she focuses on that perfectly reveals its value with positive probability. (This probability may depend on the item and its value.) Under such a learning process, thresholds of 0 and 1 are optimal for a DM who incurs a time cost of learning, but (lexicographically) prioritizes the value of her choice over this cost. More generally, however, we do not endogenize the stopping regions as resulting from optimization in a costly learning process; doing so would make for a very challenging problem involving tradeoffs that are orthogonal to the effect we identify. The key difference between our setting and one with endogenous boundaries is that we require the thresholds \underline{p}^i and \bar{p}^i to be independent of the DM’s beliefs about other items $j \neq i$. When the thresholds are interior, we interpret them as capturing bounded rationality that is particularly natural when the cost of learning is low (in which case the additional gain from precisely tailoring the thresholds is small). We conjecture that, in this case, our results approximate those arising from optimal strategies in Bayesian models with a small cost of learning.

An attention strategy α satisfies *Independence of Irrelevant Alternative i* (IIAi) if, conditional on not focusing on item i , the probabilities of focusing on each item $j \neq i$ are independent of p^i . Formally, this is the case when, for every t and $\mathbf{p}, \mathbf{q} \in (\Delta(\{0, 1\}))^I$ such that $p^j = q^j$ for all $j \neq i$ and $\alpha^i(\mathbf{p}, t), \alpha^i(\mathbf{q}, t) \neq 1$, we have that for every $j \neq i$,

$$\frac{\alpha^j(\mathbf{p}, t)}{1 - \alpha^i(\mathbf{p}, t)} = \frac{\alpha^j(\mathbf{q}, t)}{1 - \alpha^i(\mathbf{q}, t)}.$$

We say that α satisfies *Independence of Irrelevant Alternatives* (IIA) if it satisfies IIAi for all items i . Note that IIA is automatically satisfied if there are only two items. One way in which a strategy could violate IIAi would be if it focuses on the least promising of the remaining items when p^i is the lowest belief (to check whether i is indeed the worst), and on the most promising of the remaining items when p^i is the highest belief (to check whether i is indeed the best). In Subsection 3.2, we identify conditions under which the optimal strategy satisfies IIA.

We define the *interim demand for item i* as

$$D^i(\mathbf{v}; \alpha) = P_\alpha^\mathbf{v}(\mathbf{p}_\tau \in F^i);$$

this is the probability, under strategy α , that the DM chooses item i when the vector of values is \mathbf{v} .

We compare the interim demand under a baseline attention strategy to that under a manipulated strategy capturing the effect of manipulating attention. Given a *baseline strategy* α , a *target item* i , and a *manipulation length* $m \geq 1$, the *manipulated attention strategy* (in favor of i) is

$$\beta^j[\alpha, i, m](\mathbf{p}, t) = \begin{cases} 1_{j=i} & \text{if } t \leq m - 1, \text{ and } p^i > \underline{p}^i, \\ \alpha(\mathbf{p}, t) & \text{otherwise.} \end{cases}$$

Thus, under the manipulated attention strategy, the DM focuses on item i in the first m periods unless she is sufficiently certain that i is of low value, and then follows her baseline strategy in every subsequent period.⁹

The following proposition states that manipulation in favor of an item both increases and accelerates demand for this item, and decreases and decelerates demand for every other item. The result holds regardless of the underlying values: even if the target item is worse than other items, drawing attention to it is never detrimental to the likelihood that it is chosen.

Theorem 1. *If an attention strategy α is stationary, satisfies IIA i , and is non-wasteful, then for every \mathbf{v} and manipulation length $m \geq 1$,*

$$D^i(\mathbf{v}; \beta[\alpha, i, m]) \geq D^i(\mathbf{v}; \alpha)$$

and $D^j(\mathbf{v}; \beta[\alpha, i, m]) \leq D^j(\mathbf{v}; \alpha)$ for every $j \neq i$.

Moreover, $\tau^i(\alpha)$ first-order stochastically dominates $\tau^i(\beta[\alpha, i, m])$, and $\tau^j(\beta[\alpha, i, m])$ first-order stochastically dominates $\tau^j(\alpha)$ for every $j \neq i$.

The result extends immediately to manipulation windows starting in any period $t \geq 1$ provided that the manipulation does not affect the attention strategy employed before period t . This assumption corresponds to the idea that manipulation is unanticipated. In many contexts, while attention-grabbing behavior might be expected, predicting the timing of the manipulation and its target item is likely to be more difficult, which may preclude anticipatory adjustments to the attention strategy.

⁹Since Theorem 1 holds for any manipulation length, m can be randomized without affecting the conclusion.

3.1 Proof of Theorem 1

The proof relies on a technique known as “coupling” (see, e.g., Lindvall, 1992): we fix the vector of values \mathbf{v} and construct a common probability space on which we can compare the process of beliefs and items of focus $(\mathbf{p}_t, \iota_t)_t$ under the baseline attention strategy α with the process $(\hat{\mathbf{p}}_t, \hat{\iota}_t)_t$ under the manipulated strategy $\beta = \beta[\alpha, i, m]$. We construct this space in such a way that the law of $(\mathbf{p}_t, \iota_t)_t$ is $P_\alpha^\mathbf{v}$, while the law of $(\hat{\mathbf{p}}_t, \hat{\iota}_t)_t$ is $P_\beta^\mathbf{v}$.

We present here a coupling construction that suffices to prove the proposition for pure attention strategies. In Appendix A.1, we extend the construction to stochastic attention strategies, in which case the coupling argument is significantly more complex.

The probability space consists of realizations of a *learning process* $\pi = (\pi^j)_{j=1}^I$. The process π is a family of independent learning processes $\pi^j = (\pi_\kappa^j)_{\kappa=0,1,\dots}$ for each item j , where π^j is a Markov process starting at p_0^j with transitions $\phi^j(\cdot, v^j)$. The κ -th term π_κ^j of the learning process for item j specifies the belief about item j after κ periods of focus on that item. A *learning draw* is a realization of the learning process π .

We now construct, for each pure strategy $\gamma \in \{\alpha, \beta\}$, a realization of the process $(\mathbf{p}_t(\gamma), \iota_t(\gamma))_t$ as a function of the learning draw. To do this, we recursively define $(\mathbf{p}_t(\gamma), \iota_t(\gamma))_t$ as follows. For $t \geq 0$, and given the process $(\iota_s(\gamma))_{s < t}$, for each item j , let $k(j, t; \gamma)$ denote the *cumulative focus* of strategy γ on item j before period t ; that is,

$$k(j, t; \gamma) = |\{s < t : \iota_s(\gamma) = j\}|.$$

Similarly, let $k(-i, t; \gamma) = t - k(i, t; \gamma)$ denote the number of periods of focus on items other than i . Set $p_t^j(\gamma) = \pi_{k(j,t;\gamma)}^j$ for every j ; i.e., set the belief about each item j after $k(j, t; \gamma)$ periods of focus on this item to be the $k(j, t; \gamma)$ -th value of the learning process π^j . Given γ , let the focus in period t be $\iota_t(\gamma) = \gamma(\mathbf{p}_t)$. By construction, the law of the process $(\mathbf{p}_t(\gamma), \iota_t(\gamma))_t$ is $P_\gamma^\mathbf{v}$, as needed.

For notational purposes, it is convenient to extend the process $(\mathbf{p}_t(\gamma), \iota_t(\gamma))_t$ beyond the stopping time τ . To this end, if $\mathbf{p}_t(\gamma) \in F$, then we set $\iota_t(\gamma) = \emptyset$ and $\mathbf{p}_{t+1}(\gamma) = \mathbf{p}_t(\gamma)$. From this point forward, we write (\mathbf{p}_t, ι_t) for $(\mathbf{p}_t(\alpha), \iota_t(\alpha))$, $(\hat{\mathbf{p}}_t, \hat{\iota}_t)$ for $(\mathbf{p}_t(\beta), \iota_t(\beta))$, and similarly for k , and τ . We fix a learning draw and compare the two corresponding realizations of the processes $(\mathbf{p}_t, \iota_t)_t$ and $(\hat{\mathbf{p}}_t, \hat{\iota}_t)_t$.

For any $n \geq 0$, let $\mathbf{t}(n)$ be the n -th period such that $\iota_t \neq i$; that is, $\mathbf{t}(n) = \min\{t : k(-i, t) \geq n\}$, where potentially $\mathbf{t}(n) = \infty$, and similarly, $\hat{\mathbf{t}}(n) = \min\{t : \hat{k}(-i, t) \geq n\}$.

Lemma 1 (Coupling Lemma). *The baseline and manipulated processes coincide when re-*

stricted to periods of focus on items other than the target i . That is, for every n ,

$$\left(\mathbf{P}_{t(0)}^{-i}, \iota_{t(0)}, \dots, \mathbf{P}_{t(n)}^{-i}, \iota_{t(n)}\right) = \left(\hat{\mathbf{P}}_{t(0)}^{-i}, \hat{\iota}_{t(0)}, \dots, \hat{\mathbf{P}}_{t(n)}^{-i}, \hat{\iota}_{t(n)}\right).$$

Proofs omitted in the main text can be found in the appendix.

For $t \geq m$, we say that the baseline and the manipulated processes *meet* (in period t) if the cumulative focus on the target item i up to time t is the same for both processes, i.e. if $k(i, t) = \hat{k}(i, t)$. Note that for the two processes to meet, the definition considers only the cumulative focus on the *target* item. In general, this does not imply that the cumulative focus on any other item coincides under the two processes at time t (see counterexamples 4.2 and 4.3). The Coupling Lemma shows that when stationarity and IIA i hold, meeting of the two processes implies that the cumulative focus is the same for *every* item. According to the next result, the processes then coincide in every subsequent period.

Lemma 2 (Meeting Lemma). *If the baseline and manipulated processes meet in period t , then $(\mathbf{p}_s, \iota_s) = (\hat{\mathbf{p}}_s, \hat{\iota}_s)$ for all $s \geq t$.*

The Meeting Lemma implies the next result.

Lemma 3 (Attention Lemma). *In every period, the cumulative focus on the target item i is at least as large under the manipulated process as under the baseline process, and the cumulative focus on any $j \neq i$ is at least as large under the baseline process as under the manipulated process. That is, for every $t \geq 1$ and $j \neq i$,*

$$\begin{aligned} \hat{k}(i, t) &\geq k(i, t) \\ \text{and } k(j, t) &\geq \hat{k}(j, t). \end{aligned}$$

For any attention strategy γ , let $\tau^{oo}(\gamma) = \tau(\gamma)$ if $\mathbf{p}_{\tau(\gamma)} \in F^{oo}$ (and thus the outside option is chosen) and $\tau^{oo}(\gamma) = \infty$ (indicating that the outside option is not chosen) otherwise.

Lemma 4 (Outside Option Lemma). *For any two non-wasteful attention strategies γ and γ' and any learning draw, $\tau^{oo}(\gamma) = \tau^{oo}(\gamma')$.*

The Outside Option Lemma implies in particular that the outside option is chosen under the baseline process if and only if it is chosen under the manipulated process. Thus attention manipulation shifts demand within items without affecting the total demand across all items.

Lemma 5 (Choice Lemma). *For any learning draw,*

1. *if the target item i is chosen under the baseline process, then i is also chosen under the manipulated process, and no later than under the baseline process; and*

2. if some item $j \neq i$ is chosen under the manipulated process, then j is also chosen under the baseline process, and no later than under the manipulated process.

Proof. Statement 1: Consider any learning draw such that the target item i is chosen in period τ under the baseline process. Then, by the Outside Option Lemma, the outside option is not chosen under the manipulated process since it is not chosen under the baseline process. Suppose for contradiction that $j \neq i$ is chosen at some time $\hat{\tau} \leq \tau$ under the manipulated process. Then $\hat{p}_{\hat{\tau}}^j \geq \bar{p}^j$. By the Attention Lemma, the cumulative focus on items $-i$ by period $\hat{\tau}$ under the baseline process is at least as large as that under the manipulated process; that is, $k(-i, \hat{\tau}) \geq \hat{k}(-i, \hat{\tau})$. By the Coupling Lemma, there exists a period $t \leq \hat{\tau}$ such that $\mathbf{p}_t^{-i} = \hat{\mathbf{p}}_{\hat{\tau}}^{-i}$, and hence the baseline process stops with the choice of j in period t , which establishes the contradiction since stopping in F^i and F^j are mutually exclusive for $j \neq i$. Therefore, it cannot be that, under the manipulated process, an item $j \neq i$ is chosen at a time $\hat{\tau} \leq \tau$. By the Attention Lemma, $\hat{k}(i, t) \geq k(i, t)$ for all t . Hence, there exists $\hat{\tau} \leq \tau$ such that $\hat{p}_{\hat{\tau}}^i = p_{\hat{\tau}}^i \geq \bar{p}^i$ (since the manipulated process does not stop with the choice of $j \neq i$ or the outside option before $\hat{\tau}$). Thus, the manipulated process stops at time $\hat{\tau} \leq \tau$ with the choice of i , as needed.

Statement 2: The proof of the second statement is symmetric to that of the first. Accordingly, consider any draw such that, under the manipulated process, an item $j \neq i$ is chosen in period $\hat{\tau}$. Then, by the Outside Option Lemma, the baseline process does not choose the outside option. Suppose for contradiction that, under the baseline process, the target item i is chosen in some period $\tau \leq \hat{\tau}$. Then $p_{\tau}^i \geq \bar{p}^i$. By the Attention Lemma, $\hat{k}(i, \tau) \geq k(i, \tau)$. Thus, there exists a period $t \leq \tau$ such that $\hat{p}_t^i = p_{\tau}^i \geq \bar{p}^i$ and hence under the manipulated process, i is chosen in period $t \leq \hat{\tau}$, which establishes the contradiction. Therefore, it cannot be that, under the baseline process, i is chosen at a time $\tau \leq \hat{\tau}$. By the Attention Lemma, $k(-i, t) \geq \hat{k}(-i, t)$ for all t . By the Coupling Lemma, the beliefs \mathbf{p}^{-i} and $\hat{\mathbf{p}}^{-i}$ coincide when restricted to the periods of focus on items $-i$. Hence, there exists $\tau \leq \hat{\tau}$ such that $\mathbf{p}_{\tau}^{-i} = \hat{\mathbf{p}}_{\hat{\tau}}^{-i}$, and the baseline process stops in period τ with the choice of item j , as needed. \square

Theorem 1 follows from the Choice Lemma by taking expectations across learning draws.

