

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

December 3, 2020

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. A similar result applies to manipulation of choices in bandit problems.

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 discussions with Ian Krajbich, Arina Nikandrova, Marco Ottaviani, Ryan Webb, Andy Zapechelnuk, and participants at various seminars and conferences. 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 it directly affects 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. We consider a decision-maker who 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 noisy information about its value. Due to cognitive (or other) limitations, the decision-maker can focus on only one item at a time; while she pays attention to a given item, her assessment of its value evolves stochastically, while her assessments of the other items remain the same. In our main model, once her assessment of an item is sufficiently high (according to an exogenously fixed threshold), she stops and chooses that item, and if her assessments of all items are sufficiently low, she stops and chooses an outside option. 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. Starting from a given strategy governing the decision-maker’s attention, we introduce an attention-grabbing manipulation that induces the decision-maker to focus more on one “target” item (perhaps only for a limited time). We show that, under general conditions, such a manipulation increases demand for the target item.

The decision-maker’s learning is governed by an *attention strategy* that maps assessments 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 assessments 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; the baseline strategy is the one the decision-maker employs in the absence of manipulation, while the manipulated strategy induces the decision-maker to increase her focus on a target

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).

item at each profile of assessments and in each period, but agrees with the baseline strategy conditional on not focusing on the target item.

We show that manipulation of attention increases demand and decreases the time to decision in favor of the target item (in the sense of first-order stochastic dominance). This increase in demand comes at the expense of each of the other items, for which demand decreases and the time to decision increases. These results hold for any realization of the items' values. In particular, manipulating attention increases demand even if the target item is worse than the other items and the outside option.

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 the sequences 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 that, for some attention strategies, temporary 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 on the attention strategy. First, the attention strategy should be stationary: focus in each period must depend only on the current assessment, not on the current time. The second assumption is a form of *independence of irrelevant alternatives* (IIA): conditional on not focusing on the target 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.

A similar but simpler result applies to manipulation of choices in multi-armed bandit problems. In this framework, the decision-maker chooses an arm in each period and revises her assessment of an arm after each time she chooses it. We show that, under general

conditions, manipulating choice toward a target arm never reduces the number of times it is chosen up to any period. Even if the manipulation ends after a limited time, the decision-maker’s subsequent choices never overcome the initial increase in the choice of the target arm. In particular, introducing one arm before the others cannot lead to a reduction in the total number of times that arm is chosen up to any given period.

To formalize these intuitions, we rely on a technique known in probability theory as *coupling* (as in, e.g., Lindvall, 1992). 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*; that is, a draw is a state in the joint probability space. Given any draw, the decision-maker’s strategy determines the realization of the learning process. The strategy affects only the timing of the process for each item, not the signal realizations or their ordering. We show that, with this coupling, 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 5.1 and Appendix B. Both assumptions, though restrictive, are automatically satisfied if the attention strategy is optimized given the stopping rules in our main model: we prove that, for a general class of learning costs, optimal attention 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 motivates the stopping rule employed by the decision-maker in our main model as a form of satisficing behavior. 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. Our stopping rule implicitly assumes that the value of the outside option is not so low that the decision-maker would reject it in favor of an item she has not inspected, nor so high that she would choose it in favor of such an item. If there were no outside option (or the value of the outside option was low), 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. Moreover, 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.⁴

Our results can operate on a variety of timescales. Over a short time horizon, we envision a consumer choosing among several brands at the point of purchase. In this case, learning may consist primarily of introspection about the current desirability of each item, including recall of past experience consuming the item. Attention can be manipulated through placement of products on the shelf or webpage, or through previous exposure to advertising that increases the likelihood of a particular brand coming to mind (e.g., through the recency effect). Over a longer time horizon, a decision-maker could be repeatedly choosing among brands of some item and learning about her taste for each brand through her experience consuming it. In this case, manipulation in our model corresponds to any change that boosts consumption (temporarily or permanently) of a particular brand, such as introductory sale pricing. Our results also apply to changes in the time at which items are introduced: one item entering the market before the others is equivalent to manipulation in favor of that item at the beginning of the process. Clinical trials, which we discuss in Section 6.1, are another application of our model with a relatively long time horizon.

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 assessment of the item. The decision-maker need not be Bayesian; we can, for instance, interpret her assessments 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. We also allow for general attention strategies, whether optimal with respect to some cost structure or not.

When applied to advertising, our model has distinctive predictions compared to theories of informative advertising (Stigler, 1961; Telser, 1964), signaling (Nelson, 1974), complementary advertising (Stigler and Becker, 1977; Becker and Murphy, 1993), and persuasive advertising (Braithwaite, 1928; Kaldor, 1950). Unlike these theories, manipulation in our model is not associated with a systematic change in the consumer’s assessments of the values of the items. Data combining consumers’ choices and beliefs can therefore be used to test among these theories. For example, Atalay et al. (2012) find experimentally that placement of a product in the horizontal center of a display is associated with increased likelihood of choice but not with inferences about the quality of the corresponding brand. They argue that the effect of horizontal placement is due to greater attention being paid to central items, as in our model. Our theory also provides novel predictions regarding timing. First, manipulation reduces the time to choose the target item and increases the time for other items. Second,

⁴This is reminiscent of the experimental 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.

as noted above, earlier introduction of an item is associated with an increase in demand.

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). Manipulation of attention in our model could be exogenous in this sense, while the baseline attention strategy could be endogenous.

Our model builds on a long tradition in statistics and economic theory originating with 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 (2019) 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.⁵ Krajbich, Armel, and Rangel (2010) explicitly incorporate attention choice in this model and introduce an exogenous bias in the accumulated signal toward

⁵Ratcliff’s model is essentially equivalent to a sequential sampling model along the lines of Wald (1945).

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 drift-diffusion 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). In addition to applying our main result to bandit problems, we exploit this connection to show that optimal attention strategies in our model of one-shot choice satisfy IIA by using the Gittins index characterization.

2 Simplified setting

Before presenting the model in its full generality, we introduce the main ideas in two particularly simple environments in which temporarily manipulating attention toward an item increases its demand. These two environments differ in how choices of items relate to learning, but the mechanism behind the effect of manipulation is the same.

In both environments, a decision-maker (DM) learns sequentially about two items $j \in \{1, 2\}$ with unknown values $v^j \in \{-1, 1\}$. In each period $t = 0, 1, \dots$, the DM focuses on an item $\iota_t \in \{1, 2\}$. Focusing on an item generates a signal x_t about the value of that item which is independent of the value of the other item. Conditional on the values, signals are independent across periods. In this section, the signal takes a simple form: for each item, the possible signal realizations are -1 and 1 , and for each value v , the realized signal about an item of value v is equal to v with probability $\lambda > 1/2$; that is, $\Pr(x_t = 1 \mid v^{\iota_t} = 1) = \lambda = \Pr(x_t = -1 \mid v^{\iota_t} = -1)$.

For each j and t , let

$$p_t^j = \sum_{s < t: \iota_s = j} \log \frac{\Pr(x_s \mid v^j = 1)}{\Pr(x_s \mid v^j = -1)} = \log \left(\frac{\lambda}{1 - \lambda} \right) \sum_{s < t: \iota_s = j} x_s. \quad (1)$$

Thus p_t^j is the log-likelihood ratio (LLR) comparing $v^j = 1$ to $v^j = -1$ given the observed signals about item j up to the beginning of period t . We write \mathbf{p}_t for the pair (p_t^1, p_t^2) . Attention allocation is governed by a (pure) *attention strategy* $\alpha : \mathbb{R}^2 \rightarrow \{1, 2\}$ that specifies

the item of focus $\iota_t = \alpha(\mathbf{p}_t)$ for each pair \mathbf{p}_t of LLRs.⁶

2.1 Manipulation of attention

In order to understand how manipulation of attention affects choice, we first examine how it affects subsequent attention allocation.⁷ To this end, given a *baseline strategy* β , we introduce a *manipulated strategy* μ constructed from β by making item 1 the item of focus in the initial period and then returning to β in all subsequent periods. That is, the item of focus $\mu(\mathbf{p}, t)$ in period t for LLRs \mathbf{p} is given by

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

For each t , let k_t denote the *cumulative focus* on item 1 before period t under the baseline strategy β ; that is,

$$k_t = |\{s < t : \beta(\mathbf{p}_s) = 1\}|.$$

Similarly, let \hat{k}_t denote the analogous cumulative focus under the manipulated strategy μ .

While manipulation clearly increases attention to item 1 in the first period, subsequent allocation of attention could, in principle, more than compensate for this effect and cause item 1 to eventually receive less attention than it would without manipulation. The following result indicates that this cannot happen: manipulation never decreases the attention allocated to item 1.

Proposition 1. *For each $t \geq 1$ and all pairs of values \mathbf{v} , \hat{k}_t (weakly) first-order stochastically dominates k_t .*⁸

This proposition is a special case of Proposition 3; accordingly, we omit its proof.

To compare the allocation of attention under the two strategies, we “couple” the resulting processes governing the LLRs; that is, we define them on a common probability space in a way that enables us to make comparisons realization-by-realization. In this section, we describe this construction informally; a precise treatment is provided in Section 3.1.

⁶One could equivalently formulate the model in terms of a Bayesian decision-maker whose attention strategy is a function of posterior beliefs about the values.

⁷We do not explicitly model how attention is manipulated. 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).

⁸Recall that a random variable σ (weakly) first-order stochastically dominates another random variable σ' if $\Pr(\sigma \leq t) \leq \Pr(\sigma' \leq t)$ for every t . We henceforth drop the “weakly” qualifier.

For this construction, fix \mathbf{v} , and imagine that there is a large (countably infinite) deck of cards for each item, with each card showing a signal realization of -1 or 1 (with probabilities as described above). Beginning with all cards in each deck face down, in each period t , the baseline attention strategy selects a deck from which to draw the next card according to the item of focus $\iota_t = \beta(\mathbf{p}_t)$. The DM updates the relevant LLR based on the signal shown on the card that was drawn. Now consider the effect of manipulation 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, suppose that, absent manipulation, the DM first draws from deck 2. Manipulation can have persistent and complicated effects on attention allocation by taking the DM to different path of LLRs \mathbf{p}_t . And yet, as Proposition 1 indicates, the subsequent allocation of attention can never overtake the direct effect of manipulation: the cumulative attention devoted to the target item can only increase as a result of manipulation.