3.2 IIA and stationarity

We now provide an argument in support of the IIA and stationarity assumptions based on optimization of the attention strategy. We fix, for each item i , the belief-updating process ϕ^i and the stopping thresholds \underline{p}^i and \bar{p}^i , and let the DM control her attention strategy α . We assume in this section that beliefs follow a Markov process (unconditional on \mathbf{v}), as is the

case if they are obtained through Bayesian updating based on observed signals. Until she stops learning, the DM pays a uniformly bounded flow cost $c(p_t^{l_t}, l_t) \geq 0$ in each period t ; this flow cost may depend on the item l_t of current focus and on the belief $p_t^{l_t}$ in the current period. After she terminates with an item i or with the outside option, the DM receives a one-time payoff v^i or z , respectively.

Given the stopping thresholds, the DM chooses a non-wasteful strategy α to maximize the expected discounted value

$$V(\alpha) = \mathbb{E} \left[\delta^\tau v^{i^*} - \sum_{t=0}^{\tau} \delta^t c(p_t^{l_t}, l_t) \right], \quad (2)$$

where i^* is the chosen item or the outside option (i.e., the item for which $\mathbf{p}_\tau \in F^{i^*}$), $v^{i^*} \in \{0, z, 1\}$ is its value, and $\delta \in (0, 1)$ is a discount factor. The expectation is with respect to the ex ante law $P_\alpha^{\text{ea}} = \mathbb{E} P_\alpha^\mathbf{v} = \sum_{\mathbf{v}} \left(\prod_j p_0^j(v^j) \right) P_\alpha^\mathbf{v}$ governing the evolution of beliefs.

We rely here on the theory of multi-armed bandits to show that a Gittins index strategy is optimal: for each item i , there exists a Gittins index function $G^i(p_t^i)$ that depends only on the belief about item i , such that the optimal strategy consists in each period of focusing on an item with the highest Gittins index. When ties are broken with uniform randomization, such a strategy satisfies IIA and stationarity.

Proposition 2. *There exists a non-wasteful strategy that maximizes the objective (2) and satisfies IIA and stationarity.*

Note that, since this result holds for all fixed stopping thresholds, it also holds if the thresholds are chosen optimally (subject to the constraint that the thresholds for each item i are independent of the beliefs about items other than i).

The main challenge in proving this result is that the cost in (2) exhibits interdependence across items because whether the decision process stops with the DM choosing the outside option depends on the whole profile of beliefs. Since the theory of multi-armed bandits applies to problems with flow payoffs that are independent across objects, we need to construct an auxiliary multi-armed bandit problem with this property. The construction is based on Lemma 4 (the Outside Option Lemma), which states that, conditional on the outside option being chosen, the stopping time is independent of the attention strategy. This construction applies because the stopping thresholds are restricted to be independent of the DM's beliefs about the other items; if both the attention strategy and the stopping region are optimized as in Nikandrova and Pancs (2018) and Ke and Villas-Boas (2019), then there need not exist an optimal Gittins index strategy and IIA and stationarity are not guaranteed.

4 Examples and counterexamples

4.1 Example: manipulation of the fastest strategy

To illustrate the quantitative impact of manipulating attention, we return to the simple setting from Section 2 and examine the fastest attention strategy for the given stopping thresholds. This strategy takes a simple form that allows for analytical characterization of the impact of manipulation.

Recall from Section 2 that if the DM has beliefs \mathbf{p}_t and focuses on item $j \in \{1, 2\}$ in period t , she updates her belief about j to $p_t^j[-]$ with probability $(1 - \lambda)v^j + \lambda(1 - v^j)$ and to $p_t^j[+]$ with probability $\lambda v^j + (1 - \lambda)(1 - v^j)$, where $p[-]$ and $p[+]$ are specified in (1). For simplicity, assume that the stopping thresholds \bar{p} and \underline{p} are the same for the two items, and that each of these thresholds can be reached exactly through some sequence of signals.¹⁰ Thus, the belief process for each item takes realizations in the set $\{\underline{p}, \underline{p}[+], \dots, \bar{p}[-], \bar{p}\}$.

The strategy α^* depicted in Figure 2A focuses on whichever item the DM views as more promising. Accordingly, for each item j and $i \neq j$,

$$\alpha^{*j}(\mathbf{p}) = \begin{cases} 1 & \text{if } p^j > p^i, \\ 1/2 & \text{if } p^j = p^i, \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

The next result states that α^* is the fastest attention strategy in this environment. Hence α^* is optimal for a DM who, given the stopping region F , minimizes a monotone time cost.¹¹ To state this formally, define the stochastic stopping time $\tau^{\text{ea}}(\alpha)$ for strategy α to be the minimal time t at which $\mathbf{p}_t(\alpha) \in F$ under the ex ante law $P_\alpha^{\text{ea}} = \mathbb{E} P_\alpha^{\mathbf{v}}$.

Proposition 3. *For any strategy α , $\tau^{\text{ea}}(\alpha)$ weakly first-order stochastically dominates $\tau^{\text{ea}}(\alpha^*)$.*

As with our main results, the proof makes use of coupling, although the particular construction is distinct from our main one. A different coupling construction is necessary because the strategy α^* is not the fastest one in every learning draw: there exist draws in which focusing on the more promising item leads to a long sequence of contradictory signals. To prove Proposition 3, we construct a distinct probability space in which a learning process

¹⁰On its own, the latter assumption is without loss of generality since moving a threshold within an interval in which no belief is reached with positive probability has no effect. When combined with the commonality of the thresholds, this assumption places a restriction on the prior belief vector.

¹¹Given the structure of this example, minimizing time cost is equivalent to maximizing the value $V(\alpha)$ in (2). This equivalence follows from the commonality of the thresholds across items together with the Outside Option Lemma.




Figure 2: A: the fastest attention strategy α^* . B: the ex ante demand $D_{\text{ea}}^1(p^1, p^2)$ for item 1 as a function of p^1 . The belief p^2 is fixed at 0.7 (solid curve), or 0.3 (dashed curve), respectively. The stopping thresholds are $\underline{p} \approx .99$ and $\bar{p} \approx .01$. The depicted (single-period) manipulation effect e^1 is evaluated for $p^1 = 0.5$, and $p^2 = .7$. The curve for $p^2 = 0.3$ is included for ease of comparison with the corresponding curve in Figure 3 when there is no outside option.

generated by any strategy γ follows the ex ante law P_γ^{ea} , and α^* is always at least as fast as any other strategy.

How much does the manipulation of attention affect demand? To quantify the effect, define the *ex ante demand (for item i)* by

$$D_{\text{ea}}^i(\mathbf{p}_0) = P_{\alpha^*}^{\text{ea}}(\mathbf{p}_\tau \in F^i) = \mathbb{E}_{\mathbf{p}_0} D^i(\mathbf{v}; \alpha^*).$$

Thus $D_{\text{ea}}^i(\mathbf{p}_0)$ is the ex ante probability that a DM with prior belief \mathbf{p}_0 chooses item i under the fastest strategy α^* . We define the *manipulation effect* $e^i(\mathbf{p})$ to be the change in the ex ante demand for i resulting from a single-period manipulation of the fastest strategy α^* that targets the item i ; that is,

$$e^i(\mathbf{p}) = \mathbb{E} [D^i(\mathbf{v}; \beta[\alpha^*, 1, 1]) - D^i(\mathbf{v}, \alpha^*)] = \mathbb{E} D_{\text{ea}}^i(\tilde{p}^i, p^{-i}) - D_{\text{ea}}^i(p^i, p^{-i}), \quad (4)$$

where the first expectation is with respect to \mathbf{v} distributed according to the prior \mathbf{p} , \tilde{p}^i is

the belief resulting from a single-period update of p^i , and the second expectation is over the possible values $p^i[+]$ and $p^i[-]$ of \tilde{p}^i . The magnitude of the effect is therefore determined by the curvature of $D_{\text{ea}}^i(p^i, p^{-i})$ with respect to p^i around \mathbf{p} ; see Figure 2B. Note that, by Theorem 1, $e^i(\mathbf{p})$ is nonnegative for all \mathbf{p} .

Proposition 4. *The manipulation effect $e^i(\mathbf{p})$ is*

1. *positive whenever manipulation affects attention (when $p^i \leq p^{-i}$),*
2. *decreasing in p^{-i} on $\{p^i[+], \dots, \bar{p}\}$, and*
3. *increasing in p^i on $\{\underline{p}, p[+], \dots, p^{-i}[-]\}$.*
4. *nonvanishing as $\underline{p}^i \rightarrow 0$ and $\bar{p}^i \rightarrow 1$ (in the region where the effect is positive).*

This result implies that manipulation has the strongest impact when it targets an item with a small a priori disadvantage relative to the other item. The proof in the appendix provides an explicit expression for the ex ante demand and the manipulation effect.

While the analysis above applies for exogenously fixed stopping beliefs, similar results hold in two related settings in which the stopping region is chosen optimally. First, consider a variant of the model in which learning terminates at an exogenous stochastic deadline T . At the termination time, the DM makes an optimal choice given her posterior belief. The optimal strategy in this problem solves

$$\max_{\alpha} \mathbb{E} \max \{z, p_T^1(\alpha), p_T^2(\alpha)\}, \quad (5)$$

where the expectation is with respect to T and the beliefs $p_T^i(\alpha)$.

Proposition 5. *The attention strategy α^* defined in (3) solves Problem (5).*

Second, Ke and Villas-Boas (2019) examine a related continuous-time setting with exponential discounting and show that the analogue of the strategy α^* (together with a stopping region they characterize) is optimal. Despite the differences among these models, all three converge to the same demands in the limit as learning becomes arbitrarily precise, i.e., in the limit as $\underline{p} \rightarrow 0$ and $\bar{p} \rightarrow 1$ in our model, as the support of T tends to ∞ in the stochastic deadline model, and as the discount factor tends to 1 in the model of Ke and Villas-Boas. Therefore, all three exhibit exactly the same manipulation effect in their respective limits.

4.2 Counterexample: failure of IIA

The conclusions of Theorem 1 do not generally hold if the attention strategy does not satisfy IIA. To illustrate this, consider an example with three items and a Bayesian DM with prior belief $p_0^i = 1/2$ for each item $i \in \{1, 2, 3\}$. If the DM's belief about an item i is p_t^i in period t , focusing on i in t leads to belief $p_{t+1}^i = p_t^i[-]$ or $p_t^i[+]$ as in Section 2. The stopping boundaries are $\underline{p} = \frac{1}{2}[-][-]$ and $\bar{p} = \frac{1}{2}[+][+]$ for each item. Let α be a stationary pure attention strategy that satisfies

$$\alpha(\mathbf{p}) = \begin{cases} 2 & \text{if } \mathbf{p} = \mathbf{p}_0, \\ 1 & \text{if } p^2 \neq p_0^2 \text{ and } p^1 > \underline{p}, \\ 3 & \text{if } \mathbf{p} \neq \mathbf{p}_0, p^2 = p_0^2, \text{ and } p^3 > \underline{p}. \end{cases}$$

When following such a strategy, the DM first focuses on item 2, and then, from the second period onwards, focuses on item 1 until p^1 reaches \underline{p} or \bar{p} . This leads to item 1 being chosen with probability $1/2$.

Now consider the manipulated strategy $\beta = \beta[\alpha, 1, 1]$ obtained by directing attention to item 1 in the first period. From the second period onwards, the DM focuses on item 3 until p^3 reaches \underline{p} or \bar{p} , and thus β chooses item 1 only if p^3 reaches \underline{p} and p^1 reaches \bar{p} . Therefore, the manipulated DM chooses the target item 1 with probability at most $1/4$.

The strategy α violates the IIA i assumption for $i = 1$ since the allocation of attention between items 2 and 3 at $\mathbf{p}^{-1} = (p_0^2, p_0^3)$ depends on the belief about item 1. A failure of IIA can cause the Meeting Lemma to be violated: the baseline and manipulated processes may meet but not coincide thereafter. Recall that two processes meet at some time if, by that time, they focus on the target item for the same number of periods. In this example, the baseline and manipulated processes meet after the first two periods since both focus on the target item (item 1) once. Yet the baseline beliefs \mathbf{p}_2^{-1} and the manipulated beliefs $\hat{\mathbf{p}}_2^{-1}$ about the non-target items differ, and hence the continuations of the processes differ as well.

4.3 Counterexample: failure of stationarity

We now show that the conclusions of Theorem 1 can fail if the attention strategy satisfies IIA but is non-stationary. Suppose there are three items with prior beliefs $p_0^i = 1/2$ for each $i \in \{1, 2, 3\}$. The stopping thresholds \underline{p} and \bar{p} and the signal technology for items 1 and 2 are the same as in the previous counterexample. On the other hand, the value of item 3 is perfectly revealed as soon as the DM focuses on it for a single period. The stopping thresholds are \underline{p} and \bar{p} for each item.

Let α be a non-stationary pure attention strategy that satisfies the following for $p^1 > \underline{p}$:

$$\alpha(\mathbf{p}, t) = \begin{cases} 2 & \text{if } t = 0, \\ 3 & \text{if } t = 1, p^1 \neq p_0^1, \text{ and } p^3 > \underline{p}, \\ 1 & \text{otherwise.} \end{cases}$$

Under α , the DM first focuses on item 2, and then focuses on item 1 until p^1 reaches \underline{p} or \bar{p} . Hence she chooses item 1 with probability $1/2$. Under the manipulated strategy $\beta = \beta[\alpha, 1, 1]$, the DM first focuses on item 1, and then focuses on item 3—thereby learning its value—in the second period. Thus, under β , the DM chooses item 1 only if item 3 has value 0 *and* p^1 reaches \bar{p} , which occurs with probability $1/4$.

As in the previous counterexample, Theorem 1 does not apply because the Meeting Lemma fails. The baseline and manipulated processes meet after two periods since each focuses on item 1 for exactly one of those periods. However, the beliefs about items 2 and 3 differ between these two processes at $t = 2$, which causes the continuation of the processes to differ. When IIA*i* and stationarity are satisfied, the Meeting Lemma follows from the fact that, for each draw, the accumulated focus $k(i, t)$ on the target item i is a sufficient statistic for the distribution of beliefs \mathbf{p}_t about all items in period t .

5 No outside option: forced choice

While our main results focus on choice in the presence of a known outside option, many experiments employ a “forced choice” design in which there is no such option and subjects must eventually choose one of the items. In this section, we show that, in our framework, but absent an outside option, the effect of attention manipulation depends on the value of the target item. As learning becomes precise (in the sense that the stopping thresholds approach 0 and 1), manipulating attention toward a good item increases its demand, while manipulating attention toward a bad item has the opposite effect. This result parallels the experimental findings of Armel et al. (2008) and Milosavljevic et al. (2012), who manipulate focus in forced choice problems involving either desirable or aversive items.