The key observation is that, once the ordering of the cards is fixed, we only need to keep track of how many cards the baseline and manipulated strategies draw from each deck, and can disregard the order in which the decks are chosen. 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, LLRs and attention allocation coincide under the two processes following the period in which they meet. Either way, regardless of the ordering of the cards, for any baseline attention strategy, and up to any period t , cumulative focus on item 1 is at least as large under the manipulated strategy as under the baseline strategy, and correspondingly, the cumulative focus on item 2 is at least as large under the baseline strategy as under the manipulated one.

Our interest in how manipulation affects attention allocation is driven by its consequent effect on choice. The following two subsections present settings in which increased attention leads to an unambiguous increase in demand.

2.2 One-shot choice

First, we consider a DM who decides when to stop learning and make a one-shot choice among the two items $j \in \{1, 2\}$ and an outside option with a known value. We focus on a simple stopping rule that adapts Simon’s (1955) model of satisficing to allow for gradual learning and the presence of an outside option. If, at any point in the process, the DM has

collected enough evidence that one of the items is of high value, she stops and chooses that item. If, on the other hand, she has collected enough evidence that both items are of low value, she stops and chooses neither item (i.e., chooses the outside option). Accordingly, we introduce thresholds $\underline{p} < 0 < \bar{p}$ and define stopping regions $F^j = \{\mathbf{p} : p^j \geq \bar{p}\}$ for $j = 1, 2$, $F^{oo} = \{\mathbf{p} : p^j \leq \underline{p} \text{ for } j = 1, 2\}$ and $F = F^1 \cup F^2 \cup F^{oo}$. Learning stops in period $\tau = \min\{t : \mathbf{p}_t \in F\}$ with the DM choosing item j if $\mathbf{p}_\tau \in F^j$ and the outside option if $\mathbf{p}_\tau \in F^{oo}$. (Stopping in F^1 , F^2 , and F^{oo} are mutually exclusive.) See Figure 1 for an illustration.⁹

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

The above rules specify, for any given pair of values \mathbf{v} , the joint stochastic process of LLRs and focus of attention $(\mathbf{p}_t, \iota_t)_t$; we denote by $P_\alpha^\mathbf{v}$ their joint law. For any strategy α and pair of values \mathbf{v} , we let the *interim demand* for item j ,

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

be the probability that the DM stops with the choice of j when the true values are \mathbf{v} .

We say that an attention strategy α is *non-wasteful* if $\alpha(\mathbf{p}) \neq j$ whenever \mathbf{p} is such that $p^j \leq \underline{p}$ and $p^{j'} \geq \underline{p}$ for $j' \neq j$. Non-wasteful strategies do not focus on an item that the DM deems to have low value.

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

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

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

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

and

2. *accelerates the choice of item 1 and decelerates the choice of item 2; that is, $\tau^1(\beta)$*

⁹This process is an extension of the sequential probability ratio test of Wald (1945) from two to multiple hypotheses. It also extends the drift-diffusion model of Ratcliff (1978) to allow for choice among three options (including the outside option). In a non-Bayesian, prior-free setting, results on optimal stopping for multiple-hypothesis sequential sampling are scarce, presumably due to the lack of an obvious objective and the complexity of the problem. An exception is the asymptotic literature originating in Chernoff (1959), who shows that a simple threshold stopping rule is approximately optimal as the cost of sampling vanishes.

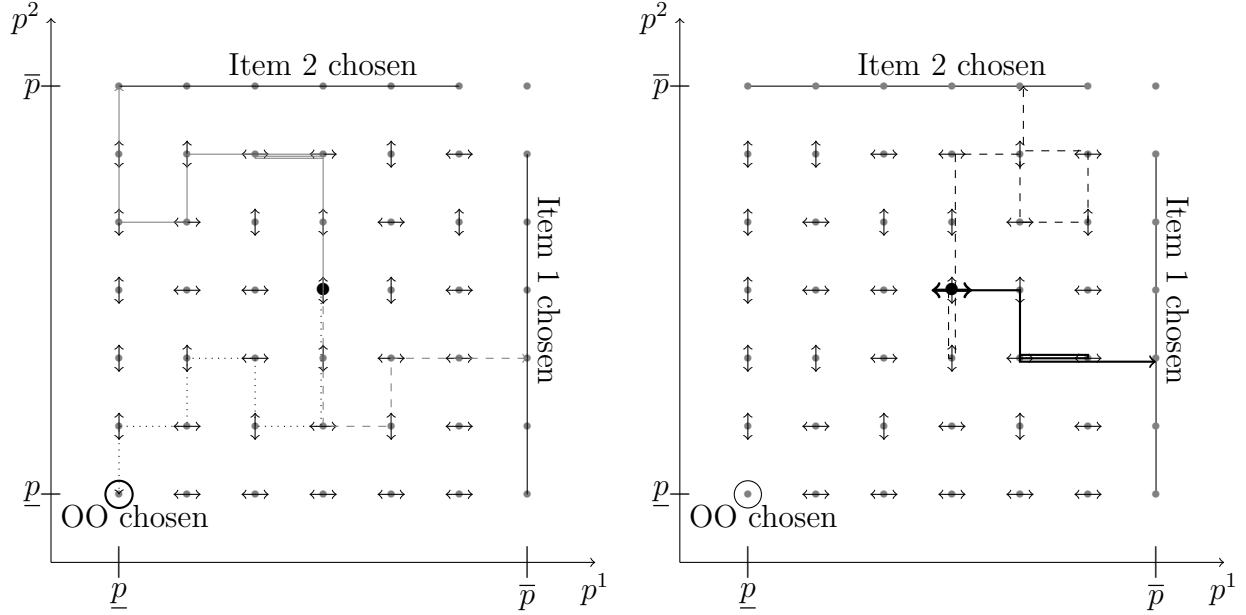


Figure 1: *Left*: One possible strategy, together with three possible paths of log-likelihood ratios. Each dot depicts a pair (p^1, p^2) that can be reached following some sequence of signals. The arrows at each dot specify the direction of movement, horizontal or vertical, according to whether the strategy focuses on item 1 or 2, respectively. The bold dot identifies the starting point $(0, 0)$. At each step of the learning process, (p^1, p^2) moves to one of the two adjacent points in the direction of the arrows. The dashed path 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. *Right*: Manipulation of attention toward item 1 in the first period can alter the entire path and the resulting choice. The dashed path results from the baseline strategy and the solid path from the manipulated strategy for a given realization of the sequence of signals for each item (in accordance with the coupling construction described in Section 2.1).

first-order stochastically dominates $\tau^1(\mu)$ and $\tau^2(\mu)$ first-order stochastically dominates $\tau^2(\beta)$.

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

Proposition 2 is a special case of Proposition 4 in the next section, and thus we provide here only an informal sketch of the argument.

Suppose that the baseline strategy leads to the choice of item 1; that is, the LLR for item 1 reaches \bar{p} in some period τ before that for item 2 does. Following the decks of cards analogy described above, consider a fixed realization of the sequence of the cards in the two decks. Note first that, for this realization, regardless of the attention strategy, the LLR for

item 1 can never reach \underline{p} before it reaches \bar{p} . In particular, under the manipulated strategy, the outside option is not chosen.

As explained above, compared to the baseline strategy, the manipulated strategy must be at least as far ahead with deck 1 and no further ahead with deck 2 in each period $t < \tau$. Thus it must be that under the manipulated strategy, the LLR for item 1 reaches \bar{p} before the ratio for 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. Since this argument applies for every ordering of the cards in the two decks, it also holds when averaging across sequences conditional on the items' values (or, for that matter, conditional on any other event). The argument for the statements about item 2 is symmetric.

2.3 Repeated choice

We show here that manipulation of attention has an unambiguous effect on choice in bandit problems that is similar to the one-shot case. Suppose the DM chooses in each period $t = 0, 1, \dots$ one of two arms $j = 1, 2$ of fixed types $v^j \in \{-1, 1\}$. For example, the two arms could represent different brands of a product for which the DM has unit demand in each period, and v^j could represent the DM's taste for brand j . Experiencing a brand allows the consumer to learn about her taste for it. The choice of arm ι_t in period t results in a stochastic outcome $x_t \in \{-1, 1\}$, where $\Pr(x_t = 1 \mid v^{\iota_t} = 1) = \lambda = \Pr(x_t = -1 \mid v^{\iota_t} = -1)$. In this case, there is no longer any distinction between attention and choice: the DM learns about an arm in a given period if and only if that is the arm she chooses.

Suppose the DM employs a strategy $\beta(\mathbf{p}_t)$ that depends only on the pair of LLRs. Consider a manipulated strategy μ that chooses arm 1 in the first period and is identical to β thereafter. It follows immediately from Proposition 1 that, regardless of the realization of \mathbf{v} , up to any given period, the distribution of the number of times arm 1 is chosen under μ first-order stochastically dominates that under β .

3 Attention and its manipulation

We now extend the setting to allow for more than two items, general signal structures, and stochastic attention strategies. Each item $j \in \mathcal{I} = \{1, \dots, I\}$ has a fixed unknown value $v^j \in V$, where $V \subset \mathbb{R}$ is finite. As above, we let $\mathbf{v} = (v^1, \dots, v^I)$.

The DM forms an *assessment* $p_t^j \in A$ about each item j in each period t , where A is a countable set. For instance, if the set of values V is binary, the assessment p_t^j could be the

log-likelihood ratio comparing the probabilities of the realized signals across the two values. Similarly, for a Bayesian DM, p_t^j could be the posterior belief about item j (in which case $A \subset \Delta(V)$), or the posterior expected value of item j at the beginning of period t (in which case $A \subset \mathbb{R}$). The model also allows for the case in which an assessment of an item consists of the sequence of signals received about it.¹⁰ Finally, in a neuroeconomic context, p_t^j could represent the level of neural activity in an area of the brain that accumulates evidence in favor of choosing j . We write $\mathbf{p}_t = (p_t^j)_j$ for the profile of assessments in period t .

In each period $t = 0, 1, \dots$, the DM focuses on a single item. 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 $\alpha: A^I \times \mathbb{N} \rightarrow \Delta(\mathcal{I})$ that specifies a probability distribution over items of focus as a function of the assessments at the beginning of a period together with the current time. One can think of $\alpha(\mathbf{p}, t)$ as a vector $(\alpha^j(\mathbf{p}, t))_j$, where $\alpha^j(\mathbf{p}, t)$ is the probability with which the DM focuses on item j in period t . An attention strategy is *stationary* if it does not depend on time, in which case we simply write $\alpha(\mathbf{p})$ for $\alpha(\mathbf{p}, t)$.¹¹ 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 assessments as a function of the attention strategy is as follows. The DM begins with a vector $\mathbf{p}_0 = (p_0^1, \dots, p_0^I)$. If the assessments at the beginning of period t are \mathbf{p}_t and the DM focuses on item j in period t , then the assessment of item j follows a stochastic transition $\phi^j: A \times V \rightarrow \Delta(A)$ that depends on p_t^j and on v^j , while the assessments 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}^j is distributed according to $\phi^j(p_t^j, v^j)$ and $p_{t+1}^{j'} = p_t^{j'}$ for all $j' \neq j$. By fixing the DM's assessment of items she is not currently focusing on, we are implicitly assuming that she treats information extracted in each period as informative only about the item of current focus.

An attention strategy α naturally induces, for each \mathbf{v} , processes of items of focus $(\iota_t)_t$ and of assessments $(\mathbf{p}_t)_t$. 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 assessments \mathbf{p}_{t+1} are 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 $(\mathbf{p}_t, \iota_t)_t$.

An attention strategy α satisfies *Independence of Irrelevant Alternative j* (IIA $_j$) if, conditional on not focusing on item j , the probabilities of focusing on each item $j' \neq j$ are independent of p^j . Formally, this is the case when, for every t and \mathbf{p}, \mathbf{q} such that $p^{j'} = q^{j'}$

¹⁰In each of these cases we implicitly assume that the set of signal realizations is countable, which ensures that the set of assessments is countable as well.

¹¹Note that when an assessment consists of all past signals about an item, the time t can be retrieved from the vector of assessments \mathbf{p}_t and every strategy can be formulated as a stationary strategy.

for all $j' \neq j$ and $\alpha^{j'}(\mathbf{p}, t), \alpha^{j'}(\mathbf{q}, t) \neq 1$, we have that for every $j' \neq j$,

$$\frac{\alpha^{j'}(\mathbf{p}, t)}{1 - \alpha^{j'}(\mathbf{p}, t)} = \frac{\alpha^{j'}(\mathbf{q}, t)}{1 - \alpha^{j'}(\mathbf{q}, t)}.$$

We say that α satisfies *Independence of Irrelevant Alternatives* (IIA) if it satisfies IIA $_j$ for all items j . Note that IIA is automatically satisfied if there are only two items. It is also satisfied whenever there is an index function $G^j: A \rightarrow \mathbb{R}$ for each item j such that the DM allocates attention to the item with highest index when there is one, and randomizes uniformly in case of ties; Sections 5.2 and 4 present contexts in which such indexes emerge as a property of optimal strategies. Finally, when the index functions are positive, IIA is satisfied if the DM allocates attention to each item with probability proportional to its index. One way in which a strategy could violate IIA $_j$ would be for it to focus on the least promising of the remaining items when p^j is the lowest assessment (to check whether j is indeed the worst), and on the most promising of the remaining items when p^j is the highest assessment (to check whether j is indeed the best).

Given a baseline strategy β , an attention strategy μ is a *manipulated strategy with target item i* if

$$\mu^i(\mathbf{p}, t) \geq \beta^i(\mathbf{p}, t) \text{ for all } \mathbf{p}, t, \quad (2)$$

and, for all items $j \neq i$ and all \mathbf{p} and t such that $\mu^i(\mathbf{p}, t) < 1$,

$$\frac{\mu^j(\mathbf{p}, t)}{1 - \mu^i(\mathbf{p}, t)} = \frac{\beta^j(\mathbf{p}, t)}{1 - \beta^i(\mathbf{p}, t)}. \quad (3)$$

Together, these two conditions say that the effect of manipulation is to (weakly) increase attention toward the target item, with no effect on the relative attention devoted to other items (conditional on not inspecting the target item).

This formulation includes as a special case that manipulation induces the DM to focus on the target item for a fixed number of periods and then follows the baseline strategy thereafter. Equivalently, β could be the strategy the DM uses when all items are introduced simultaneously, while μ is the one she uses if the target item is introduced earlier than the others. Another special case is that both β and μ are stationary, with μ focusing on the target item with higher probability than β does at every \mathbf{p} (and the two strategies agreeing conditional on not focusing on the target item). For example, β could be the strategy that the DM employs in the absence of any campaign to market the target item, while μ is the strategy that results from such a campaign.

Given an attention strategy α , for each $t \geq 0$ and each item j , let $k(j, t; \alpha)$ denote the

cumulative focus of strategy α on item j before period t ; that is,

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

The following result generalizes Proposition 1.

Proposition 3. *Suppose that the baseline strategy β is stationary and satisfies IIA i , and that μ is a manipulated strategy with target item i . For each $t \geq 1$ and each \mathbf{v} , $k(i, t; \mu)$ first-order stochastically dominates $k(i, t; \beta)$, and, for each $j \neq i$, $k(j, t; \beta)$ first-order stochastically dominates $k(j, t; \mu)$.*

This proposition is an immediate corollary of Theorem 1 below, which follows from a coupling argument that we now describe.

3.1 Coupling construction

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 $(\mathbf{p}_t(\beta), \iota_t(\beta))_t$ under the baseline attention strategy β to the process $(\mathbf{p}_t(\mu), \iota_t(\mu))_t$ under the manipulated strategy μ . The construction formalizes the informal exposition in Section 2 based on decks of cards. It relies on a technique known as “coupling” (see, e.g., Lindvall, 1992): we construct the common space in such a way that the law of $(\mathbf{p}_t(\beta), \iota_t(\beta))_t$ is $P_\beta^\mathbf{v}$, while the law of $(\mathbf{p}_t(\mu), \iota_t(\mu))_t$ is $P_\mu^\mathbf{v}$. We present here a coupling construction that suffices for pure baseline and manipulated attention strategies. In Appendix A.1, we extend the construction to stochastic attention strategies.

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 assessment of item j after κ periods of focus on that item. A *learning draw* is a realization of the learning process π . For pure attention strategies, we refer to a learning draw simply as a *draw*; with stochastic attention strategies, a draw also specifies the realized items of focus (as described in Appendix A.1).

We now construct, for each pure strategy $\alpha \in \{\beta, \mu\}$, a realization of the process $(\mathbf{p}_t(\alpha), \iota_t(\alpha))_t$ as a function of the learning draw. To do so, we recursively define $(\mathbf{p}_t(\alpha), \iota_t(\alpha))_t$ as follows. Set $p_t^j(\alpha) = \pi_{k(j,t;\alpha)}^j$ for every j ; i.e., set the assessment of each item j after $k(j, t; \alpha)$ periods of focus on this item to be the $k(j, t; \alpha)$ -th value of the learning process π^j . Given α , let the focus in period t be $\iota_t(\alpha) = \alpha(\mathbf{p}_t)$. By construction, the law of the process $(\mathbf{p}_t(\alpha), \iota_t(\alpha))_t$ is $P_\alpha^\mathbf{v}$, as needed.

Our main technical insight is that, given the coupled probability space, manipulating focus alters the cumulative attention in an unambiguous direction regardless of how it affects subsequent behavior. Below, we use this insight to derive implications of attention manipulation for choice.

Theorem 1 (Attention Theorem). *In every draw, in every period, and for any stationary baseline attention strategy that satisfies IIA_i, 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 other item is at least as large under the baseline process as under the manipulated process. That is, for every $t \geq 1$ and $j \neq i$,*

$$k(i, t; \mu) \geq k(i, t; \beta)$$

and $k(j, t; \beta) \geq k(j, t; \mu)$.

The proof of this theorem is presented in Appendix A.1.

Unlike Proposition 3, this result does not describe a property of the baseline and manipulated processes on their own; it relies on our coupling construction. There are many other ways to couple the two processes in which the conclusion of this theorem does not hold, such as the coupling in which the two processes are independent; that is, for alternative couplings, there can be states of the world in which the inequalities in the theorem are violated. However, that there exists one coupling for which the result holds is sufficient to prove Proposition 3.

4 Multi-armed bandits

Theorem 1 applies to bandit problems much more generally than the special case described in Section 2.3.

In each period t , the DM chooses an item $j \in \mathcal{I}$, which, in this subsection, we refer to as an arm. Each arm j is of a fixed type $v^j \in V$. Whenever an arm j is chosen in some period t , it generates a stochastic outcome according to a discrete distribution that may depend on v^j and the sequence of previous outcomes obtained from arm j , but does not depend on t (or on the timing of previous choices of arms). For each j and t , the DM forms an index $p_t^j \in \mathbb{R}$ as a function of the sequence of outcomes obtained from arm j in periods 0 through $t - 1$. For example, p_t^j could be a Gittins index.

We define the cumulative demand $d_t^j(\mathbf{v}; \alpha)$ to be the number of periods up to t in which the DM pulls arm j under the strategy α . Let β^* be the strategy that always chooses the

arm with the highest index, with ties broken by uniform randomization among the index-maximizing arms.

Corollary 1. *Given the baseline strategy β^* , let μ be a manipulated strategy with target arm i . Relative to the baseline process, regardless of the profile \mathbf{v} of types, manipulation increases cumulative demand for the target arm i up to any period t and decreases the cumulative demand for each other arm. That is, $d_t^i(\mathbf{v}; \mu)$ first-order stochastically dominates $d_t^i(\mathbf{v}; \beta^*)$, and $d_t^j(\mathbf{v}; \beta^*)$ first-order stochastically dominates $d_t^j(\mathbf{v}; \mu)$ for all \mathbf{v} and all $j \neq i$.*

Since β^* is stationary and satisfies IIA, this result follows immediately from the Attention Theorem with assessments given by the sequence of outcomes obtained from each arm.