In our model of forced choice, the DM selects exactly one item from the set $\{1, \dots, I\}$. Each item i has value $v^i \in \{0, 1\}$. The DM learns about these values using an attention strategy α , with beliefs transitioning according to a stochastic rule ϕ^i , just as in Section 3. Relative to our main model, the DM’s choice differs in the stopping rule: here, the DM stops and chooses an item if she is sufficiently certain either that its value is high, or that the values of all other items are low. More precisely, for each item i , there are thresholds

\underline{p}^i and \bar{p}^i . Letting $\mathbf{p} = (\underline{p}^1, \bar{p}^1, \underline{p}^2, \bar{p}^2, \dots, \underline{p}^I, \bar{p}^I)$, the DM stops learning and chooses item i whenever the belief \mathbf{p} lies in the set

$$F_{\text{no}}^i(\mathbf{p}) = \{\mathbf{p} : p^i \geq \bar{p}^i \text{ or } p^j \leq \underline{p}^j \text{ for all } j \neq i\}.$$

Let $F_{\text{no}}(\mathbf{p}) = \bigcup_{i=1}^I F_{\text{no}}^i(\mathbf{p})$, and define the stopping time $\tau(\alpha, \mathbf{p}) = \min\{t : \mathbf{p}_t(\alpha) \in F_{\text{no}}(\mathbf{p})\}$; to simplify notation, we write τ in place of $\tau(\alpha, \mathbf{p})$ when the arguments are clear from the context.

We identify the impact of attention manipulation conditional on the value of the target item. To this end, let $P_\alpha^{v^i} = \mathbb{E} P_\alpha^{(v^i, v^{-i})}$, where the expectation is over $v^{-i} = (v^j)_{j \neq i}$ according to the prior distribution \mathbf{p}^{-i} . Thus $P_\alpha^{v^i}$ describes the evolution of beliefs conditional on the value v^i of item i . The *demand for item i* , $\tilde{D}^i(v^i; \alpha, \mathbf{p}) = P_\alpha^{v^i}(\mathbf{p}_\tau \in F_{\text{no}}^i(\mathbf{p}))$, is the probability that i is chosen conditional on its value.

The next result states that, in the limit as learning becomes perfectly precise, manipulating attention increases demand for a high-value target item and decreases demand for a low-value target item. This result holds for the same class of attention strategies as in Theorem 1.¹²

Proposition 6. *Suppose the attention strategy α is stationary, satisfies IIAi, and is non-wasteful. Then, in the limit as all of the thresholds \underline{p}^j tend to 0 and \bar{p}^j tend to 1, manipulating attention toward item i weakly increases its demand when its value is high and weakly decreases its demand when its value is low. That is, for each vector \mathbf{p} of thresholds, letting $\alpha_{\mathbf{p}}$ be an arbitrary stationary, non-wasteful attention strategy that satisfies IIAi,*

$$\liminf_{\mathbf{p} \rightarrow (0,1,\dots,0,1)} \left(\tilde{D}^i(1; \beta[\alpha_{\mathbf{p}}, i, m], \mathbf{p}) - \tilde{D}^i(1; \alpha_{\mathbf{p}}, \mathbf{p}) \right) \geq 0, \quad (6)$$

$$\text{and} \quad \limsup_{\mathbf{p} \rightarrow (0,1,\dots,0,1)} \left(\tilde{D}^i(0; \beta[\alpha_{\mathbf{p}}, i, m], \mathbf{p}) - \tilde{D}^i(0; \alpha_{\mathbf{p}}, \mathbf{p}) \right) \leq 0. \quad (7)$$

The proof in the appendix uses the same coupling construction as in the proof of Theorem 1: the coupling is across learning processes π^j that focus exclusively on items $j = 1, \dots, I$, respectively. Whenever the learning process for the target item i reaches its lower threshold \underline{p}^i and the learning process for at least one other item $j \neq i$ reaches its upper threshold \bar{p}^j , no attention strategy leads to item i being chosen. Hence choice is manipulable only if (i) the learning process π^i for the target item reaches its upper threshold \bar{p}^i , or

¹²Proposition 6 does not generally hold outside of the limit. For a counterexample, consider a binary choice with an attention strategy that focuses on item 2 whenever it is not excluded by the non-wastefulness condition. Suppose $v^1 = 0$. Then manipulation targeting item 1 increases its demand if there is some chance that the belief p^1 reaches \bar{p}^1 during the manipulation window.

(ii) the learning processes for all items reach their lower thresholds. As the lower thresholds approach 0 and the upper thresholds approach 1, the probability of event (i) vanishes when $v^i = 0$ and, likewise, the probability of event (ii) vanishes when $v^i = 1$. Finally, the manipulation effect is nonnegative in all draws in (i), and nonpositive in all draws in (ii).

The presence of an outside option affects demand in two distinct ways: directly, by allowing the DM not to choose any item, and indirectly, by affecting the stopping region. It turns out that the direct channel is irrelevant for the manipulation effect; the difference between the two settings is driven by the indirect channel.

To disentangle the two effects, consider an alternative model in which there is no outside option, and the stopping region is that from our main setting. Thus learning stops with the choice of item i whenever $\mathbf{p}_\tau \in F^i = \{\mathbf{p} : p^i \geq \bar{p}^i\}$ for some i , and with an equal probability of choosing any item whenever $\mathbf{p}_\tau \in F^{oo} = \{\mathbf{p} : p^j \leq \underline{p}^j \text{ for all } j\}$. The demand for item i is therefore

$$\begin{aligned} D_{\text{alt}}^i(\mathbf{v}; \alpha) &= P_\alpha^\mathbf{v}(\mathbf{p}_\tau \in F^i) + \frac{1}{I} P_\alpha^\mathbf{v}(\mathbf{p}_\tau \in F^{oo}) \\ &= D^i(\mathbf{v}; \alpha) + \frac{1}{I} P_\alpha^\mathbf{v}(\mathbf{p}_\tau \in F^{oo}), \end{aligned} \tag{8}$$

where $D^i(\mathbf{v}; \alpha)$ is the interim demand for item i from our main setting with an outside option. By Lemma 4 (Outside Option Lemma), the second summand in (8) is independent of the attention strategy. It follows that

$$D_{\text{alt}}^i(\mathbf{v}; \beta[\alpha, i, m]) - D_{\text{alt}}^i(\mathbf{v}; \alpha) = D^i(\mathbf{v}; \beta[\alpha, i, m]) - D^i(\mathbf{v}; \alpha),$$

and thus, for each \mathbf{v} , the manipulation effect in this alternative model is the same as in our main model with an outside option; the presence of the outside option affects the impact of manipulation only through the stopping region.¹³

To quantify the effect of attention manipulation in the forced choice model, consider a binary forced choice problem with the signal technology from Section 2. We derive the fastest attention strategy in this problem, which turns out to have a simple structure that allows us to analytically quantify the manipulation effect and compare it with that of the corresponding problem with an outside option. Consider threshold beliefs \underline{p} and $\bar{p} = 1 - \underline{p}$ —the same for the two items—such that each can be reached exactly through some sequence of signals starting from the prior belief. Thus, the belief process for each item takes realizations in the set $\{\underline{p}, \underline{p}[+], \dots, \bar{p}[-], \bar{p}\}$, which is symmetric around $1/2$.

¹³In Section 6, we extend the model to allow for multiple values. There, we abstract from the presence of the outside option, and instead assume that the DM wishes to learn the value of the chosen item. The

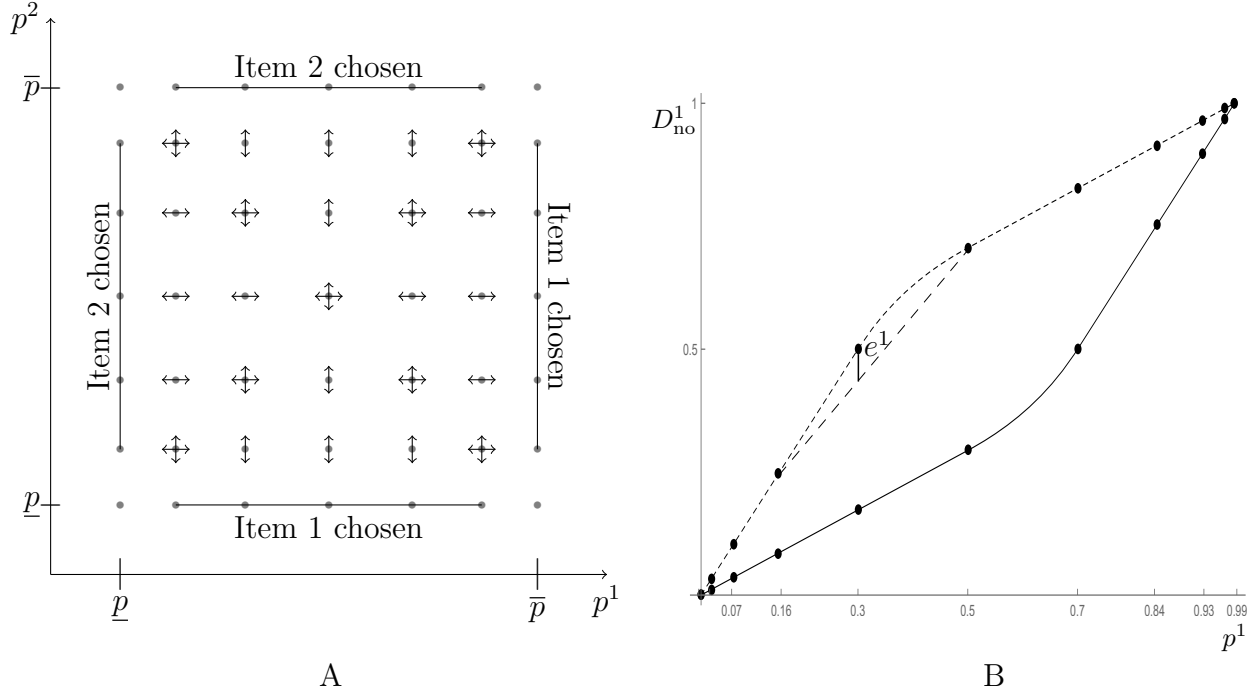


Figure 3: A: the fastest attention strategy α^{**} . B: demand $D_{\text{no}}^1(p^1, p^2)$ for item 1 as a function of p^1 . The belief p^2 is fixed at .7 (solid curve), or .3 (dashed curve), respectively. The stopping thresholds are $\underline{p} \approx .99$ and $\bar{p} \approx .01$. The depicted manipulation effect e^1 is evaluated for $p^1 = p^2 = .3$.

Let α^{**} be the strategy that focuses on the item about which the DM is more certain; that is, for $j \neq i$, let

$$\alpha^{**j}(\mathbf{p}) = \begin{cases} 1 & \text{if } |p^j - 1/2| > |p^i - 1/2|, \\ 1/2 & \text{if } |p^j - 1/2| = |p^i - 1/2|, \\ 0 & \text{otherwise.} \end{cases} \quad (9)$$

The strategy α^{**} is represented graphically in Figure 3A. Our next result states that α^{**} is the fastest attention strategy in this setting. Let $\tau_{\text{no}}^{\text{ea}}(\alpha)$ denote the time to decision under the ex ante law $P_{\alpha}^{\text{ea}} = \mathbb{E} P_{\alpha}^{\text{v}}$.

Proposition 7. *For any strategy α , $\tau_{\text{no}}^{\text{ea}}(\alpha)$ weakly first-order stochastically dominates $\tau_{\text{no}}^{\text{ea}}(\alpha^{**})$.*

To compare the manipulation effect for the fastest strategy in this setting to that in our main setting with an outside option, we compute the ex ante demand $D_{\text{no}}^1(\mathbf{p}) = P_{\alpha^{**}}^{\text{ea}}(\mathbf{p}_{\tau} \in F_{\text{no}}^1)$ in Appendix A.3 and plot it in Figure 3B as a function of p^1 . For $p^1 = p^2 = 0.3$, the figure demand defined in Section 6 is thus an extension of that in the alternative model described here.

shows that manipulation of attention toward item 1 decreases the demand for this item: at this prior, it is relatively likely that both π^1 and π^2 reach \underline{p} before \bar{p} . In any such draw, the DM eliminates the first item for which her belief reaches \underline{p} , and thus manipulating attention toward an item increases its chance of being eliminated. More generally, when $p^{-i} < 1/2$, the ex ante demand $D_{\text{no}}^i(\mathbf{p})$ is concave in p^i , and thus any manipulation with target item i decreases its demand. Similarly, when $p^{-i} > 1/2$, $D_{\text{no}}^i(\mathbf{p})$ is convex in p^i , and manipulation increases demand.

While manipulation can either increase or decrease demand for the target item when the set of items is fixed, it tends only to increase demand as the set of items grows large. Given any choice set I and prior \mathbf{p} , consider the enlarged choice problem obtained by an n -fold replication yielding choice set I^n and prior \mathbf{p}^n (i.e., the prior belief associated with any replication of item i is the same as for i itself). Since the negative manipulation effect arises only from draws in which the learning process for *every* item reaches its lower stopping threshold, as n grows large, this effect vanishes uniformly across all lengths of the manipulation window. In contrast, the positive manipulation effect arises from draws in which the learning processes for the target item and at least one other item reach the upper stopping threshold. Consequently, when n is sufficiently large, manipulation increases demand for the target item if the manipulation window is long enough.

6 Multiple values

We now extend the model from Section 3 to allow for the values of the items to lie in a finite set $V \subset \mathbb{R}$. Instead of comparing items to an outside option, we assume that the DM chooses an item i only if the belief p^i lies in a fixed set. Our main interpretation is that the DM wants to be informed about the chosen item, which may be useful for subsequent decisions in which the optimal action depends on the chosen item's value; accordingly, she stops learning only once she is sufficiently certain of this value. For example, one can think of the previous setup as involving a choice among items followed by a subsequent decision of whether to trade the chosen item for the outside option. Even in the absence of an outside option, an investor choosing among projects may not be content merely to learn that one project is likely to be better than the alternatives if information about the chosen project will help with other investment decisions.

As before, an attention strategy $\alpha(\mathbf{p}, t)$ specifies the probability distribution over the focus in period t at beliefs $\mathbf{p} \in (\Delta(V))^I$. For $p^j \in \Delta(V)$, $\phi^j(p^j, v^j) \in \Delta(\Delta(V))$ describes the distribution over beliefs about item j after one period of focus on this item starting from belief p^j when the true value is v^j . We extend the stopping rule from Section 3 as follows.

For each item j and value v , there is a nonempty *sufficient certainty region* $C^j(v) \subset \Delta(V)$. When $p^j \in C^j(v)$, we say that the DM is *sufficiently certain* that item j has value v . Let $C^j = \bigcup_{v \in V} C^j(v)$ denote the set of beliefs at which the DM is sufficiently certain of the value of item j . Similarly, for each j and v , there exists a nonempty *dominated region* $L^j(v) \subset \Delta(V)$. When $p^j \in L^j(v)$, we say that item j is *v-dominated*, with the interpretation that the DM is sufficiently certain that the value of item j is at most v . If for some v and i , $p^j \in L^j(v)$ and $p^i \in C^i(v)$, we say that j is *dominated*.

For any $\mathcal{I}^* \subseteq \mathcal{I}$ and any v , define the *stopping region* to be

$$F(\mathcal{I}^*, v) = \{\mathbf{p} : p^j \in C^j(v) \text{ for all } j \in \mathcal{I}^* \text{ and } p^j \in L^j(v) \text{ for all } j \in \mathcal{I} \setminus \mathcal{I}^*\}.$$

Thus $F(\mathcal{I}^*, v)$ consists of those beliefs at which the DM is sufficiently certain that items in \mathcal{I}^* have value v and all other items are v -dominated. Let $F^{\mathcal{I}^*} = \bigcup_v F(\mathcal{I}^*, v)$; whenever $\mathbf{p} \in F^{\mathcal{I}^*}$, we say that the DM *deems items in \mathcal{I}^* optimal*. Finally, let $F = \bigcup_{\mathcal{I}^*} F^{\mathcal{I}^*}$. The DM stops learning and makes a choice as soon as $\mathbf{p}_t \in F$; accordingly, $\tau = \min\{t : \mathbf{p}_t \in F\}$ is the stopping time for the learning process. (Again, we allow for the possibility that the process does not stop.) This stopping rule formalizes the assumption that the DM only stops when she is sufficiently certain of the value of the chosen item (and that it is optimal).

One example of particular interest is that of a Bayesian DM who receives signals that eventually, with some probability, perfectly reveal the value of the item she focuses on (as with Poisson learning). The DM may then always choose an optimal item, and stop learning as soon as she is certain of an optimal choice and of its value. In this case, the sets $C^i(v)$ are singletons consisting of the belief that attaches probability one to item i having value v , and the sets $L^i(v)$ consist of those beliefs that attach probability one to the event that $v^i \leq v$.

We impose two restrictions on the attention strategy α :

1. *non-wastefulness I*: if $\mathbf{p} \notin F$ and $p^j \in C^j$, then $\alpha^j(\mathbf{p}, t) = 0$; and
2. *non-wastefulness II*: if $\mathbf{p} \notin F$, and, for some v , $p^j \in L^j(v)$ and $p^{j'} \in C^{j'}(v)$ for some j' , then $\alpha^{j'}(\mathbf{p}, t) = 0$.

The first of these properties states that the DM does not focus on an item if she is sufficiently certain of its value. The second states that she does not focus on any item that is dominated. For example, in the case with perfectly revealing signals described above, these two properties are satisfied if she always focuses on an item for which the maximum of the support of the belief is the highest. One such strategy, which also satisfies stationarity and IIA, is when $\alpha(\mathbf{p})$ selects uniformly from $\arg \max_j \bar{v}(p^j)$, where $\bar{v}(p^j) = \max\{v \in V : p^j(v) > 0\}$.¹⁴

¹⁴In particular, non-wastefulness I holds because if the DM is certain of an item in $\arg \max_j \bar{v}(p^j)$ then

Unlike the setting with binary values, with multiple values, the DM may deem more than one item optimal at the stopping time. Her choice then depends on a tie-breaking rule. Accordingly, let $\sigma(j, \mathcal{I}^*)$ be the probability that the DM chooses item j if she deems the items in \mathcal{I}^* optimal. Assume that (i) $\sigma(j, \mathcal{I}^*) = 0$ if $j \notin \mathcal{I}^*$ (*optimality*), and (ii) $\sigma(j, \mathcal{I}^*) \geq \sigma(j, \mathcal{J}^*)$ if $\mathcal{I}^* \subseteq \mathcal{J}^*$ (*monotonicity*). For instance, these two properties hold if the DM selects uniformly from \mathcal{I}^* . The interim demand for item j under attention strategy α is

$$D^j(\mathbf{v}; \alpha) = \sum_{\mathcal{I}^*} P_{\alpha}^{\mathbf{v}}(\mathbf{p}_{\tau} \in F^{\mathcal{I}^*}) \sigma(j, \mathcal{I}^*).$$

As before, to compare the stopping times under the baseline and manipulated processes, let $\tau^j = \tau$ if item j is chosen, and $\tau^j = \infty$ otherwise. Given any α and \mathbf{v} , τ^j has distribution function

$$H^j(t; \mathbf{v}, \alpha) = \sum_{\tau' \leq t, \mathcal{I}^*} P_{\alpha}^{\mathbf{v}}(\mathbf{p}_{\tau'} \in F^{\mathcal{I}^*}) \sigma(j, \mathcal{I}^*).$$

Let $\beta[\alpha, i, m] = \beta$ denote the manipulated strategy constructed from the baseline strategy α by making the target item i the item of focus in periods $t = 0, \dots, m-1$ whenever the DM is not sufficiently certain about the value of i and i is not dominated; that is, $\beta^j(\mathbf{p}, t) = 1_{j=i}$ if $t \leq m-1$, $p^i \notin C^i$, and there is no item j and value v such that $p^j \in C^j(v)$ and $p^i \in L^i(v)$, and $\beta(\mathbf{p}, t) = \alpha(\mathbf{p}, t)$ otherwise. Note that β inherits non-wastefulness I and II from α .

The following result generalizes Theorem 1.

Theorem 2. *If an attention strategy α satisfies stationarity, IIAi, and non-wastefulness properties I and II, then for every \mathbf{v} and $m \geq 1$,*

$$D^i(\mathbf{v}; \beta[\alpha, i, m]) \geq D^i(\mathbf{v}; \alpha), \tag{10}$$

$$D^j(\mathbf{v}; \beta[\alpha, i, m]) \leq D^j(\mathbf{v}; \alpha) \text{ for all } j \neq i. \tag{11}$$

Moreover, $\tau^i(\alpha)$ first-order stochastically dominates $\tau^i(\beta[\alpha, i, m])$ and $\tau^j(\beta[\alpha, i, m])$ first-order stochastically dominates $\tau^j(\alpha)$ for every $j \neq i$.

7 Discussion

In light of our results, it is natural to ask when the effect of manipulation is largest. In general, the answer to this question depends on the attention strategy employed by the DM. We established in subsection 4.1 that, with two items and an outside option, if the DM uses the fastest strategy, manipulating attention has the greatest impact when the decision

the process stops.

makers belief about the target item is just below her belief about the other item, and when she is close to being ready to choose one of the items.

Our model offers a number of testable predictions about attention and choice. First, attention manipulation increases demand for the target item when an outside option is available, but can either increase or decrease demand depending on the value of the target item when there is no outside option. In binary choice problems, this prediction holds for every baseline attention strategy. Second, when an outside option is available, manipulation accelerates the choice of the target item (conditional on its being chosen). Third, in binary-choice problems with an outside option, if strategies minimize learning time, then focus on an item at any point in the process is positively associated with its choice relative to the other item. Each of these predictions is amenable to experimental tests using standard methods to identify focus and to manipulate attention either by boosting the salience of the target item or by introducing the target item for inspection earlier than the other items.

References

- Armel, K. C., A. Beaumel, and A. Rangel (2008). Biasing simple choices by manipulating relative visual attention. *Judgment and Decision Making* 3(5), 396–403.
- Austen-Smith, D. and C. Martinelli (2018). Optimal exploration. Working paper.
- Bagwell, K. (2007). The economic analysis of advertising. In M. Armstrong and R. Porter (Eds.), *Handbook of Industrial Organization*, Volume 3, pp. 1701–1844. Elsevier.
- Bogacz, R., E. Brown, J. Moehlis, P. Holmes, and J. D. Cohen (2006). The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review* 113(4), 700.
- Chandon, P., J. W. Hutchinson, E. T. Bradlow, and S. H. Young (2009). Does in-store marketing work? Effects of the number and position of shelf facings on brand attention and evaluation at the point of purchase. *Journal of Marketing* 73(6), 1–17.
- Che, Y.-K. and K. Mierendorff (2017). Optimal sequential decision with limited attention. Working paper.
- Dukas, R. (2002). Behavioural and ecological consequences of limited attention. *Philosophical Transactions of the Royal Society of London B: Biological Sciences* 357(1427), 1539–1547.
- Eliasz, K. and R. Spiegler (2011). On the strategic use of attention grabbers. *Theoretical Economics* 6(1), 127–155.

- Fogg-Meade, E. (1901). The place of advertising in modern business. *Journal of Political Economy* 9(2), 218–242.
- Gittins, J. and D. Jones (1974). A dynamic allocation index for the sequential design of experiments. In J. Gani (Ed.), *Progress in Statistics*, pp. 241–266. Amsterdam, NL: North-Holland.
- Ke, T. T., Z.-J. M. Shen, and J. M. Villas-Boas (2016). Search for information on multiple products. *Management Science* 62(12), 3576–3603.
- Ke, T. T. and J. M. Villas-Boas (2019). Optimal learning before choice. *Journal of Economic Theory* 180, 383–437.
- Krajbich, I., C. Armel, and A. Rangel (2010). Visual fixations and the computation and comparison of value in simple choice. *Nature Neuroscience* 13(10), 1292.
- Krajbich, I. and A. Rangel (2011). Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions. *Proceedings of the National Academy of Sciences* 108(33), 13852–13857.
- Lindvall, T. (1992). *Lectures on the Coupling Method*. Wiley.
- Lou, D. (2014). Attracting investor attention through advertising. *Review of Financial Studies* 27(6), 1797–1829.
- Mandelbaum, A., L. A. Shepp, and R. J. Vanderbei (1990). Optimal switching between a pair of Brownian motions. *Annals of Probability* 18(3), 1010–1033.
- Mayer, A. R., J. M. Dorflinger, S. M. Rao, and M. Seidenberg (2004). Neural networks underlying endogenous and exogenous visual-spatial orienting. *Neuroimage* 23(2), 534–541.
- Milosavljevic, M., V. Navalpakkam, C. Koch, and A. Rangel (2012). Relative visual saliency differences induce sizable bias in consumer choice. *Journal of Consumer Psychology* 22(1), 67–74.
- Morris, S. and P. Strack (2017). The Wald problem and the equivalence of sequential sampling and static information costs. Working paper.
- Nikandrova, A. and R. Páncs (2018). Dynamic project selection. *Theoretical Economics* 13(1), 115–143.

- Orquin, J. L. and S. M. Loose (2013). Attention and choice: A review on eye movements in decision making. *Acta Psychologica* 144(1), 190–206.
- Platt, M. L. and P. W. Glimcher (1999). Neural correlates of decision variables in parietal cortex. *Nature* 400(6741), 233.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review* 85(2), 59.
- Seasholes, M. S. and G. Wu (2007). Predictable behavior, profits, and attention. *Journal of Empirical Finance* 14(5), 590–610.
- Spelke, E., W. Hirst, and U. Neisser (1976). Skills of divided attention. *Cognition* 4(3), 215–230.
- Tsitsiklis, J. N. (1994). A short proof of the gittins index theorem. *Annals of Applied Probability* 4(1), 194–199.
- Wald, A. (1945). Sequential tests of statistical hypotheses. *Annals of Mathematical Statistics* 16(2), 117–186.
- Weber, R. (1992). On the Gittins index for multiarmed bandits. *Annals of Applied Probability* 2(4), 1024–1033.
- Yorzinski, J. L., G. L. Patricelli, J. S. Babcock, J. M. Pearson, and M. L. Platt (2013). Through their eyes: Selective attention in peahens during courtship. *Journal of Experimental Biology* 216(16), 3035–3046.

A Proofs

A.1 Proofs for Section 3

To accommodate stochastic attention strategies, we introduce the *attention process* describing, for each time t , (i) for each belief vector $\mathbf{p} = (p^i, \mathbf{p}^{-i})$, whether the DM focuses on the target item i , and (ii) for each \mathbf{p}^{-i} , which item she focuses on when she does not focus on i . The attention process is given by a family of random variables $(a_{\mathbf{p},t})_{\mathbf{p},t}, (b_{\mathbf{p}^{-i},t})_{\mathbf{p}^{-i},t}$ where all draws are independent and independent of the learning process. The random variable $a_{\mathbf{p},t}$ takes values in $\{0, 1\}$; the probability that it takes the value 1 is $\alpha^i(\mathbf{p})$. The random variable $b_{\mathbf{p}^{-i},t}$ takes values in $\mathcal{I} \setminus \{i\}$. For a fixed value of \mathbf{p}^{-i} , if there exists p^i such that $\alpha^i(p^i, \mathbf{p}^{-i}) \neq 1$, then the probability that $b_{\mathbf{p}^{-i},t}$ takes the value $j \neq i$ is $\alpha^j(p^i, \mathbf{p}^{-i}) / (1 - \alpha^i(p^i, \mathbf{p}^{-i}))$, where

we note that the particular value of p^i in the formula is irrelevant since α satisfies IIAi.¹⁵ A realization of the attention process is called an *attention draw*. We refer to the pair of the learning and attention draws simply as a *draw*.

We now recursively construct, for stochastic attention strategies $\gamma \in \{\alpha, \beta\}$, the processes $(\mathbf{p}_t(\gamma), \iota_t(\gamma))_t$ as functions of the learning and attention processes. As in the construction for the pure strategies in subsection 3.1, let the belief $p_t^j(\gamma) = \pi_{k(j,t;\gamma)}^j$ where $k(j,t;\gamma)$ is the cumulative focus on item j in periods $0, \dots, t-1$. We proceed to construct the focus $\iota_t(\gamma)$. Suppose that $\mathbf{p}_t \notin F$. Given a sequence of beliefs and focus items $(\mathbf{p}_s(\gamma), \iota_s(\gamma))_{s < t}$ and any vector \mathbf{p}^{-i} of beliefs about $j \neq i$, we let $\mu(\mathbf{p}^{-i}, t; \gamma)$ be the total number of periods starting with belief \mathbf{p}^{-i} about items $j \neq i$ in which the DM has not focused on the target item i under the strategy γ before time t ; that is,

$$\mu(\mathbf{p}^{-i}, t; \gamma) = |\{s < t : \mathbf{p}_s^{-i}(\gamma) = \mathbf{p}^{-i}, \iota_s(\gamma) \neq i\}|.$$

Let $\iota_t(\gamma) = i$ if $a_{\mathbf{p}_t(\gamma), t} = 1$ and $\iota_t(\gamma) = b_{\mathbf{p}_t^{-i}(\gamma), \mu(\mathbf{p}^{-i}(\gamma), t; \gamma)}$ otherwise. Note a subtlety in the construction: the element $a_{\mathbf{p}, t}$ of the attention draw deciding whether the focus at \mathbf{p} is on the target item i depends on t , while the element $b_{\mathbf{p}^{-i}, \mu(\mathbf{p}^{-i}, t; \gamma)}$ of the draw deciding the item of focus conditional on its *not* being i depends on $\mu(\mathbf{p}^{-i}, t; \gamma)$. This asymmetry in the construction is exploited in the proof of Lemma 2.