One application of this corollary is to a DM who is choosing an experience good in each period t that pays a stochastic flow utility which depends on the unknown quality of its brand. If the DM employs a Gittins index policy (with ties broken uniformly), any marketing intervention that attracts the DM to a target brand (say, by providing a free sample), weakly increases the cumulative demand for its target up to any point in time at the expense of the cumulative demand for each other brand. This result holds regardless of how the change in her behavior as a result of the intervention affects her subsequent choices.

In a similar vein to multi-armed bandits, the Attention Theorem also applies to models with many consumers who act sequentially, with each choosing an item based on feedback from past consumers (for example, in the form of product reviews). In this case, assessments consist of the history of feedback. Our result implies that manipulation of demand toward a product (through promotions, shelf placement, etc.) cannot backfire, and can only have a positive effect on cumulative demand for the target item.

5 Manipulation of one-shot choice

We now extend the example in Section 2.2 to the more general setting of Section 3. For simplicity, we suppose that assessments are real-valued, that is, $A \subset \mathbb{R}$.¹² The time at which the DM stops learning is governed by thresholds \underline{p}^j and \bar{p}^j for each j satisfying $\underline{p}^j < p_0^j < \bar{p}^j$. If the DM's assessment p_t^j satisfies $p_t^j \leq \underline{p}^j$ then we say that she *rejects item j* and similarly, if $p_t^j \geq \bar{p}^j$ then she *accepts item j* . The DM learns until she accepts an item or rejects all items (in which case she chooses the outside option). We thus define the stopping region F according to $F = \left(\bigcup_j F^j\right) \cup F^{oo}$, where $F^j = \{\mathbf{p} : p^j \geq \bar{p}^j\}$ and $F^{oo} = \{\mathbf{p} : p^j \leq \underline{p}^j \text{ for all } j\}$.

¹²The results in this section extend to arbitrary countable sets A of assessments together with functions $f^j : A \rightarrow \mathbb{R}$ and a stopping rule based on comparing $f(p_t^j)$ to thresholds $\underline{p}, \bar{p} \in \mathbb{R}$ in the same way as for p_t^j in this section. Thus, for example, assessments could be posterior beliefs and f^j the function mapping each belief to the corresponding expected value.

The DM makes her choice at the stopping time $\tau = \min\{t : \mathbf{p}_t \in F\}$: either $\mathbf{p}_\tau \in F^j$ for some j , in which case this j is chosen, or $\mathbf{p}_\tau \in F^{oo}$, in which case the outside option is chosen. (These cases are all mutually exclusive.) For any $j \in \mathcal{I} \cup \{oo\}$, let the *stopping time* τ^j for j be equal to τ if j is chosen and ∞ otherwise. (We allow for the possibility that learning does not stop, in which case $\tau = \tau^j = \infty$ for all j .)

An attention strategy α is *non-wasteful* if $\alpha^j(\mathbf{p}, t) = 0$ for all \mathbf{p} such that $p^j \leq \underline{p}^j$ and $p^{j'} > \underline{p}^{j'}$ for some $j' \neq j$. A non-wasteful strategy never focuses on an item that the DM has rejected.

We define the *interim demand* for item $j \in \mathcal{I} \cup \{oo\}$ under attention strategy α to be

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

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

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 demand for the outside option and the timing of its choice are unaffected by the manipulation. 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.

Proposition 4. *If a baseline attention strategy β is stationary, satisfies IIAi, and is non-wasteful, then for every \mathbf{v} and every non-wasteful manipulated strategy μ with target i ,*

$$D^i(\mathbf{v}; \mu) \geq D^i(\mathbf{v}; \beta), \tag{4}$$

$$D^j(\mathbf{v}; \mu) \leq D^j(\mathbf{v}; \beta) \text{ for every } j \neq i, \tag{5}$$

$$\text{and } D^{oo}(\mathbf{v}; \mu) = D^{oo}(\mathbf{v}; \beta). \tag{6}$$

Moreover, $\tau^i(\beta)$ first-order stochastically dominates $\tau^i(\mu)$, $\tau^j(\mu)$ first-order stochastically dominates $\tau^j(\beta)$ for every $j \neq i$, and $\tau^{oo}(\beta)$ has the same distribution as $\tau^{oo}(\mu)$.

While the non-wasteful assumption seems natural for the baseline attention strategy, in some contexts, it may be possible to manipulate attention toward items the DM would otherwise reject as being of low value. If the baseline attention strategy is non-wasteful but the manipulated strategy is not, (4) and (5) continue to hold, though (6) may not; in this case, manipulation can reduce demand for the outside option in favor of the target item. This result follows from the same reasoning as that underlying Proposition 4.

To prove Proposition 4, we start by examining the choice of the outside option.

Lemma 1 (Outside Option Lemma). *For any two non-wasteful attention strategies α and α' and any learning draw, $\tau^{oo}(\alpha) = \tau^{oo}(\alpha')$.*

The Outside Option Lemma, whose proof may be found in Appendix A.2, states that whether and when the outside option is chosen depends only on the learning draw, not on the attention strategy.¹³ For the outside option to be chosen, the assessment of each item j must reach the lower threshold \underline{p}^j before reaching the upper threshold \bar{p}^j . For any fixed draw and any particular item, the attention strategy affects only the time at which assessments about that item are attained, not their order. Thus if all items are rejected under a strategy α , they are also rejected after exactly the same number of inspections of each item under any other strategy α' .

Proof of Proposition 4. To simplify notation, let $\mathbf{p}_t = \mathbf{p}_t(\beta)$ and $\hat{\mathbf{p}}_t = \mathbf{p}_t(\mu)$. Equation (6) follows from the Outside Option Lemma by taking expectations across learning draws. For inequality (4), it suffices to show that, for any given draw, if the target item i is chosen under the baseline strategy, then i is also chosen under the manipulated strategy, and is chosen no later than under the baseline strategy.

To see this, consider any learning draw such that the target item i is chosen in period τ under the baseline strategy. Then, by the Outside Option Lemma, the outside option is not chosen under the manipulated strategy since it is not chosen under the baseline strategy. Suppose for contradiction that $j \neq i$ is chosen at some time $\hat{\tau} \leq \tau$ under the manipulated strategy. Then $\hat{p}_{\hat{\tau}}^j \geq \bar{p}^j$. By the Attention Theorem, the cumulative focus on item j under the baseline strategy is at least as large in each period as that under the manipulated strategy. Thus, there exists a period $t \leq \hat{\tau}$ such that $p_t^j = \hat{p}_{\hat{\tau}}^j \geq \bar{p}^j$, and hence the process under the baseline strategy stops with the choice of j in period t , which establishes the contradiction since stopping in F^i and F^j are mutually exclusive. Therefore, it cannot be that, under the manipulated strategy, an item $j \neq i$ is chosen at a time $\hat{\tau} \leq \tau$. By the Attention Theorem, $k(i, t; \mu) \geq k(i, t; \beta)$ 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. Again, (4) follows from taking expectations across learning draws. The proof of (5) is analogous and is relegated to Appendix A.2. \square

¹³When applied to advertising, this result relates to the question of “combativeness” (Bagwell, 2007). If advertising of a product decreases demand for its substitutes, e.g., as documented for sodas by Gasmi et al. (1992), then regulation may be justified to mitigate this negative externality. In our model, by the Outside Option Lemma, advertising is purely combative: the total demand for all items is unaffected by manipulation.

5.1 Counterexample: failure of IIA

The conclusions of Proposition 4 do not generally hold if the attention strategy does not satisfy IIA. In terms of the deck of cards analogy from Section 2.1, the problem with a failure of IIA is that, for a given draw, it may happen at some point that the number of cards viewed from the target item's deck is equal under the two processes, while the numbers for other decks differ, and consequently the beliefs also differ. In that case, the baseline process can “overtake” the manipulated process in terms of cumulative focus on the target item. IIA i ensures that this does not happen: whenever the two processes have the same cumulative focus on the target item i , they have the same cumulative focus on each other item.

To illustrate, consider an example with three items having binary values $v^j \in \{-1, 1\}$. Focusing on item $\iota_t \in \{1, 2, 3\}$ in period t generates a binary signal $x_t \in \{-1, 1\}$ equal to the true value of ι_t with probability $\lambda > 1/2$ as in Section 2. The DM is Bayesian with the prior belief that the items are independent and each has probability $1/2$ of being of high value. The assessment p_t^j is the posterior belief at the beginning of period t that $v^j = 1$. Thus, if the DM has assessment \mathbf{p}_t and focuses on item j , 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

$$\begin{aligned} p[+] &= \frac{\lambda p}{\lambda p + (1 - \lambda)(1 - p)} \\ \text{and } p[-] &= \frac{(1 - \lambda)p}{(1 - \lambda)p + \lambda(1 - p)}, \end{aligned} \tag{7}$$

as implied by Bayes' Law.

The stopping thresholds are $\underline{p} = \frac{1}{2}[-][-]$ and $\bar{p} = \frac{1}{2}[+][+]$ for each item. Let the baseline strategy β be any stationary pure attention strategy that satisfies

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

Any such strategy 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.

When following strategy β , 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 ex ante probability $1/2$.

Now consider the manipulated strategy μ obtained by directing attention to item 1 in the first period and following β thereafter. From the second period onwards, μ focuses on item 3 until p^3 reaches \underline{p} or \bar{p} , and thus μ leads to the choice of 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 $1/4$.

One can similarly construct a counterexample with an attention strategy that satisfies IIA but is non-stationary; see Appendix B.

5.2 IIA and stationarity

IIA and stationarity are satisfied by many simple strategies, such as the one that always focuses on the most promising item (i.e., the one with the highest expected value) and randomizes uniformly in case of a tie. In this section, we provide an argument in support of the IIA and stationarity assumptions based on optimization of the attention strategy for a Bayesian DM. Accordingly, we assume that the DM has a prior $p_0^j \in \Delta(V)$ for each item j , with the values of the items being independent.

We fix the stopping rule for each item and let the DM control her attention strategy α . Each period of focus on an item j generates a conditionally independent signal whose distribution depends on v^j . Given this signal, the DM updates her belief about item j according to the transition rule $\phi^j(p_t^j, v^j)$ described by Bayes' Law. Until she stops learning, the DM pays a uniformly bounded flow cost $c(p_t^{\iota_t}, \iota_t) \geq 0$ in each period t ; this flow cost may depend on the item ι_t of current focus and on the belief $p_t^{\iota_t}$ in the current period. After she terminates in some period t with an item j or with the outside option, the DM receives a one-time payoff of $E[v^j | p_t^j]$ or v^{oo} , respectively.