By construction, conditional on $(\mathbf{p}_0(\gamma), \iota_0(\gamma), \dots, \mathbf{p}_{t-1}(\gamma), \iota_{t-1}(\gamma), \mathbf{p}_t(\gamma))$, $\iota_t(\gamma)$ is distributed according to $\alpha(\mathbf{p}_t(\gamma))$. Conditional on $(\mathbf{p}_0(\gamma), \iota_0(\gamma), \dots, \mathbf{p}_t(\gamma), \iota_t(\gamma))$, $p_{t+1}^j(\gamma) = p_t^j(\gamma)$ for all $j \neq \iota_t(\gamma)$, and $p_{t+1}^{\iota_t(\gamma)}(\gamma)$ is distributed according to the transition probability $\phi^{\iota_t(\gamma)}(p_t^{\iota_t(\gamma)}(\gamma), v^{\iota_t(\gamma)})$. Therefore, the law of the process $(\mathbf{p}_t(\gamma), \iota_t(\gamma))_t$ is P_γ^v , as needed.

Proof of Lemma 1 (Coupling Lemma). We prove the result by induction. The property holds for $n = 0$ since $\mathbf{p}_{\mathfrak{t}(0)}^{-i} = \hat{\mathbf{p}}_{\mathfrak{t}(0)}^{-i} = \mathbf{p}_0^{-i}$. Assume it holds for n . For item $j = \iota_{\mathfrak{t}(n)} = \hat{\iota}_{\mathfrak{t}(n)}$, we have $p_{\mathfrak{t}(n+1)}^j = \pi_{k(j, \mathfrak{t}(n)+1)}^j = \pi_{\hat{k}(j, \hat{\mathfrak{t}}(n)+1)}^j = \hat{p}_{\mathfrak{t}(n+1)}^j$. For items $j \neq i, \iota_{\mathfrak{t}(n)}$ that do not receive attention in periods $\mathfrak{t}(n)$ and $\hat{\mathfrak{t}}(n)$, respectively, we have $p_{\mathfrak{t}(n+1)}^j = p_{\mathfrak{t}(n)}^j = \hat{p}_{\mathfrak{t}(n)}^j = \hat{p}_{\hat{\mathfrak{t}}(n+1)}^j$. We thus have $\mathbf{p}_{\mathfrak{t}(n+1)}^{-i} = \hat{\mathbf{p}}_{\hat{\mathfrak{t}}(n+1)}^{-i}$.

It remains to show that the items of attention in period $\mathfrak{t}(n+1)$ of the baseline process and in period $\hat{\mathfrak{t}}(n+1)$ of the manipulated process coincide. By the definitions of the attention processes, we have

$$\iota_{\mathfrak{t}(n+1)} = b_{\mathbf{p}_{\mathfrak{t}(n+1)}^{-i}, \mu(\mathbf{p}_{\mathfrak{t}(n+1)}^{-i}, \mathfrak{t}(n+1))} = b_{\hat{\mathbf{p}}_{\hat{\mathfrak{t}}(n+1)}^{-i}, \mu(\hat{\mathbf{p}}_{\hat{\mathfrak{t}}(n+1)}^{-i}, \mathfrak{t}(n+1))} = b_{\hat{\mathbf{p}}_{\hat{\mathfrak{t}}(n+1)}^{-i}, \hat{\mu}(\hat{\mathbf{p}}_{\hat{\mathfrak{t}}(n+1)}^{-i}, \hat{\mathfrak{t}}(n+1))} = \hat{\iota}_{\hat{\mathfrak{t}}(n+1)}$$

since, by the induction hypothesis, $\mu(\hat{\mathbf{p}}_{\hat{\mathfrak{t}}(n+1)}^{-i}, \mathfrak{t}(n+1)) = \hat{\mu}(\hat{\mathbf{p}}_{\hat{\mathfrak{t}}(n+1)}^{-i}, \hat{\mathfrak{t}}(n+1))$. \square

¹⁵The specification of $b_{\mathbf{p}^{-i}, t}$ when no such p^i exists is immaterial for our purposes.

The next proof exploits the subtlety in the coupling construction mentioned above. The subtlety ensures that if the processes meet in period t , they coincide thereafter. The event that they meet in period t implies that the two processes have visited each vector \mathbf{p}^{-i} in the same number of periods, but it does not ensure that the number of periods they visited each vector \mathbf{p} is the same. To this end, the draw a_t in the coupling construction does not depend on the number of periods that \mathbf{p} was visited (as it could be different for the two processes), but rather on t .

Proof of Lemma 2 (Meeting Lemma). Suppose the processes meet in period t , and thus $k(-i, t) = \hat{k}(-i, t)$. Hence, by the Coupling Lemma,

$$\left(\mathbf{P}_{\iota(0)}^{-i}, \iota(0) \cdots, \mathbf{P}_{\iota(k(-i,t))}^{-i}, \iota(k(-i,t)) \right) = \left(\hat{\mathbf{P}}_{\hat{\iota}(0)}^{-i}, \hat{\iota}(0), \cdots, \hat{\mathbf{P}}_{\hat{\iota}(\hat{k}(-i,t))}^{-i}, \hat{\iota}(\hat{k}(-i,t)) \right).$$

Therefore, for every $j \in \mathcal{I}$, $k(j, t) = \hat{k}(j, t)$, so that the cumulative focus on each item before time t is the same in the two processes, which implies that $\mathbf{p}_t = \hat{\mathbf{p}}_t$. The baseline process focuses on the target item i at time t if and only if the manipulated process focuses on i at t , since both processes focus on i at t if $a_{\mathbf{p}_t, t} = a_{\hat{\mathbf{p}}_t, t} = 1$. The baseline process focuses on item $j \neq i$ at time t if and only if the manipulated process focuses on j at t , since the baseline process focuses on j at t if $b_{\mathbf{p}_t^{-i}, \mu(\mathbf{p}_t^{-i}, t)} = j$, the manipulated process focuses on j in t if $b_{\hat{\mathbf{p}}_t^{-i}, \hat{\mu}(\hat{\mathbf{p}}_t^{-i}, t)} = j$, and $\mu(\mathbf{p}_t^{-i}, t) = \hat{\mu}(\hat{\mathbf{p}}_t^{-i}, t)$ by the Coupling Lemma. Thus $\iota_t = \hat{\iota}_t$.

Note that $k(i, t+1) = \hat{k}(i, t+1)$. Therefore, if the two processes meet in period $t \geq m$, they also meet in period $t+1$, and hence in every period $s \geq t$. Thus, $(\mathbf{p}_t, \iota_s) = (\hat{\mathbf{p}}_t, \hat{\iota}_s)$ for all $s \geq t$. \square

Proof of Lemma 3 (Attention Lemma). The statement obviously holds for every $t \leq m$. Since $\hat{k}(i, m) = m \geq k(i, m)$, and $\hat{k}(i, t) = k(i, t)$ implies $\hat{k}(i, s) = k(i, s)$ for every $s > t$, we have $\hat{k}(i, t) \geq k(i, t)$ and $\hat{k}(-i, t) \leq k(-i, t)$ for every $t > m$. The Coupling Lemma implies that items $j \neq i$ are explored in the same order under both processes, which in turn implies that, for every t and every $j \neq i$, $\hat{k}(j, t) \leq k(j, t)$, as needed. \square

Proof of Lemma 4 (Outside Option Lemma). Fix a learning draw such that $\tau^{oo}(\alpha)$ is finite; that is, strategy α leads to the outside option being chosen. For each item i , let $\kappa^i = \min\{\kappa : \pi_\kappa^i \leq \underline{p}^i\}$ be the number of steps needed for the learning process π^i to reach the threshold \underline{p}^i . Since α stops with the choice of the outside option once all beliefs reach their respective lower thresholds, $\tau^{oo}(\alpha) = \sum_i \kappa^i$, which is independent of α . \square

Proof of Proposition 2. We construct a bandit problem that is a special case of the model studied in Tsitsiklis (1994). In the model of Tsitsiklis (1994) (described in the language of

our model), when the DM selects an item $i = \iota_t$ with belief p_t^i in period t , she receives a stochastic flow reward $R(p_t^i, i)$ and the selected item remains active for $T^i(p_t^i)$ periods. In period $t+T^i(p_t^i)$, the belief about item i transitions to a new belief $p_{t+T^i(p_t^i)}^i = \psi^i(p_t^i)$ according to some stochastic function ψ^i , while the beliefs for all other items remain unchanged (that is, $p_{t+T^i(p_t^i)}^j = p_t^j$ for all $j \neq i$). Finally, the DM selects a new item $\iota_{t+T^i(p_t^i)}$ in period $t+T^i(p_t^i)$. Let t_k denote the random sequence of periods in which the DM selects items: thus, $t_0 = 0$ and $t_{k+1} = t_k + T(p_{t_k}^{\iota_{t_k}})$ for each $k \geq 0$. The DM maximizes

$$\Gamma(\alpha) = \mathbb{E} \sum_{k=0}^{\infty} \delta^{t_k} R(p_{t_k}^{\iota_{t_k}}, \iota_{t_k}). \quad (12)$$

Let $\tilde{\phi}^i(p_{t+1}^i | p_t^i)$ denote the ex ante Markov transition probabilities of the DM's belief about item i when she focuses on i in period t , and define the stochastic transition function ψ by $\Pr(\psi(p_t^i) = p) = \tilde{\phi}^i(p | p_t^i)$ for each p . Fix the reward function to be

$$R(p, i) = \begin{cases} -c(p, i) & \text{if } \underline{p}^i < \psi^i(p) < \bar{p}^i, \\ -c(p, i) + \psi^i(p) & \text{if } \psi^i(p) \geq \bar{p}^i, \\ -\bar{c} & \text{otherwise,} \end{cases}$$

where \bar{c} is a uniform upper bound on $c(p_t^{\iota_t}, \iota_t)$. If item i is selected in period t_k , we set the activity length $T^i(p_{t_k}^i)$ to be 1 if $p_{t_{k+1}}^i < \bar{p}^i$ and ∞ otherwise.¹⁶

This problem differs from the original problem in that the flow cost from focusing on item i in period t depends only on p_t^i and is independent of p_t^j for $j \neq i$. The flow payoffs in the original problem are interdependent in that whether the DM stops upon reaching $p_t^i \leq \underline{p}$ depends on whether $p_t^j \leq \underline{p}$ for all other items $j \neq i$. This difference allows us to apply Tsitsiklis's result from the theory of multi-armed bandits (see also Weber (1992)): an optimal solution to the bandit problem (12) is a Gittins index strategy where the index for an item depends only on the current belief. In particular, when ties for the highest Gittins index are broken by uniform randomization, this describes an optimal strategy that satisfies IIA and stationarity.

It remains to connect the bandit problem (12) to the original problem (2). Note that any solution to the bandit problem (i) chooses an item i forever once $p_t^i \geq \bar{p}^i$, and (ii) always chooses an item i such that $p_t^i > \underline{p}$ if such an item exists.

Note that $\Gamma(\alpha) = V(\alpha) - \bar{c} \mathbb{E} \sum_{t=1}^{\infty} 1_{t \geq \tau^{\circ\circ}(\alpha)}$, where the expectation is with respect to $\tau^{\circ\circ}(\alpha)$. By the Outside Option Lemma, $\tau^{\circ\circ}(\alpha)$ is identical for all α satisfying (i) and (ii).

¹⁶This correlation among the active-period length, reward, and belief transition complies with Tsitsiklis's setting, as he allows for arbitrary correlations among these random variables.

Therefore, there exists a constant K_δ such that

$$\Gamma(\alpha) = C(\alpha) + K_\delta. \quad (13)$$

Thus, a non-wasteful strategy α solves the auxiliary bandit problem if and only if it solves the original problem, as needed. \square

A.2 Proofs for Section 4

Proof of Proposition 3. This proof makes use of coupling. We again construct a common probability space on which we can compare the processes of beliefs $\mathbf{p}_t(\alpha)$ and $\mathbf{p}_t(\alpha^*)$, where, by construction, the beliefs under strategy α follow the law P_α^{ea} while those under α^* follow $P_{\alpha^*}^{\text{ea}}$. However, the particular construction differs from that in the proof of Theorem 1.

Let $\Pi = \{\underline{p}[+], \underline{p}[+][+], \dots, \bar{p}[-][-], \bar{p}[-]\}$, which is the set of transient beliefs that the DM may attain for either item. For each $\pi \in \Pi$ and $\kappa = 0, 1, \dots$, let $\ell(\pi, \kappa)$ be an i.i.d. random variable that attains values $\pi[+]$ and $\pi[-]$ with probabilities $\frac{\pi - \pi[-]}{\pi[+] - \pi[-]}$ and $\frac{\pi[+] - \pi}{\pi[+] - \pi[-]}$, respectively. An *updating draw* is a collection $(\ell(\pi, \kappa))_{\pi, \kappa}$ of realizations of $\ell(\pi, \kappa)$, one for each pair $(\pi, \kappa) \in \Pi \times \mathbb{N}$.

We interpret $\ell(\pi, \kappa)$ as the updated belief of a DM who learns for one period about an item i starting at the belief $p^i = \pi$ where κ is a counter indicating the total number of times the DM has focused on an item with associated belief π . We now construct, for each fixed updating draw and any attention strategy γ , the process of beliefs $\mathbf{p}_t(\gamma)$. In this construction, we use an auxiliary counter $\mathbf{k}_t(\gamma)$ that takes values in $\mathbb{N}^{|\Pi|}$. Define the joint process of $\mathbf{p}_t(\gamma)$ and $\mathbf{k}_t(\gamma)$ as follows. Let $k_0^\pi(\gamma) = 0$ for all π and $\mathbf{p}_0(\gamma) = \mathbf{p}_0$. In each period t , the focus of attention ι_t in t is chosen according to the attention strategy $\gamma(\mathbf{p}_t, t)$. Recursively define

$$k_{t+1}^\pi(\gamma) = \begin{cases} k_t^\pi(\gamma) + 1 & \text{if } \pi = p_{t+1}^{\iota_t}, \\ k_t^\pi(\gamma) & \text{otherwise} \end{cases}$$

and

$$p_{t+1}^j(\gamma) = \begin{cases} \ell\left(\pi, k_t^{p_t^{\iota_t}}(\gamma)\right) & \text{if } j = \iota_t, \\ p_t^j(\gamma) & \text{otherwise.} \end{cases}$$

By construction, for each strategy γ , the beliefs $\mathbf{p}_t(\gamma)$ follow the law P_γ^{ea} .

For each transient belief $\pi \in \Pi$, we introduce the belief process $a_t(\pi)$ that would result from learning about a single item starting from the prior belief π , making use of a counter $h_t^{\tilde{\pi}}$. Formally, let $a_0(\pi) = \pi$, $h_0^{\tilde{\pi}} = 0$ for all $\tilde{\pi} \in \Pi$, and recursively define $a_{t+1}(\pi) =$

$\ell(a_t(\pi), h_t^{a_t(\pi)})$, and, for each $\tilde{\pi} \in \Pi$,

$$h_{t+1}^{\tilde{\pi}} = \begin{cases} h_t^{\tilde{\pi}} + 1 & \text{if } \tilde{\pi} = a_t(\pi), \\ h_t^{\tilde{\pi}} & \text{otherwise.} \end{cases}$$

For each strategy γ , let $M_t(\gamma) = \max_{i=1,2} p_t^i(\gamma)$, $\overline{M}_t(\gamma) = \max_{s=0,\dots,t} M_s(\gamma)$, and $\underline{M}_t(\gamma) = \min_{s=0,\dots,t} M_s(\gamma)$. The strategy γ stops by period t , i.e., $\tau^{\text{ea}}(\gamma) \leq t$, if

$$\underline{p} = \underline{M}_t(\gamma) \quad \text{or} \quad \overline{M}_t(\gamma) = \overline{p}.$$

We will prove that, for every prior, in each updating draw, if a strategy α stops by t , then the strategy α^* also stops by t . We proceed by induction on t . To see that the statement holds for $t = 1$, note that if $p_0^1 \neq p_0^2$, then $M_0(\alpha) = M_0(\alpha^*) = \max\{p_0^1, p_0^2\}$ and $M_1(\alpha^*) = a_1(\max\{p_0^1, p_0^2\})$ while $M_1(\alpha)$ equals $a_1(\max\{p_0^1, p_0^2\})$ or $\max\{p_0^1, p_0^2\}$. (The latter case arises when α focuses on the item with the lower belief in period 0.) Thus

$$\underline{M}_1(\alpha^*) \leq \underline{M}_1(\alpha) \leq \overline{M}_1(\alpha) \leq \overline{M}_1(\alpha^*),$$

as needed.

Suppose the statement holds for $t - 1$. If a strategy α stops by period t then, since the induction hypothesis applies regardless of the prior, the strategy $\tilde{\alpha}$ stops by t , where $\tilde{\alpha}(\mathbf{p}, 0) = \alpha(\mathbf{p}, 0)$ and $\tilde{\alpha}(\mathbf{p}, t) = \alpha^*(\mathbf{p})$ for $t > 0$. Therefore, to close the induction step, it suffices to prove that if $\tilde{\alpha}$ stops by t then α^* stops by t . This is immediate if $p_0^1 = p_0^2$. Accordingly, suppose that $p_0^1 \neq p_0^2$ and, without loss of generality, take $p_0^1 < p_0^2$. If $\tilde{\alpha}$ focuses on item 2 in period 0, then the belief processes are the same under $\tilde{\alpha}$ and α^* . Thus it suffices to show that if $\beta = \beta[\alpha^*, 1, 1]$ stops by t then α^* stops by t .¹⁷

To prove the last implication, we distinguish two sets of updating draws. The first set consists of those for which $a_s(p_0^2) > p_0^1$ in every period $s = 0, 1, \dots, t - 1$. For any (ordered) belief pair $\mathbf{p} = (p^1, p^2)$, write $\langle \mathbf{p} \rangle$ for the unordered pair $\{p^1, p^2\}$. (By considering the unordered pairs of beliefs we eliminate the need to keep track of which item has the higher belief and which is randomly chosen at a tie.) For each updating draw in this first set, for each $s = 1, 2, \dots, t$,

$$\langle \mathbf{p}_s(\alpha^*) \rangle = \{p_0^1, a_s(p_0^2)\} \quad \text{and} \quad M_s(\alpha^*) = a_s(p_0^2),$$

¹⁷Recall that $\beta[\alpha^*, 1, 1]$ is the strategy constructed from α^* by making item 1 the item of focus at $t = 0$ and following α^* for every $t > 0$.

and

$$\langle \mathbf{p}_s(\beta) \rangle = \{a_1(p_0^1), a_{s-1}(p_0^2)\} \quad \text{and} \quad M_s(\beta) = a_{s-1}(p_0^2)$$

since $a_s(p_0^2) > p_0^1$ and $a_s(p_0^2) \geq a_1(p_0^1)$ for all $s = 1, 2, \dots, t-1$. Therefore, α^* updates the belief $a_s(p_0^2)$ in all periods $s = 0, \dots, t$, and β updates the belief $a_{s-1}(p_0^2)$ in all periods $s = 1, \dots, t$. Thus, for each updating draw in this first set,

$$\underline{M}_t(\alpha^*) \leq \underline{M}_t(\beta) \leq \overline{M}_t(\beta) \leq \overline{M}_t(\alpha^*),$$

and the induction step holds.

The second set of updating draws consists of those for which there exists a period $s \in \{1, 2, \dots, t-1\}$ in which $a_s(p_0^2) = p_0^1$. (Note that the second set is complementary to the first set.) Let s^* be the minimal such period. For each draw in this second set, we have

$$\langle \mathbf{p}_{s^*+1}(\alpha^*) \rangle = \{p_0^1, a_1(p_0^1)\}$$

since $\langle \mathbf{p}_{s^*}(\alpha^*) \rangle = \{p_0^1, p_0^1\}$ and the belief p_0^1 is updated once by s^* . We also have

$$\langle \mathbf{p}_{s^*+1}(\beta) \rangle = \{a_1(p_0^1), a_{s^*}(p_0^2)\}$$

since the belief p_0^1 is updated once in period 0 and in each period $s = 1, \dots, s^*$, $\langle \mathbf{p}_s(\beta) \rangle = \{a_1(p_0^1), a_{s-1}(p_0^2)\}$ and $a_{s-1}(p_0^2) \geq a_1(p_0^1)$. Thus the strategy β updates the belief $a_{s-1}(p_0^2)$ in all periods $s = 1, \dots, s^*$. Since $a_{s^*}(p_0^2) = p_0^1$, we have

$$\langle \mathbf{p}_{s^*+1}(\beta) \rangle = \{a_1(p_0^1), p_0^1\} = \langle \mathbf{p}_{s^*+1}(\alpha^*) \rangle.$$

Therefore, in each updating draw from the second set, $\langle \mathbf{p}_s(\alpha^*) \rangle = \langle \mathbf{p}_s(\beta) \rangle$ for all $s \geq s^* + 1$. In particular, α^* and β stop in the same period, concluding the proof of the induction step. \square

Proof of Proposition 4. Without loss of generality, we prove the statements for $i = 1$. We claim that

$$D_{\text{ea}}^1(\mathbf{p}) = \begin{cases} \frac{\bar{p}-p^1}{2(\bar{p}-p^2)} \left(1 - \frac{(\bar{p}-p^2)^2}{(\bar{p}-\underline{p})^2}\right) + \frac{p^1-p^2}{\bar{p}-p^2} & \text{if } p^1 \geq p^2, \\ \frac{\bar{p}-p^2}{2(\bar{p}-p^1)} \left(1 - \frac{(\bar{p}-p^1)^2}{(\bar{p}-\underline{p})^2}\right) & \text{if } p^1 \leq p^2. \end{cases} \quad (14)$$

If $p_0^1 = p_0^2 = p$, then α^* stops with beliefs $(p_\tau^1, p_\tau^2) = (\underline{p}, \underline{p})$ —in which case the outside option is chosen—with probability $\left(\frac{\bar{p}-p}{\bar{p}-\underline{p}}\right)^2$. By symmetry, conditional on not choosing the outside

option, the DM chooses item 1 with probability 1/2. Thus

$$D_{\text{ea}}^1(p, p) = \frac{1}{2} \left(1 - \left(\frac{\bar{p} - p}{\bar{p} - \underline{p}} \right)^2 \right).$$

Now consider prior beliefs such that $p_0^1 > p_0^2$. The strategy α^* initially focuses on item 1 until $p_t^1 = \bar{p}$ or $p_t^1 = p_0^2$. In the former case, which occurs with probability $\frac{p_0^1 - p_0^2}{\bar{p} - p_0^2}$, the DM chooses item 1. In the latter case, which occurs with probability $\frac{\bar{p} - p_0^1}{\bar{p} - p_0^2}$, the DM chooses item 1 with probability $D_{\text{ea}}^1(p_0^2, p_0^2)$. Therefore, for $p^1 > p^2$,

$$D_{\text{ea}}^1(p^1, p^2) = \frac{p^1 - p^2}{\bar{p} - p^2} + \frac{\bar{p} - p^1}{\bar{p} - p^2} D_{\text{ea}}^1(p^2, p^2),$$

in agreement with (14). Finally, consider prior beliefs such that $p_0^1 < p_0^2$. The strategy α^* initially focuses on item 2 until $p_t^2 = \bar{p}$ or $p_t^2 = p_0^1$. In the former case, the DM chooses item 2. In the latter case, which occurs with probability $\frac{\bar{p} - p_0^2}{\bar{p} - p_0^1}$, the DM chooses item 1 with probability $D_{\text{ea}}^1(p_0^1, p_0^1)$. Thus, for $p^1 < p^2$ we have

$$D_{\text{ea}}^1(p^1, p^2) = \frac{\bar{p} - p^2}{\bar{p} - p^1} D_{\text{ea}}^1(p^1, p^1),$$

as needed.

Statement 1 of the proposition follows from the fact that $D_{\text{ea}}^1(p^1, p^2)$ is convex in p^1 and strictly convex in p^1 for $[p, p^2]$.

Substituting (14) into the definition of $e^i(\mathbf{p})$ gives, for $p^1 < p^2$,

$$e^1(\mathbf{p}; \bar{p}) = \frac{(1 - 2\lambda)^2 (1 - p^1)^2 (p^1)^2 (\bar{p} - p^2)}{(\bar{p} - p^1)(\bar{p}p^1 - \bar{p} + \lambda p^1 - \bar{p}\lambda(2p^1 - 1))(p^1 - \bar{p}p^1 - \lambda p^1 + \bar{p}\lambda(2p^1 - 1))}. \quad (15)$$

Statement 4 of the proposition follows immediately from this expression. Statements 2 and 3 follow from the monotonicity of the right-hand side of (15) with respect to p^2 and p^1 , respectively. Monotonicity with respect to p^2 is straightforward. For p^1 , one may verify that $\frac{\partial^2}{\partial p^1 \partial \bar{p}} \ln e^1(\mathbf{p}; \bar{p}) < 0$ for $0 < p^1 < \bar{p} < 1$. Thus, it suffices to show that $\ln e^1(\mathbf{p}; \bar{p})$ for $\bar{p} = 1$ increases in p^1 . At $\bar{p} = 1$, (15) simplifies to

$$e^1(\mathbf{p}; 1) = \frac{(2\lambda - 1)^2 (p^1)^2 (1 - p^2)}{(1 - \lambda)\lambda(1 - p^1)},$$

which is indeed increasing in p^1 . □

Proof of Proposition 5. It suffices to prove the statement for all deterministic deadlines T ,

for which we proceed by induction on T . The statement holds vacuously for $T = 0$. Assume for induction that it holds for $T - 1$. Without loss of generality, consider a prior such that $p_0^1 < p_0^2$. Since the induction hypothesis applies regardless of the prior, it suffices to prove that the strategy α^* achieves at least as high a payoff as does $\beta = \beta[\alpha^*, 1, 1]$, recalling that $\beta(\mathbf{p}, 0) = 1$ and $\beta(\mathbf{p}, t) = \alpha^*(\mathbf{p})$ for $t > 0$.

Consider the same coupling construction and the same two sets of updating draws as in the proof of Proposition 3. The first set consists of draws for which $a_s(p_0^2) > p_0^1$ in every period $s = 0, 1, \dots, T$, where $a_s(p_0^2)$ is the belief about item 2 resulting from updating the belief about item 2 in all periods up to s . The payoff achieved by β in this set of draws is $E \max \{z, a_{T-1}(p_0^2)\}$, where the expectation is with respect to $a_{T-1}(p_0^2)$. The strategy α^* achieves a payoff $E \max \{z, a_T(p_0^2)\}$ that is at least as large since $\max \{z, p^2\}$ is convex in p^2 .

The second set of updating draws (complementary to the first set) consists of those for which there exists a period $s \in \{1, 2, \dots, T\}$ in which $a_s(p_0^2) = p_0^1$. Let s^* be the minimal such period. The same argument as in the proof of Proposition 3 implies that the belief $\mathbf{p}_{s^*+1}(\beta)$ formed in period $s^* + 1$ under strategy β is identical to the belief $\mathbf{p}_{s^*+1}(\alpha^*)$ formed in period $s^* + 1$ under α^* . Hence, the two strategies achieve the same payoff in each draw from this set. \square

Proof of Proposition 6. We use the same coupling construction as in the proof of Theorem 1, again denoting by \mathbf{p}_t and $\hat{\mathbf{p}}_t$ the belief sequences generated by strategies α and $\beta = \beta[\alpha, i, m]$, respectively. Similarly, we denote variables generated with strategy α by $\tau, k(i, t)$, etc., and variables generated by β by $\hat{\tau}, \hat{k}(i, t)$, etc.

For each item i and learning process π^i , we say that π^i reaches \bar{p}^i if there exists κ^i such that $\pi_{\kappa^i}^i \geq \bar{p}^i$ and $\underline{p}^i < \pi_{\kappa}^i$ for all $\kappa < \kappa^i$. Similarly, we say that π^i reaches \underline{p}^i if there exists κ^i , such that $\pi_{\kappa^i}^i \leq \underline{p}^i$ and $\pi_{\kappa}^i < \bar{p}^i$ for all $\kappa < \kappa^i$. Let \mathbf{r} be any element of $\{\underline{p}^1, \bar{p}^1\} \times \{\underline{p}^2, \bar{p}^2\} \times \dots \times \{\underline{p}^I, \bar{p}^I\}$ and let $r^i \in \{\underline{p}^i, \bar{p}^i\}$ denote its i th component. That is, \mathbf{r} is a vector that specifies for each item i an upper or lower threshold r^i . We define for each \mathbf{r} a set $\Pi(\mathbf{r})$ consisting of those learning draws $\pi = (\pi^1, \dots, \pi^I)$ for which the process π^i reaches the threshold r^i for each i . Note that the sets $\Pi(\mathbf{r})$ are disjoint and their union over all \mathbf{r} has probability 1.

We distinguish three sets of vectors \mathbf{r} . The first set, R_1 , consists of all \mathbf{r} such that $r^i = \underline{p}^i$ for the target item i , and there exists $j \neq i$ for which $r^j = \bar{p}^j$. Consider $\mathbf{r} \in R_1$ and any draw $\pi \in \Pi(\mathbf{r})$. Regardless of her attention strategy, the DM does not choose the target item i in π (and instead chooses some j for which $r^j = \bar{p}^j$).