Termination is based on thresholds \underline{v}^j and \bar{v}^j for each item j , with $\underline{v}^j < v^{oo} < \bar{v}^j$. The DM stops and chooses item j as soon as $E[v^j | p_t^j] \geq \bar{v}^j$. She never focuses on an item j if $E[v^j | p_t^j] \leq \underline{v}^j$, and she stops and chooses the outside option as soon as $E[v^j | p_t^j] \leq \underline{v}^j$ for all items j . Given these restrictions, the DM chooses a strategy α to maximize the expected discounted value

$$U(\alpha) = E \left[\delta^\tau V^{j^*} - \sum_{t=0}^{\tau} \delta^t c(p_t^{\iota_t}, \iota_t) \right], \quad (8)$$

where $\delta \in (0, 1)$ is a discount factor, τ is the stopping time, j^* is the chosen item or the outside option, and $V^{j^*} = E[v^{j^*} | p_\tau^{j^*}]$ if j^* is an item and is equal to v^{oo} if j^* is the outside option. The expectation is with respect to the ex ante law $P_\alpha^{\text{ea}} = E_\mathbf{v} P_\alpha^\mathbf{v}$ 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 j , there exists a Gittins index function $G^j(p_t^j)$ that depends only on the assessment of item j , 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 5. *There exists a strategy that maximizes the objective (8) and satisfies IIA and stationarity.*

Note that, since this result holds for all fixed stopping thresholds, it also holds if the threshold values \underline{v}^j and \bar{v}^j are chosen optimally (within the particular family of stopping rules considered here).

The proof of Proposition 5 is relegated to Appendix A.2. The proof constructs a bandit problem that is equivalent to the optimization of (8). The main challenge in this construction is to ensure that optimal choices respect the exogenous stopping rule. For the choice of an item, the corresponding arm becomes “sticky” once its threshold is reached; the outside option has no corresponding arm, and instead is captured by all arms incurring a permanent high cost. This part of the construction is based on the Outside Option Lemma, which states that whether and when the outside option is chosen is independent of the attention strategy, and therefore appears as a constant in the objective function.

Our bandit construction applies because the stopping thresholds are restricted to be independent of the DM’s expectations 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 is not guaranteed if there are more than two items. (As noted above, with only two items, every strategy trivially satisfies IIA.)

5.3 Example: manipulation of the fastest strategy

To illustrate the quantitative impact of manipulating attention, we consider, in a simple setting along the lines of that in Section 2, the fastest attention strategy for any given stopping thresholds. This strategy takes a simple form that allows for analytical characterization of the impact of manipulation.

As in Section 2, the DM chooses among two items with unknown values $v^j \in \{-1, 1\}$ and an outside option of known value. The DM is Bayesian with, for each item j , prior belief p_0^j that $v^j = 1$; values are independent across items. The assessment of an item j is the posterior belief p_t^j . Signals are the same as in Section 2. Thus, if the DM has belief \mathbf{p}_t at the beginning of period t and focuses on item $j \in \{1, 2\}$, she updates her belief about j to $p_t^j[-]$ or $p_t^j[+]$ as described by (7).

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. Thus, the belief process for each item takes realizations in the set $\{\underline{p}, \underline{p}[+], \dots, \bar{p}[-], \bar{p}\}$.

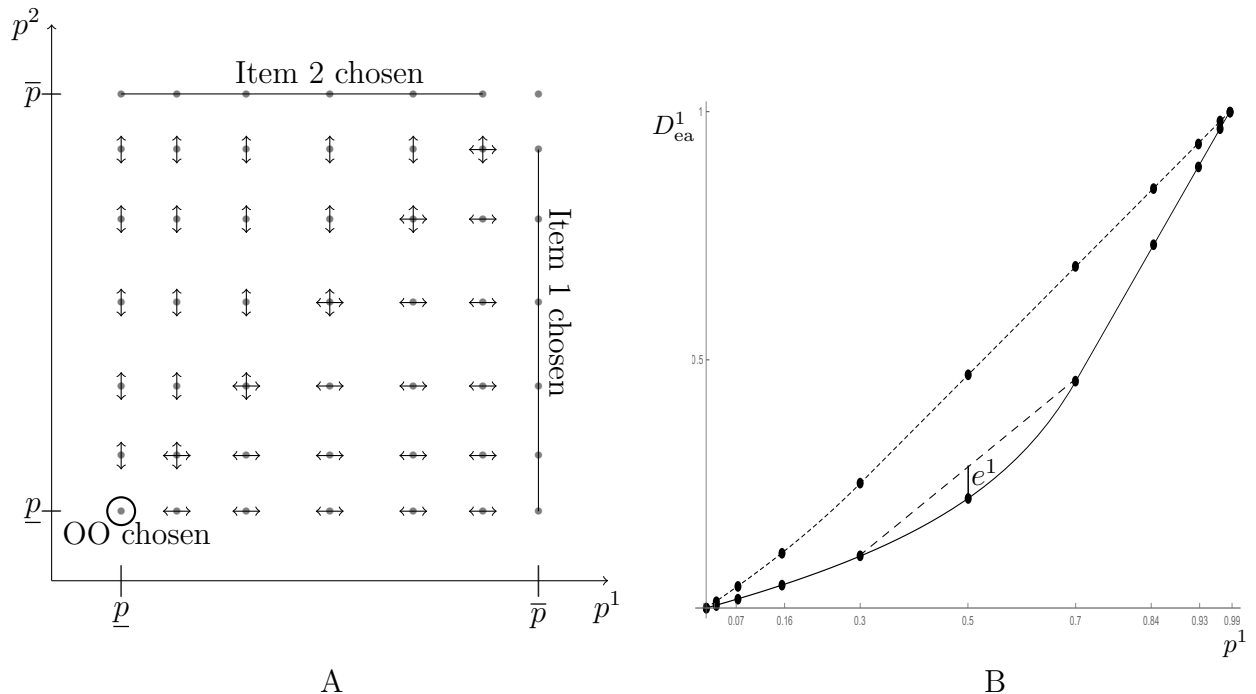


Figure 2: A: the fastest attention strategy α^* . B: the ex ante demand $D_{ea}^1(p^1, p^2)$ for item 1 as a function of p^1 . The belief p^2 is fixed at 0.7 (solid curve), or 0.3 (dashed curve). The stopping thresholds are $\bar{p} \approx .99$ and $\underline{p} \approx .01$. The depicted (single-period) 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 comparison.

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

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

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.¹⁴

¹⁴Given the stopping rule in this example, minimizing the total expected time cost is equivalent to maximizing the value $U(\alpha)$ in (8) when $c(p_t^i, \iota_t)$ is constant. This equivalence follows from the commonality of the thresholds across items together with the Outside Option Lemma.

To state this formally, define the (ex ante) 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 6. *For any strategy α , $\tau^{\text{ea}}(\alpha)$ first-order stochastically dominates $\tau^{\text{ea}}(\alpha^*)$.*

As with our main results, the proof of this result in Appendix A.2 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 6, we construct a distinct probability space in which a learning process generated by any strategy α follows the ex ante law P_α^{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 j by

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

Thus $D_{\text{ea}}^j(\mathbf{p}_0)$ is the ex ante probability that a DM with prior belief \mathbf{p}_0 chooses item j under the fastest strategy α^* . Given a prior belief \mathbf{p} , 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,

$$e^i(\mathbf{p}) = \mathbb{E} [D^i(\mathbf{v}; \mu) - D^i(\mathbf{v}, \alpha^*)] = \mathbb{E} D_{\text{ea}}^i(\tilde{p}^i, p^{-i}) - D_{\text{ea}}^i(p^i, p^{-i}), \quad (10)$$

where the first expectation is with respect to \mathbf{v} distributed according to the prior \mathbf{p} , μ is the manipulated strategy that focuses on item i in the first period and follows α^* thereafter, \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 Proposition 4, $e^i(\mathbf{p})$ is nonnegative for all \mathbf{p} .

Proposition 7. *The manipulation effect $e^i(\mathbf{p})$ is*

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

Proposition 7, which is proved in Appendix A.2, implies that manipulation has the strongest impact when it targets an item with a small a priori disadvantage relative to the other item. Equation (18) in the appendix provides an explicit expression for the size of the manipulation effect.

While the analysis above applies for exogenously fixed stopping beliefs, similar results hold in a related setting in which the stopping region is chosen optimally when the DM is patient. Ke and Villas-Boas (2019) examine a continuous-time model with exponential discounting and show that the analogue of the strategy α^* (together with a stopping region they characterize) is optimal. Despite the differences between these models, they both 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, and as the discount factor tends to 1 in the model of Ke and Villas-Boas. Therefore, they both exhibit exactly the same manipulation effect in their respective limits.

5.4 No outside option: choice by elimination and approval

While our main focus is on choice in the presence of a known outside option, many laboratory experiments employ a “forced choice” design in which there is no such option and subjects must eventually choose one of the items. Formally, we model the difference between these two cases by varying the stopping rule: in the presence of an outside option, the DM chooses an item only if she is sufficiently convinced that it is good, whereas with no outside option, she may choose an item simply because she is convinced the others are bad. In this section, we show that, with no outside option, the direction of the effect of attention manipulation generally depends on the value of the target item. With binary values and a Bayesian DM, as stopping thresholds become precise, 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 $\mathcal{I} = \{1, \dots, I\}$. Each item j has value $v^j \in \{-1, 1\}$. The DM learns about these values using an attention strategy α , with the signal structure and assessments p_t^j defined as in Section 2. (In particular, each p_t^j is a log-likelihood ratio.) Relative to our main model, the decision process 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 j , there are thresholds \underline{p}^j and \bar{p}^j . 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 j whenever \mathbf{p} lies in the set

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

Let $F_{\text{no}}(\mathbf{p}) = \bigcup_{j=1}^I F_{\text{no}}^j(\mathbf{p})$, 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. As before, we define the interim demand for item j , $\tilde{D}^j(\mathbf{v}; \alpha, \mathbf{p}) = P_{\alpha}^{\mathbf{v}}(\mathbf{p}_{\tau} \in F_{\text{no}}^j(\mathbf{p}))$, to be the probability that j is chosen conditional on the items' values.