The second set, R_2 , consists of all \mathbf{r} such that $r^i = \bar{p}^i$ for the target item i . Consider $\mathbf{r} \in R_2$ and any draw $\pi \in \Pi(\mathbf{r})$. We claim that if α leads to the choice of the target item i in draw π , then so does β . The claim is immediate if π^j reaches \underline{p}^j for each $j \neq i$; accordingly,

suppose that π^j reaches \bar{p}^j for some $j \neq i$. If, under α , the DM chooses i in period τ , then $p_\tau^i \geq \bar{p}^i$ (since it is not the case that the belief for each other item reaches its lower threshold). Assume for contradiction that, under β , the DM chooses $j \neq i$ in some period $\hat{\tau} \leq \tau$, and hence $\hat{p}_{\hat{\tau}}^j \geq \bar{p}^j$. By the Attention Lemma, $k(j, \hat{\tau}) \geq \hat{k}(j, \hat{\tau})$ and thus $p_{\tau'}^j \geq \bar{p}^j$ for some $\tau' \leq \hat{\tau} \leq \tau$. Therefore, α leads to the choice of $j \neq i$ in period $\tau' \leq \tau$, contradicting that α leads to the choice of i in period τ . Again by the Attention Lemma, $k(i, \tau) \leq \hat{k}(i, \tau)$, hence $\hat{p}_\tau^i \geq \bar{p}^i$, and β leads to the choice of i in period τ or earlier.

The third set, R_3 , is the singleton $\{(\underline{p}_1, \dots, \underline{p}_I)\}$. Consider $\mathbf{r} = (\underline{p}_1, \dots, \underline{p}_I)$ and a draw $\pi \in \Pi(\mathbf{r})$. We claim that if β leads to the choice of the target item i in draw π , then so does α . If, under β , the DM chooses i in period $\hat{\tau}$, then $\hat{p}_{\hat{\tau}}^j \leq \underline{p}^j$ for all $j \neq i$. Assume for contradiction that, under α , the DM chooses $j^* \neq i$ in some period $\tau \leq \hat{\tau}$, and hence $p_\tau^{j^*} \leq \underline{p}^{j^*}$ for all $j \neq j^*$. By the Attention Lemma, $k(i, \tau) \leq \hat{k}(i, \tau)$ and thus $\hat{p}_{\tau'}^i \leq \underline{p}^i$ for some $\tau' \leq \tau \leq \hat{\tau}$. Non-wastefulness of β together with the Coupling Lemma imply that $\mathbf{p}_\tau = \hat{\mathbf{p}}_\tau$, and thus β leads to the choice of j^* in period $\tau \leq \hat{\tau}$, contradicting that β leads to the choice of i in period $\hat{\tau}$. Again by the Attention Lemma, $k(i, \hat{\tau}) \leq \hat{k}(i, \hat{\tau})$, which, together with the Coupling Lemma, implies that $p_{\hat{\tau}}^j \leq \underline{p}^j$ for all $j \neq i$ and hence α leads to the choice of i in period $\hat{\tau}$ or earlier.

Observe that, as $\mathbf{p} \rightarrow (0, 1, \dots, 0, 1)$,

$$\Pr(\Pi(\underline{p}^1, \dots, \underline{p}^I) \mid v^i = 1; \mathbf{p}) \leq \frac{\bar{p}^i - p^i \underline{p}^i}{\bar{p}^i - \underline{p}^i p^i} \rightarrow 0$$

$$\text{and } \Pr\left(\bigcup_{\mathbf{r} \in R_2} \Pi(\mathbf{r}) \mid v^i = 0; \mathbf{p}\right) \leq \frac{p^i - \underline{p}^i}{\bar{p}^i - \underline{p}^i} \frac{1 - \bar{p}^i}{1 - p^i} \rightarrow 0$$

Therefore, the sign of $\liminf_{\mathbf{p} \rightarrow (0, 1, \dots, 0, 1)} (\tilde{D}^i(1; \beta, \mathbf{p}) - \tilde{D}^i(1; \alpha, \mathbf{p}))$ is the same as that of

$$\Pr\left(\hat{\mathbf{p}}_{\hat{\tau}} \in F_{\text{no}}^i(\mathbf{p}) \mid \bigcup_{\mathbf{r} \in R_2} \Pi(\mathbf{r}), v^i = 1\right) - \Pr\left(\mathbf{p}_\tau \in F_{\text{no}}^i(\mathbf{p}) \mid \bigcup_{\mathbf{r} \in R_2} \Pi(\mathbf{r}), v^i = 1\right).$$

Result (6) holds since for any draw in $\Pi(\mathbf{r})$ with $\mathbf{r} \in R_2$, $\mathbf{p}_\tau \in F_{\text{no}}^i(\mathbf{p})$ implies $\hat{\mathbf{p}}_{\hat{\tau}} \in F_{\text{no}}^i(\mathbf{p})$.

Similarly, the sign of $\limsup_{\mathbf{p} \rightarrow (0, 1, \dots, 0, 1)} (\tilde{D}^i(0; \beta, \mathbf{p}) - \tilde{D}^i(0; \alpha, \mathbf{p}))$ is the same as that of

$$\Pr(\hat{\mathbf{p}}_{\hat{\tau}} \in F_{\text{no}}^i(\mathbf{p}) \mid \Pi(\underline{p}^1, \dots, \underline{p}^2), v^i = 0) - \Pr(\mathbf{p}_\tau \in F_{\text{no}}^i(\mathbf{p}) \mid \Pi(\underline{p}^1, \dots, \underline{p}^2), v^i = 0).$$

Result (7) holds since for any draw in $\Pi(\underline{p}^1, \dots, \underline{p}^2)$, $\hat{\mathbf{p}}_{\hat{\tau}} \in F_{\text{no}}^i(\mathbf{p})$ implies $\mathbf{p}_\tau \in F_{\text{no}}^i(\mathbf{p})$. \square

Proof of Proposition 7. The proof is similar to that of Proposition 3. The coupling construc-

tion exploits the symmetry with respect to belief $1/2$ as follows. Let

$$\bar{\Pi} = \{\underline{p}[+], \underline{p}[+][+], \dots, \bar{p}[-][-], \bar{p}[-]\} \cap [1/2, 1)$$

and

$$\underline{\Pi} = \{\underline{p}[+], \underline{p}[+][+], \dots, \bar{p}[-][-], \bar{p}[-]\} \cap (0, 1/2).$$

An updating draw is a collection $(\ell(\pi, \kappa))_{(\pi, \kappa) \in \bar{\Pi} \times \mathbb{N}}$, where $\ell(\pi, \kappa)$ is an i.i.d. random variable that attains values $\pi[+]$ and $\pi[-]$ with probabilities $\frac{\pi - \pi[-]}{\pi[+] - \pi[-]}$ and $\frac{\pi[+] - \pi}{\pi[+] - \pi[-]}$, respectively.

For any attention strategy γ , we construct the process of beliefs $\mathbf{p}_t(\gamma)$ in each updating draw as follows. Let $k_0^\pi(\gamma) = 0$ for all $\pi \in \bar{\Pi}$ and $\mathbf{p}_0(\gamma) = \mathbf{p}_0$. In each period t , the focus of attention ι_t in t is chosen according to $\gamma(\mathbf{p}_t, t)$. We distinguish two cases: $p^{\iota_t} \geq 1/2$ and $p^{\iota_t} < 1/2$. If $p^{\iota_t} \geq 1/2$, then

$$k_{t+1}^\pi(\gamma) = \begin{cases} k_t^\pi(\gamma) + 1 & \text{if } \pi = p^{\iota_t}, \\ k_t^\pi(\gamma) & \text{if } \pi \in \bar{\Pi} \setminus \{p^{\iota_t}\}, \end{cases}$$

and

$$p_{t+1}^j(\gamma) = \begin{cases} \ell\left(p_t^j(\gamma), k_t^{p_t^j(\gamma)}(\gamma)\right) & \text{if } j = \iota_t, \\ p_t^j(\gamma) & \text{otherwise.} \end{cases}$$

If $p^{\iota_t} < 1/2$, then

$$k_{t+1}^\pi(\gamma) = \begin{cases} k_t^\pi(\gamma) + 1 & \text{if } \pi = 1 - p^{\iota_t}, \\ k_t^\pi(\gamma) & \text{if } \pi \in \bar{\Pi} \setminus \{1 - p^{\iota_t}\}, \end{cases}$$

and

$$p_{t+1}^j(\gamma) = \begin{cases} 1 - \ell\left(1 - p_t^j(\gamma), k_t^{1-p_t^j(\gamma)}(\gamma)\right) & \text{if } j = \iota_t, \\ p_t^j(\gamma) & \text{otherwise.} \end{cases}$$

By construction, for each strategy γ , the beliefs $\mathbf{p}_t(\gamma)$ follow the law P_γ^{ea} .

We again define a belief process that would result from learning about a single item. For any belief $\pi \in \underline{\Pi} \cup \bar{\Pi}$, let $\hat{\pi} = \max\{\pi, 1 - \pi\}$. Define the belief process $a_t(\pi)$ as follows: $a_0(\pi) = \hat{\pi}$ and $h_0^{\hat{\pi}} = 0$ for all $\hat{\pi} \in \bar{\Pi}$, and for $t > 0$, recursively define

$$a_{t+1}(\pi) = \max\left\{\ell\left(a_t(\pi), h_t^{a_t(\pi)}\right), 1 - \ell\left(a_t(\pi), h_t^{a_t(\pi)}\right)\right\}$$

and, for each $\tilde{\pi} \in \bar{\Pi}$,

$$h_{t+1}^{\tilde{\pi}} = \begin{cases} h_t^{\tilde{\pi}} + 1 & \text{if } \tilde{\pi} = a_t(\pi), \\ h_t^{\tilde{\pi}} & \text{otherwise.} \end{cases}$$

For any attention strategy γ , let $M_t(\gamma) = \max_{i=1,2} \hat{p}_t^i(\gamma)$ and $\bar{M}_t(\gamma) = \max_{s=0,\dots,t} M_s(\gamma)$. The strategy γ stops by period t , i.e., $\tau_{\text{no}}^{\text{ea}}(\gamma) \leq t$, if $\bar{M}_t(\gamma) = \bar{p}$. We will prove by induction on t that, for every prior, in each updating draw, if a strategy α stops by t , then the strategy α^{**} also stops by t .

To see that the statement holds for $t = 1$, note that if $p_0^1 \neq p_0^2$, then $M_0(\alpha) = M_0(\alpha^{**}) = \max\{\hat{p}_0^1, \hat{p}_0^2\}$ and $M_1(\alpha^{**}) = a_1(\max\{\hat{p}_0^1, \hat{p}_0^2\})$ while $M_1(\alpha)$ equals $a_1(\max\{\hat{p}_0^1, \hat{p}_0^2\})$ or $\max\{\hat{p}_0^1, \hat{p}_0^2\}$. (The latter case arises when α focuses on the less certain item in period 0.) Thus

$$\bar{M}_1(\alpha) \leq \bar{M}_1(\alpha^{**}),$$

as needed.

Suppose the statement holds for $t - 1$. If α stops by t then, since the induction hypothesis holds for every prior, strategy $\tilde{\alpha}$ also stops by t , where $\tilde{\alpha}(\mathbf{p}, 0) = \alpha(\mathbf{p}, 0)$ and $\tilde{\alpha}(\mathbf{p}, t) = \alpha^{**}(\mathbf{p})$ for $t > 0$. Therefore, to close the induction step, it suffices to prove that if $\tilde{\alpha}$ stops by t then α^{**} stops by t . This is immediate if $\hat{p}_0^1 = \hat{p}_0^2$ since then the two belief processes coincide. Accordingly, suppose that $p_0^1 \neq p_0^2$ and, without loss of generality, take $\hat{p}_0^1 < \hat{p}_0^2$. If $\tilde{\alpha}$ focuses on item 2 in period 0 then the belief processes coincide under $\tilde{\alpha}$ and α^{**} . Thus it suffices to show that if $\beta = \beta[\alpha^{**}, 1, 1]$ stops by t then α^{**} stops by t .

To prove the last implication, we distinguish two sets of updating draws. The first set consists of those for which $a_s(\hat{p}_0^2) > \hat{p}_0^1$ in every period $s = 0, 1, \dots, t - 1$. For each draw in this first set and each $s = 1, 2, \dots, t$,

$$\{\hat{p}_s^1(\alpha^{**}), \hat{p}_s^2(\alpha^{**})\} = \{\hat{p}_0^1, a_s(p_0^2)\} \quad \text{and} \quad M_s(\alpha^{**}) = a_s(p_0^2),$$

and

$$\{\hat{p}_s^1(\beta), \hat{p}_s^2(\beta)\} = \{a_1(p_0^1), a_{s-1}(p_0^2)\} \quad \text{and} \quad M_s(\beta) = a_{s-1}(p_0^2)$$

since $a_s(p_0^2) > \hat{p}_0^1$ and $a_s(p_0^2) \geq a_1(p_0^1)$ for all $s = 1, 2, \dots, t - 1$. Therefore, α^{**} updates $a_s(p_0^2)$ in all periods $s = 0, \dots, t$, and β updates $a_{s-1}(p_0^2)$ in all periods $s = 1, \dots, t$. Thus, for each updating draw in this first set,

$$\bar{M}_t(\beta) \leq \bar{M}_t(\alpha^{**}),$$

and the induction step holds.