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 Proposition 4.¹⁵ Its proof may be found in Appendix A.2.

Proposition 8. *Let \mathbf{p}_l be a sequence of thresholds such that $\underline{p}_l^i \rightarrow -\infty$ and $\bar{p}_l^i \rightarrow \infty$. Let (β_l, μ_l) be a sequence of baseline and manipulated strategy pairs, each with target item i , such that, for each l , β_l is stationary and satisfies IIA i , and both β_l and μ_l are non-wasteful with respect to \mathbf{p}_l . Then,*

$$\liminf_{l \rightarrow \infty} \left(\tilde{D}^i(\mathbf{v}; \mu_l, \mathbf{p}_l) - \tilde{D}^i(\mathbf{v}; \beta_l, \mathbf{p}_l) \right) \geq 0 \text{ if } v^i = 1, \quad (11)$$

$$\text{and } \limsup_{l \rightarrow \infty} \left(\tilde{D}^i(\mathbf{v}; \mu_l, \mathbf{p}_l) - \tilde{D}^i(\mathbf{v}; \beta_l, \mathbf{p}_l) \right) \leq 0 \text{ if } v^i = -1. \quad (12)$$

This result demonstrates the importance of the stopping rule in determining the effect of manipulation. When the stopping rule allows for an item to be chosen because of a negative signal about a different item, manipulation can, in some cases, reduce demand for the target item. For example, with two items and no outside option, suppose the thresholds are such that a single negative signal about either item j is enough for p^j to drop below \underline{p}^j . Then, in any draw that begins with such a signal for each item, the item chosen is whichever one is *not* examined in the first period. Therefore, manipulation does not increase demand for the target item across all draws, thresholds, and attention strategies.

The proof of Proposition 8 (in the appendix) uses the same coupling construction as in the proof of Proposition 4: 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 (before reaching its upper threshold) and the learning process

¹⁵Proposition 8 does not generally hold outside of the limit. For a counterexample, consider a binary choice with a baseline attention strategy that focuses on item 2 whenever it is not excluded by the non-wastefulness condition. Suppose $v^1 = -1$. Manipulation that induces the DM to focus on item 1 in the first period and return to the baseline strategy thereafter increases demand for item 1 if there is some chance that the assessment p^1 reaches \bar{p}^1 following a single inspection.

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, or (ii) the learning processes for all items reach their lower thresholds. As the lower threshold for item i approaches $-\infty$ and its upper threshold approaches ∞ , the probability of event (i) vanishes when the value of i is low and, likewise, the probability of event (ii) vanishes when its value is high. 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 rule. It turns out that the direct channel is irrelevant for the manipulation effect: the difference between the two settings is driven by the stopping rule.

To disentangle the two effects, consider an alternative model in which there is no outside option, and the stopping rule is that from our main setting. Thus learning stops with the choice of item j whenever $\mathbf{p}_\tau \in F^j = \{\mathbf{p} : p^j \geq \bar{p}^j\}$ for some j , 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 j is therefore

$$\begin{aligned} D_{\text{alt}}^j(\mathbf{v}; \alpha) &= P_\alpha^\mathbf{v}(\mathbf{p}_\tau \in F^j) + \frac{1}{I} P_\alpha^\mathbf{v}(\mathbf{p}_\tau \in F^{oo}) \\ &= D^j(\mathbf{v}; \alpha) + \frac{1}{|I|} D^{oo}(\mathbf{v}; \alpha), \end{aligned} \tag{13}$$

where $D^j(\mathbf{v}; \alpha)$ is the interim demand for $j \in \mathcal{I} \cup \{oo\}$ from our main setting with an outside option. By the Outside Option Lemma, the second summand in (13) is independent of the attention strategy. It follows that

$$D_{\text{alt}}^j(\mathbf{v}; \mu) - D_{\text{alt}}^j(\mathbf{v}; \beta) = D^j(\mathbf{v}; \mu) - D^j(\mathbf{v}; \beta),$$

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.

A similar argument applies if the threshold for stopping and accepting an item depends on the assessments of other items. For example, in the Bayesian continuous-time model analyzed by Ke and Villas-Boas (2019), the threshold for accepting an item is increasing in the belief about the other item (in the case of two items). At a prior lying just to the right of the threshold for accepting item 2, manipulating attention toward item 1 can have a negative effect on demand in draws that initially feature negative information about both items. In their solution, when the value of the outside option is low, there is a large range of beliefs

at which an item is chosen immediately when the other is eliminated. For higher values of the outside option, for many beliefs, when one item is eliminated, the DM continues to learn about the other one, as in our model.

6 Further applications

6.1 Clinical trials

One setting outside of marketing in which manipulation may be of concern is clinical trials (Gallo et al., 2006). Adaptive designs, in which the assignment of treatments depends on their success so far, are becoming increasingly popular in phase II clinical trials (Berry, 2006). In a “drop-the-loser” design, treatments that perform poorly are eliminated before the end of the trial. Similarly, those that perform well may be approved for phase III before the trial ends. In keeping with regulatory requirements, success or failure is typically judged based on a fixed level of statistical significance, as in the fixed thresholds in our model.¹⁶

We analyze the effect of manipulation in the administration of treatments in a stylized model of adaptive clinical trials. A laboratory runs a sequential trial of a finite set \mathcal{I} of I different drugs on at most T patients, where T is predetermined. In each period $t = 1, \dots, T$, a new patient arrives and is administered a drug $\iota_t \in \mathcal{I}$, the effect x_t of which is observed. The distribution of x_t depends on the unknown quality of the administered drug and on random unobservable characteristics of the patient. Starting from an initial score p_0^j for each drug j , the score $p_t^{\iota_t} \in \mathbb{R}$ for drug ι_t is updated to $p_{t+1}^{\iota_t} = \phi(p_t^{\iota_t}, x_t)$ according to a discrete transition map ϕ , while the scores of the remaining drugs $j \neq \iota_t$ are kept unchanged.

The choice of drug to administer to each patient is governed by a baseline strategy $\beta : \mathbb{R}^I \rightarrow \Delta(\mathcal{I})$. A drug is dropped from the trial and not administered again once its score falls below a fixed threshold \underline{p}^j ; that is, $\beta^j(\mathbf{p}, t) = 0$ for all \mathbf{p} such that $p^j \leq \underline{p}^j$. A drug is approved (for production or more advanced testing) once its score exceeds a fixed threshold \bar{p}^j . Several drugs can be approved. An approved drug is not tested further within the trial: $\beta^j(\mathbf{p}, t) = 0$ for all \mathbf{p} such that $p^j \geq \bar{p}^j$. The trial terminates at some time τ once every drug has been either dropped or approved, or all T patients have participated. Let $D^j(\mathbf{v}; \alpha) = P_\alpha^{\mathbf{v}}(p_\tau^j(\alpha) \geq \bar{p}^j)$ be the probability that j is approved under strategy α .

Corollary 2. *Suppose the baseline strategy β is stationary and satisfies IIA $_i$, and let μ be a manipulated strategy with target i . Manipulation weakly increases the probability that*

¹⁶Adaptive designs of clinical trials are increasingly popular both because of their statistical advantages and because of ethical concerns related to administration of drugs that, based on evidence accumulated during the trial, appear to be inferior. Termination rules are often set in accordance with simple statistical significance cutoffs imposed by regulatory bodies or by academic journals. See Berry (2006) for an overview.

the target drug is approved and weakly decreases the approval probability of the other drugs: $D^i(\mathbf{v}; \mu) \geq D^i(\mathbf{v}; \beta)$ and $D^j(\mathbf{v}; \beta) \geq D^j(\mathbf{v}; \mu)$ for all $j \neq i$.

One example of a stationary strategy that satisfies IIA and may be ethically appealing is the one that always administers the drug having the highest current score (among those that have not already been approved). A simple form of manipulation consists of administering the target drug to the first m patients and then continuing with the baseline strategy β .

The result is a corollary of the Attention Theorem. Since, by any period t , the target drug receives at least as much cumulative attention—i.e., is administered at least as many times—under the manipulated strategy as under the baseline strategy, it is approved under μ in any draw in which it is approved under β . Similarly, for any other drug, in any draw for which it is approved under the manipulated strategy, it is also approved under the baseline strategy.

6.2 A neuroeconomic model of choice

The fixed thresholds in our one-shot choice model differ from those in the optimal strategy in Ke and Villas-Boas (2019), which features thresholds for acceptance that are increasing in the belief about the other item. A related (but distinct) implementation of varying standards for acceptance appears in neuroscience in the form of so-called “inhibitory links.” Signals favoring the choice of each option are accumulated in distinct brain areas called “accumulators”; a given option is chosen once the firing rate of its associated accumulator exceeds a fixed threshold (Shadlen and Newsome, 2001). When a stimulus causes the accumulator associated with one item to increase, the accumulator for the other item is inhibited via the inhibitory links (see, e.g., Sterling and Laughlin, 2015).

We model the decision process from Shadlen and Newsome (2001) as follows. The DM chooses one of two items $j = 1, 2$ with values $v^j \in \{-1, 1\}$. In each period $t = 0, 1, \dots$, she inspects an item, ι_t . If the inspected item is good (i.e., $v^{\iota_t} = 1$), then she receives a stimulus with probability $\lambda > 1/2$; if it is bad, she receives a stimulus with probability $1 - \lambda$. Receipt of a stimulus for an item increases the accumulator for that item and inhibits the accumulator for the other item. More precisely, the value of the accumulator for item $j = 1, 2$ at the beginning of period t is

$$p_t^j = \sum_{s < t: \iota_s = j} x_s - \chi \sum_{s < t: \iota_s = j'} x_s,$$

where $j' \neq j$, $x_s \in \{0, 1\}$ is an indicator of whether the DM received a stimulus in period s , and $\chi > 0$ is a parameter governing the strength of inhibition.

Inspection is determined by an arbitrary baseline attention strategy β , which is a function of the pair $\mathbf{p}_t = (p_t^1, p_t^2)$. There is a fixed threshold \bar{p}^j for each j such that the process terminates with the choice of item j as soon as $p_t^j \geq \bar{p}^j$. (If neither accumulator reaches the threshold, the DM never makes a choice.) Let $F^j = \{\mathbf{p} : p^j \geq \bar{p}^j\}$, $\tau = \min_t \{\mathbf{p}_t \in F^1 \cup F^2\}$, and $D^j(\mathbf{v}; \alpha) = P_\alpha^\mathbf{v}(\mathbf{p}_\tau \in F^j)$.

Corollary 3. *Suppose the baseline strategy β is stationary, and let μ be a manipulated strategy. Manipulation weakly increases the probability that the target item is chosen and weakly decreases the probability that the other item is chosen; that is, $D^i(\mathbf{v}; \mu) \geq D^i(\mathbf{v}; \beta)$ for the target item i , and $D^j(\mathbf{v}; \beta) \geq D^j(\mathbf{v}; \mu)$ for $j \neq i$.*

As stated, this model is not a special case of our main model because the assessment p_t^j of item j can change when the other item is inspected. However, a simple reformulation enables us to apply the Attention Theorem (of which it is a corollary). Let $q_t^j = \sum_{s < t: u_s = j} x_s$, $\tilde{\beta}(q_t^1, q_t^2) = \beta(q_t^1 - \chi q_t^2, q_t^2 - \chi q_t^1)$, $\tilde{\mu}(q_t^1, q_t^2) = \mu(q_t^1 - \chi q_t^2, q_t^2 - \chi q_t^1)$, and the stopping region be $\tilde{F}^j = \{\mathbf{q} : q_t^j - \chi q_t^{-j} \geq \bar{p}^j\}$ for each j . With this reformulation, each assessment q_t^j depends only on the stimuli for item j , and the Attention Theorem applies. It follows that the cumulative attention $k(i, t, \mu)$ devoted to the target item i under the manipulated strategy, μ , is always at least as large as that under the baseline strategy, β , and the opposite is true for the other item, j . In our coupling construction, in every draw, $q_t^i(\mu) \geq q_t^i(\beta)$ and $q_t^j(\beta) \geq q_t^j(\mu)$ for each t because, in this model, the assessment of an item is nondecreasing in the number of periods of focus on that item. Therefore,

$$p_t^i(\mu) = q_t^i(\mu) - \chi q_t^j(\mu) \geq q_t^i(\beta) - \chi q_t^j(\beta) = p_t^i(\beta)$$

and

$$p_t^j(\beta) = q_t^j(\beta) - \chi q_t^i(\beta) \geq q_t^j(\mu) - \chi q_t^i(\mu) = p_t^j(\mu)$$

for each t and each draw. The first claim in the corollary now follows: for each draw in which strategy β terminates with the choice of i , so does strategy μ . The argument for the second claim is similar.

7 Summary and discussion

We show how manipulating attention toward an item can boost a DM's demand for it. Our argument has two key parts. The first is to identify the effect of manipulation on the cumulative attention paid to an item. While manipulation has the direct effect of increasing attention to the target item, it also affects attention indirectly through the DM's subsequent

choices of which item to focus on. We show that, under general conditions on the DM’s strategy, this indirect effect never overcomes the direct effect, and therefore manipulation can only increase the attention paid to its target. The second part of the argument connects cumulative attention to choice. With repeated choice, such as multi-armed bandits, this connection is immediate. In one-shot choice, we show that, in the presence of an outside option, increased attention implies increased demand when the DM chooses according to a simple satisficing rule.

Our model offers a novel combination of satisficing and gradual learning about the values of available choices. This combination necessitates some differences compared to standard models of satisficing behavior. Most notably, with gradual learning, one must take a stand on whether and when the DM rejects an item and stops learning about it. We assume that rejection occurs when the assessment falls below a fixed threshold, just as acceptance is associated with exceeding a fixed threshold.

The positive effect of attention manipulation can break down in one-shot choice if the DM sometimes chooses by elimination, i.e., as a result of rejecting other items without first verifying the value of the chosen item. Choice by elimination arises naturally in the absence of an outside option, or if the DM optimizes her stopping rule instead of satisficing. For optimal stopping by a Bayesian DM, while we know that our main result does not generally hold as stated—i.e., for every prior and realized profile of values, and a large class of attention strategies—identifying the parameters of the model for which it does is an open question.

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.
- Atalay, A. S., H. O. Bodur, and D. Rasolofoarison (2012). Shining in the center: Central gaze cascade effect on product choice. *Journal of Consumer Research* 39(4), 848–866.
- 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.
- Becker, G. S. and K. M. Murphy (1993). A simple theory of advertising as a good or bad. *Quarterly Journal of Economics* 108(4), 941–964.
- Berry, D. A. (2006). Bayesian clinical trials. *Nature Reviews Drug Discovery* 5(1), 27–36.

- 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.
- Braithwaite, D. (1928). The economic effects of advertisement. *The Economic Journal* 38(149), 16–37.
- 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 (2019). Optimal dynamic allocation of attention. *American Economic Review* 108(8), 2993–3029.
- Chernoff, H. (1959). Sequential design of experiments. *Annals of Mathematical Statistics* 30(3), 755–770.
- 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.
- Eliaz, K. and R. Spiegel (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.
- Gallo, P., C. Chuang-Stein, V. Dragalin, B. Gaydos, M. Krams, and J. Pinheiro (2006). Adaptive designs in clinical drug development: an executive summary of the PhRMA working group. *Journal of Biopharmaceutical Statistics* 16(3), 275–283.
- Gasmi, F., J. J. Laffont, and Q. Vuong (1992). Econometric analysis of collusive behavior in a soft-drink market. *Journal of Economics & Management Strategy* 1(2), 277–311.
- 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.
- Kaldor, N. (1950). The economic aspects of advertising. *Review of Economic Studies* 18(1), 1–27.
- 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.
- Nelson, P. (1974). Advertising as information. *Journal of Political Economy* 82(4), 729–754.
- Nikandrova, A. and R. Pancs (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.

- Shadlen, M. N. and W. T. Newsome (2001). Neural basis of a perceptual decision in the parietal cortex (area lip) of the rhesus monkey. *Journal of Neurophysiology* 86(4), 1916–1936.
- Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics* 69(1), 99–118.
- Spelke, E., W. Hirst, and U. Neisser (1976). Skills of divided attention. *Cognition* 4(3), 215–230.
- Sterling, P. and S. Laughlin (2015). *Principles of Neural Design*. MIT Press.
- Stigler, G. J. (1961). The economics of information. *Journal of Political Economy* 69(3), 213–225.
- Stigler, G. J. and G. S. Becker (1977). De gustibus non est disputandum. *American Economic Review* 67(2), 76–90.
- Telser, L. G. (1964). Advertising and competition. *Journal of Political Economy* 72(6), 537–562.
- 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.

Appendix

A Proofs

A.1 Proofs for Section 3

We consider a baseline strategy β and a manipulated strategy μ in favour of item i .

The main difficulty compared to the construction of the probability space in Section 3.1 is that the strategies β and μ are stochastic instead of deterministic. This requires us to amend our coupling construction in order to keep track of the realized items of focus. We thus use the learning process $\pi = (\pi^j)_{j=1}^I$ as constructed in Section 3.1, and introduce an *attention process* consisting of random variables $(a_{\mathbf{p},t})_{\mathbf{p},t}$, $(b_{\mathbf{p}^{-i},t})_{\mathbf{p}^{-i},t}$, and $(c_{\mathbf{p},t})_{\mathbf{p},t}$, all of which are independent and are independent of the learning process. The random variable $a_{\mathbf{p},t}$ takes values in $\{0, 1\}$ and $\Pr(a_{\mathbf{p},t} = 1) = \mu^i(\mathbf{p}, t)$. 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 $\beta^i(p^i, \mathbf{p}^{-i}) \neq 1$, then the probability that $b_{\mathbf{p}^{-i},t}$ takes the value $j \neq i$ is $\beta^j(p^i, \mathbf{p}^{-i}) / (1 - \beta^i(p^i, \mathbf{p}^{-i}))$, where we note that the particular value of p^i in the formula is irrelevant since β satisfies IIAi.¹⁷ The random variable $c_{\mathbf{p},t}$ takes values in $\{0, 1\}$ and $\Pr(c_{\mathbf{p},t} = 1) = \frac{\beta^i(\mathbf{p})}{\mu^i(\mathbf{p}, t)}$ whenever $\mu^i(\mathbf{p}, t) > 0$ (and is otherwise arbitrary). A realization of the attention process is called an *attention draw*. We refer to a pair of learning and attention draws simply as a *draw*.

For a strategy $\alpha = \beta$ or μ , we recursively construct the processes $(\mathbf{p}_t(\alpha), \iota_t(\alpha))_t$ as functions of the draw. Recall that $k(j, t; \alpha)$ denotes the cumulative focus on item j up to period $t - 1$ given by the realizations $(\iota_t(\alpha))_1, \dots, \iota_t(\alpha)_{t-1}$. We let $p_t^j(\alpha) = \pi_{k(j,t;\alpha)}^j$ be the belief about item j after $k(j, t; \alpha)$ inspections. We proceed to construct the focus $\iota_t(\alpha)$ whenever $\mathbf{p}_t \notin F$. Given a sequence of assessments and focus items $(\mathbf{p}_s(\alpha), \iota_s(\alpha))_{s < t}$ and any vector \mathbf{p}^{-i} of assessments about items $j \neq i$, we let $\nu(\mathbf{p}^{-i}, t; \alpha)$ be the total number of periods before time t with assessments \mathbf{p}^{-i} about items $j \neq i$ in which the DM does not focus on the target item i ; that is,

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

For the strategy μ , let $\iota_t(\mu) = i$ if $a_{\mathbf{p}_t(\mu),t} = 1$ and $\iota_t(\mu) = b_{\mathbf{p}_t^{-i}(\mu), \nu(\mathbf{p}_t^{-i}(\mu), t; \mu)}$ otherwise. For the strategy β , let $\iota_t(\beta) = i$ if $a_{\mathbf{p}_t(\beta),t} c_{\mathbf{p}_t(\beta),t} = 1$ and $\iota_t(\beta) = b_{\mathbf{p}_t^{-i}(\beta), \nu(\mathbf{p}_t^{-i}(\beta), t; \beta)}$ otherwise.