The second set of updating draws consists of those for which there exists a period $s \in$

$\{1, 2, \dots, t-1\}$ in which $a_s(p_0^2) = \hat{p}_0^1$. (Note that the second set is complementary to the first.) Let s^* be the minimal such period and observe that

$$\{\hat{p}_{s^*+1}^1(\alpha^{**}), \hat{p}_{s^*+1}^2(\alpha^{**})\} = \{\hat{p}_0^1, a_1(\hat{p}_0^1)\}$$

since $\{\hat{p}_{s^*}^1(\alpha^{**}), \hat{p}_{s^*}^2(\alpha^{**})\} = \{\hat{p}_0^1, \hat{p}_0^1\}$ and the belief with value \hat{p}_0^1 or $1 - \hat{p}_0^1$ is updated once by s^* . Also,

$$\{\hat{p}_{s^*+1}^1(\beta), \hat{p}_{s^*+1}^2(\beta)\} = \{a_1(p_0^1), a_{s^*}(p_0^2)\}$$

since the belief p_0^1 is updated once in period 0 and, for each period $s = 1, \dots, s^*$, $\{\hat{p}_s^1(\beta), \hat{p}_s^2(\beta)\} = \{a_1(p_0^1), a_{s-1}(p_0^2)\}$ and $a_{s-1}(p_0^2) \geq a_1(p_0^1)$. Since $a_{s^*}(p_0^2) = \hat{p}_0^1$, we have

$$\{\hat{p}_{s^*+1}^1(\beta), \hat{p}_{s^*+1}^2(\beta)\} = \{a_1(p_0^1), \hat{p}_0^1\} = \{\hat{p}_{s^*+1}^1(\alpha^{**}), \hat{p}_{s^*+1}^2(\alpha^{**})\}.$$

Therefore, in each updating draw from the second set, $\{\hat{p}_s^1(\beta), \hat{p}_s^2(\beta)\} = \{\hat{p}_s^1(\alpha^{**}), \hat{p}_s^2(\alpha^{**})\}$ for all $s \geq s^* + 1$. In particular, α^{**} and β stop in the same period, concluding the proof of the induction step. \square

A.3 Computation of the ex ante demand from Section 5

We characterize D_{no}^i for $i = 1$; the case of $i = 2$ is symmetric. By symmetry of α^{**} , $D_{\text{no}}^1(p^1, p^2; \alpha^{**}) = D_{\text{no}}^2(p^2, p^1; \alpha^{**}) = 1 - D_{\text{no}}^1(p^2, p^1; \alpha^{**})$, and thus it suffices to compute $D_{\text{no}}^1(p^1, p^2; \alpha^{**})$ for $p_1 \geq p_2$ only. Demand D_{no}^1 for this set of beliefs is characterized in (16)–(19) below.

Consider beliefs on the two main diagonals. First, symmetry implies that $D_{\text{no}}^1(p, p; \alpha^{**}) = 1/2$ whenever $\underline{p} < p < \bar{p}$. Second, consider prior beliefs of the form $(p_0^1, p_0^2) = (p, 1 - p)$, where $1/2 < p < \bar{p}$. Starting from such a prior, if the strategy α^{**} never leads to beliefs $(1/2, 1/2)$, then it stops either with $p_\tau^2 = \underline{p}$ or with $p_\tau^1 = \bar{p}$; in either case, the DM chooses item 1. If α^{**} does lead to beliefs $(1/2, 1/2)$, then the DM chooses item 1 with probability $1/2$. Therefore, if $1/2 < p < \bar{p}$, $D_{\text{no}}^1(p, 1 - p; \alpha^{**}) = 1 - x + x/2$ where x is the probability that the DM's beliefs reach $(1/2, 1/2)$.

The value of x may be computed as follows. Consider the event that π_κ^1 stops at \underline{p} and π_κ^2 stops at \bar{p} . The ex ante probability of this event is

$$\frac{\bar{p} - p}{\bar{p} - \underline{p}} \frac{1 - p - \underline{p}}{\bar{p} - \underline{p}} = \left(\frac{\bar{p} - p}{\bar{p} - \underline{p}} \right)^2,$$

where the equality follows from $\underline{p} = 1 - \bar{p}$. Since this event can occur only if the DM's beliefs

reach $(1/2, 1/2)$, in which case it occurs with probability $1/4$, we have

$$\left(\frac{\bar{p} - p}{\bar{p} - \underline{p}}\right)^2 = \frac{x}{4}$$

and

$$D_{\text{no}}^1(p, 1 - p; \alpha^{**}) = 1 - 2 \left(\frac{\bar{p} - p}{\bar{p} - \underline{p}}\right)^2$$

whenever $1/2 < p < \bar{p}$.

Now consider beliefs in the set $\{\mathbf{p} : p_1 \geq p_2\}$. We partition this set into four disjoint subsets:

1. For $\{\mathbf{p} : p^1 \in (\underline{p}, 1/2], p^2 \in (\underline{p}, p^1]\}$, α^{**} initially focuses on item 2 until p_t^2 reaches \underline{p} or p^1 . Thus, for \mathbf{p} in this subset,

$$D_{\text{no}}^1(\mathbf{p}; \alpha^{**}) = \frac{p^1 - p^2}{p^1 - \underline{p}} + \frac{p^2 - \underline{p}}{p^1 - \underline{p}} D_{\text{no}}^1(p^1, p^1; \alpha^{**}). \quad (16)$$

2. For $\{\mathbf{p} : p^1 \in (1/2, \bar{p}], p^2 \in (\underline{p}, 1 - p^1]\}$, α^{**} initially focuses on item 2 until p_t^2 reaches \underline{p} or $1 - p^1$. Thus, for \mathbf{p} in this subset,

$$D_{\text{no}}^1(\mathbf{p}; \alpha^{**}) = \frac{1 - p^1 - p^2}{1 - p^1 - \underline{p}} + \frac{p^2 - \underline{p}}{1 - p^1 - \underline{p}} D_{\text{no}}^1(p^1, 1 - p^1; \alpha^{**}). \quad (17)$$

3. For $\{\mathbf{p} : p^1 \in (1/2, \bar{p}], p^2 \in (1 - p^1, 1/2]\}$, α^{**} initially focuses on item 1 until p_t^1 reaches $1 - p^2$ or \bar{p} . Thus, for \mathbf{p} in this subset,

$$D_{\text{no}}^1(\mathbf{p}; \alpha^{**}) = \frac{p^1 - 1 + p^2}{\bar{p} - 1 + p^2} + \frac{\bar{p} - p^1}{\bar{p} - 1 + p^2} D_{\text{no}}^1(1 - p^2, p^2; \alpha^{**}). \quad (18)$$

4. For $\{\mathbf{p} : p^1 \in (1/2, \bar{p}], p^2 \in (1/2, p^1]\}$, α^{**} initially focuses on item 1 until p_t^1 reaches p^2 or \bar{p} . Thus, for \mathbf{p} in this subset,

$$D_{\text{no}}^1(\mathbf{p}; \alpha^{**}) = \frac{p^1 - p^2}{\bar{p} - p^2} + \frac{\bar{p} - p^1}{\bar{p} - p^2} D_{\text{no}}^1(p^2, p^2; \alpha^{**}). \quad (19)$$

A.4 Proof of Theorem 2

As in the proof of Theorem 1, for any variable η in the baseline process, let $\hat{\eta}$ denote its counterpart in the manipulated process. We construct the probability space and the baseline

and manipulated processes $(\mathbf{p}_t, \iota_t)_t$ as in the proof of Theorem 1. Note that Lemmata 1, 2, and 3 extend verbatim to the current setting.

Given any draw, let \mathcal{I}^* denote the set of items deemed optimal in the baseline process and $\hat{\mathcal{I}}^*$ the corresponding set in the manipulated process.

Lemma 6 (Choice Lemma for Multiple Values). *Let i be the target item. For any draw,*

1. *if $i \in \mathcal{I}^*$, then (i) $i \in \hat{\mathcal{I}}^*$, (ii) $\hat{\mathcal{I}}^* \subseteq \mathcal{I}^*$, and (iii) $\hat{\tau} \leq \tau$.*
2. *for any $j \neq i$, if $j \in \hat{\mathcal{I}}^*$, (i) $j \in \mathcal{I}^*$, (ii) $\mathcal{I}^* \subseteq \hat{\mathcal{I}}^*$, and (iii) $\tau \leq \hat{\tau}$.*

Proof. Statement 1: Consider any draw in which the baseline process deems the target item i optimal (i.e., $i \in \mathcal{I}^*$) with stopping time τ . First suppose $\hat{\tau} \geq \tau$. By the attention lemma, $\hat{k}(i, \tau) \geq k(i, \tau)$. Since α deems i optimal and stops at τ , the DM is sufficiently certain of i 's value at τ . Hence, by non-wastefulness I, $\hat{k}(i, \tau) = k(i, \tau)$, and the two processes meet at τ . Thus $\hat{\mathcal{I}}^* = \mathcal{I}^*$, $\hat{\tau} = \tau$, and (i–iii) hold.

Now consider any draw in which $\hat{\tau} < \tau$. We first show that, under the manipulated process, the DM does not deem any item $j \neq i$ optimal. Suppose for contradiction that $j \in \hat{\mathcal{I}}^*$. Then there exists v such that $\hat{p}_\tau^j \in C^j(v)$ and for all $j' \neq j$, we have $\hat{p}_\tau^{j'} \in C^{j'}(v)$ or $\hat{p}_\tau^{j'} \in L^{j'}(v)$. By the Attention Lemma, $k(j'', \hat{\tau}) \geq \hat{k}(j'', \hat{\tau})$ for all $j'' \neq i$. By the Coupling Lemma, there exists a period $t \leq \hat{\tau}$ such $\mathbf{p}_t^{-i} = \hat{\mathbf{p}}_\tau^{-i}$ and $k(-i, t) = \hat{k}(-i, \hat{\tau})$. Non-wastefulness I and II imply that the baseline strategy does not focus on items $j'' \neq i$ in periods $s > t$, and thus $k(-i, \hat{\tau}) = k(-i, t) = \hat{k}(-i, \hat{\tau})$. Therefore, the baseline and manipulated processes meet at $\hat{\tau}$. By the Meeting Lemma, $\mathbf{p}_\tau = \hat{\mathbf{p}}_\tau$, which implies that $\tau \leq \hat{\tau}$, contradicting the assumption that $\hat{\tau} < \tau$. Therefore, if the manipulated process stops at $\hat{\tau} < \tau$, then $\hat{\mathcal{I}}^* = \{i\}$ and properties (i–iii) again hold.

Statement 2: Consider any draw in which the manipulated process deems an item $j \neq i$ optimal (i.e., $j \in \hat{\mathcal{I}}^*$) with stopping time $\hat{\tau}$. First suppose $\tau \geq \hat{\tau}$. By the Attention Lemma, $k(-i, \hat{\tau}) \geq \hat{k}(-i, \hat{\tau})$. By the Coupling Lemma, there exists $t \leq \hat{\tau}$ such that $k(-i, t) = \hat{k}(-i, \hat{\tau})$ and $\mathbf{p}_t^{-i} = \hat{\mathbf{p}}_\tau^{-i}$. Since, under the manipulated strategy, the process stops at $\hat{\tau}$ and the DM deems j optimal, it must be that for some v , $p_\tau^j \in C^j(v)$ and for each $j' \neq j$, $p_\tau^{j'} \in C^{j'}$ or $p_\tau^{j'} \in L^{j'}(v)$. Thus, by non-wastefulness I and II, the baseline strategy focuses only on item i in each period $s \geq t$, and hence $k(-i, \hat{\tau}) = k(-i, t) = \hat{k}(-i, \hat{\tau})$. Therefore, the two processes meet at $\hat{\tau}$, $\hat{\mathcal{I}}^* = \mathcal{I}^*$, $\hat{\tau} = \tau$, and (i–iii) hold.

Now consider any draw in which $\tau < \hat{\tau}$. We first show that, under the baseline process, the DM does not deem item i optimal. Suppose for contradiction that $i \in \mathcal{I}^*$. Then, under the baseline process, in period τ the DM is sufficiently certain of the value of i . By the Attention Lemma, $\hat{k}(i, \tau) \geq k(i, \tau)$. Non-wastefulness I implies that $\hat{k}(i, \tau) = k(i, \tau)$.

Therefore, the two processes meet at τ . By the Meeting Lemma, $\mathbf{p}_\tau = \hat{\mathbf{p}}_\tau$, which implies that $\hat{\tau} \leq \tau$, contradicting the assumption that $\tau < \hat{\tau}$. Therefore, if the baseline process stops at $\tau < \hat{\tau}$, then $i \notin \mathcal{I}^*$.

Next, observe that $k(-i, \tau) = \hat{k}(-i, \hat{\tau})$. Otherwise, one of the two processes $\gamma \in \{\alpha, \beta\}$ focuses on items other than i in fewer periods by $\tau(\gamma)$ than the other process γ' does by $\tau(\gamma')$; that is, $k(-i, \tau(\gamma'), \gamma') > k(-i, \tau(\gamma), \gamma)$. By the Coupling Lemma, there exists a period t such that the process γ' does not stop by t , $k(-i, t; \gamma') = k(-i, \tau(\gamma); \gamma)$ and $\mathbf{p}_t^{-i}(\gamma') = \mathbf{p}_{\tau(\gamma)}^{-i}(\gamma)$. Since γ stops at $\tau(\gamma)$ and the DM deems $j \neq i$ optimal, it must be that, for some v , $p_{\tau(\gamma)}^j \in C^j(v)$ and for each item $j' \neq i$, either $p_{\tau(\gamma)}^{j'} \in C^{j'}$ or $p_{\tau(\gamma)}^{j'} \in L^{j'}(v)$. Therefore, non-wastefulness I and II imply that the strategy γ' focuses only on item i in each period $s > t$. Hence $k(-i, \tau(\gamma'), \gamma') = k(-i, t; \gamma') = \hat{k}(-i, \tau(\gamma), \gamma)$.

Therefore, $\mathbf{p}_\tau^{-i} = \hat{\mathbf{p}}_\tau^{-i}$. Since i is not the unique optimal item under either process, there exists some v and a nonempty set $\mathcal{I}^{-i,*} \subseteq \mathcal{I} \setminus \{i\}$ such that $p_\tau^j = \hat{p}_\tau^j \in C^j(v)$ for each $j \in \mathcal{I}^{-i,*}$, and $p_\tau^j = \hat{p}_\tau^j \in L^j(v)$ for each $j \neq i$ such that $j \notin \mathcal{I}^{-i,*}$. Since the DM does not deem i optimal under the baseline process, $p_\tau^i \in L^i(v)$ and $\mathcal{I}^* = \mathcal{I}^{-i,*}$. The DM may deem i optimal under the manipulated process, and hence either $\hat{\mathcal{I}}^* = \mathcal{I}^{-i,*}$ or $\hat{\mathcal{I}}^* = \mathcal{I}^{-i,*} \cup \{i\}$. Therefore, either $\hat{\mathcal{I}}^* = \mathcal{I}^*$ or $\hat{\mathcal{I}}^* = \mathcal{I}^* \cup \{i\}$, and properties (i-iii) hold.

For (10) and (11), observe that optimality and monotonicity of σ and statements 1 and 2 of Lemma 6 imply, respectively, that $\sigma(i, \hat{\mathcal{I}}^*) \geq \sigma(i, \mathcal{I}^*)$ and $\sigma(j, \hat{\mathcal{I}}^*) \leq \sigma(j, \mathcal{I}^*)$ in each draw. For the statements about first-order stochastic dominance, let $h^j(t) = 1_{\tau \leq t} \sigma(j, \mathcal{I}^*)$ and $\hat{h}^j(t) = 1_{\hat{\tau} \leq t} \sigma(j, \hat{\mathcal{I}}^*)$. Statement 1 of Lemma 6 implies that $h^i(t) \leq \hat{h}^i(t)$ for the target item i and any period t . Statement 2 of the lemma implies that $h^j(t) \geq \hat{h}^j(t)$ for all items $j \neq i$ and any period t . Taking expectations across draws yields the result. \square