By construction, conditional on $(\mathbf{p}_0(\alpha), \iota_0(\alpha), \dots, \mathbf{p}_{t-1}(\alpha), \iota_{t-1}(\alpha), \mathbf{p}_t(\alpha))$, $\iota_t(\alpha)$ is distributed according to $\alpha(\mathbf{p}_t(\alpha), t)$. Conditional on $(\mathbf{p}_0(\alpha), \iota_0(\alpha), \dots, \mathbf{p}_t(\alpha), \iota_t(\alpha))$, $p_{t+1}^j(\alpha) =$

¹⁷The specification of $b_{\mathbf{p}^{-i},t}$ when no such p^i exists is immaterial for our purposes.

$p_t^j(\alpha)$ for all $j \neq \iota_t(\alpha)$, and $p_{t+1}^{\iota_t(\alpha)}(\alpha)$ is distributed according to $\phi^{\iota_t(\alpha)}\left(p_t^{\iota_t(\alpha)}(\alpha), v^{\iota_t(\alpha)}\right)$. Therefore, the law of the process $(\mathbf{p}_t(\alpha), \iota_t(\alpha))_t$ is P_α^ν for both $\alpha = \beta, \mu$, and this is indeed a valid coupling.

Let $k(-i, t; \alpha) = t - k(i, t; \alpha)$ denote the number of periods of focus on items other than the target i . To simplify notation, let $k(j, t) = k(j, t; \beta)$ and $\hat{k}(j, t) = k(j, t; \mu)$ for $j \in \mathcal{I} \cup \{-i\}$. For any $n \geq 0$, let $\mathbf{t}(n)$ be the n -th period such that $\iota_t \neq i$ under the baseline strategy; that is, $\mathbf{t}(n) = \min\{t : k(-i, t) \geq n\}$ (where potentially $\mathbf{t}(n) = \infty$). Similarly, let $\hat{\mathbf{t}}(n) = \min\{t : \hat{k}(-i, t) \geq n\}$. Let $(\mathbf{p}_t, \iota_t)_t = (\mathbf{p}_t(\beta), \iota_t(\beta))_t$ and $(\hat{\mathbf{p}}_t, \hat{\iota}_t)_t = (\mathbf{p}_t(\mu), \iota_t(\mu))_t$.

Lemma 2 (Coupling Lemma). *The baseline and manipulated processes coincide when restricted to periods of focus on items other than the target, i . That is, for every n such that $\mathbf{t}(n), \hat{\mathbf{t}}(n) < \infty$,*

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

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

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

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

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

Proof of the Attention Theorem. We prove by induction that $\hat{k}(i, t) \geq k(i, t)$ and $\hat{k}(-i, t) \leq k(-i, t)$ for all periods t . This statement obviously holds for period $t+1$ if $\hat{k}(i, t) > k(i, t)$.

Suppose that $k(i, t) = \hat{k}(i, t)$ in period t . By the Coupling Lemma,

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

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 $\mathbf{p}_t = \hat{\mathbf{p}}_t$. By the coupling construction, for each \mathbf{p} and t , the baseline process focuses on the target item i at time t

only if the manipulated process focuses on i at t . Thus, again, the statement holds in period $t + 1$.

Additionally, 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

A.2 Proofs for Section 5

The next two proofs use the probability space arising from the coupling construction in Section A.1.

Proof of the 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 j , let $\kappa^j = \min\{\kappa : \pi_\kappa^j \leq \underline{p}^j\}$ be the number of steps needed for the learning process π^j to reach the threshold \underline{p}^j . Since α stops with the choice of the outside option once all beliefs reach their respective lower thresholds, $\tau^{oo}(\alpha) = \sum_j \kappa^j$, which is independent of α . In particular, the distribution of $\tau^{oo}(\alpha)$, as well as whether the outside option is chosen (i.e., $\tau^{oo}(\alpha) < \infty$), are independent of α , and so is the probability of this event. \square

Proof of Proposition 4, inequality (5). 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 Theorem, there exists a period $t \leq \tau$ such $\hat{p}_t^i = p_\tau^i \geq \bar{p}^i$ 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 or the outside option is chosen at a time $\tau \leq \hat{\tau}$. By the Attention Theorem, $k(-i, t) \geq \hat{k}(-i, t)$ 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 $\mathbf{p}_\tau^{-i} = \hat{\mathbf{p}}_\tau^{-i}$, and the baseline process stops in period τ with the choice of item j , as needed. \square

Proof of Proposition 5. We construct a bandit problem and apply a result from Tsitsiklis (1994) showing the existence of an optimal (stationary) Gittins index strategy; with an appropriate choice of tie-breaking rule, such strategies satisfy IIA. In the model of Tsitsiklis (1994) (described in the language of our model), when the DM selects an item $j = \iota_t$ with belief p_t^j in period t , she receives a stochastic flow reward $R(p_t^j, j)$ and the selected item remains active for $T^j(p_t^j)$ periods. In period $t + T^j(p_t^j)$, the belief about item j transitions to a new belief $p_{t+T^j(p_t^j)}^j = \psi^j(p_t^j)$ according to some stochastic function ψ^j , while the beliefs

about all other items remain unchanged (that is, $p_{t+T^j(p_t^j)}^{j'} = p_t^{j'}$ for all $j' \neq j$). Then, the DM selects a new item $\iota_{t+T^j(p_t^j)}$ in period $t + T^j(p_t^j)$, and so on. . . Let t_k denote the random sequence of periods in which the DM selects items: $t_0 = 0$ and $t_{k+1} = t_k + T(p_{t_k}^{\iota_{t_k}})$ for each $k \geq 0$. The DM's problem is to maximize the expected discounted total reward:

$$\Gamma(\alpha) = \mathbb{E} \sum_{k=0}^{\infty} \delta^{t_k} R(p_{t_k}^{\iota_{t_k}}, \iota_{t_k}). \quad (14)$$

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

$$R(p, j) = \begin{cases} -c(p, j) & \text{if } \underline{v}^j < \mathbb{E}[v^j | p^j = \psi^j(p)] < \bar{v}^j, \\ -c(p, j) + \mathbb{E}[v^j | p^j = \psi^j(p)] & \text{if } \mathbb{E}[v^j | p^j = \psi^j(p)] \geq \bar{v}^j, \\ -\bar{c} & \text{otherwise,} \end{cases}$$

where \bar{c} is a uniform upper bound on $c(p_{t_k}^{\iota_{t_k}}, \iota_{t_k})$. If item j is selected in period t_k , we set the activity length $T^j(p_{t_k}^j)$ to be 1 if $\mathbb{E}[v^j | p_{t_{k+1}}^j] < \bar{v}^j$ and ∞ otherwise.¹⁸

This problem differs from the original problem (8) in that the flow payoff from focusing on item j in period t depends only on p_t^j and is independent of $p_t^{j'}$ for $j' \neq j$. The flow payoffs in the original problem are interdependent because whether the DM stops upon her expectation of v^j reaching \underline{v}^j depends on whether expectations for other items have reached their respective thresholds. This difference allows us to apply Theorem 2.2 from Tsitsiklis (1994) (see also Weber, 1992): an optimal solution to the bandit problem (14) is a Gittins index strategy where the index for an item depends only on the current belief about the item. 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 (14) to the original problem (8). Note that any optimal strategy in the bandit problem (i) chooses an item j forever once $\mathbb{E}[v^j | p_t^j] \geq \bar{v}^j$, and (ii) always chooses an item j such that $\mathbb{E}[v^j | p_t^j] > \underline{v}^j$ provided such an item exists.

Note that $\Gamma(\alpha) = U(\alpha) - \bar{c} \mathbb{E} \sum_{t=1}^{\infty} 1_{t \geq \tau^{\circ\circ}(\alpha)}$, 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). Therefore, there exists a constant K_δ such that

$$\Gamma(\alpha) = U(\alpha) + K_\delta. \quad (15)$$

¹⁸This correlation among the active-period length, reward, and assessment transition complies with Tsitsiklis's setting, as he allows for arbitrary correlations among these random variables.

Thus, a non-wasteful strategy α solves the auxiliary bandit problem (14) if and only if it solves the original problem (8), as needed. \square

Proof of Proposition 6. This proof makes use of coupling: we construct a common probability space on which we can compare the processes of beliefs $\mathbf{p}_t(\alpha)$ across attention strategies α . However, the idea of the proof and the construction of the probability space significantly differ from that in the proof of Proposition 4.

In this construction, an outcome in the probability space specifies, for every possible belief, the sequence of belief transitions on any of the items starting at this belief. Let $\Pi = \{\underline{p}[+], \underline{p}[+][+], \dots, \bar{p}[-][-], \bar{p}[-]\}$ denote the set of beliefs that the DM may attain for either item. For each $\pi \in \Pi$ and $\kappa = 0, 1, \dots$, let $(\ell(\pi, \kappa))_{\pi, \kappa}$ be an i.i.d. family of random variables with values in $\{\pi[+], \pi[-]\}$ such that the respective probabilities that $\ell(\pi, \kappa)$ takes each value are $\frac{\pi - \pi[-]}{\pi[+] - \pi[-]}$ and $\frac{\pi[+] - \pi}{\pi[+] - \pi[-]}$. An element of the probability space, which we refer to as an *updating draw*, is a realization of the family of random variables $(\ell(\pi, \kappa))_{\pi, \kappa}$. We interpret $\ell(\pi, \kappa)$ as the updated belief of a DM who learns for one period about an item $j = 1, 2$ starting at the belief $p^j = \pi$, where κ indicates the total number of times the DM has focused on an item with associated belief π earlier on.

We now construct, for each fixed updating draw and any attention strategy α , the process of beliefs $\mathbf{p}_t(\alpha)$ and verify that they follow the law P_α^{ea} . Let $\mathbf{k}_t(\alpha) = (k_t^\pi(\alpha))_{\pi \in \Pi}$ be an auxiliary counter that takes values in $\mathbb{N}^{|\Pi|}$. We interpret $k_t^\pi(\alpha)$ as the number of times the belief π has been updated up to period t under strategy α . Define the joint process of the beliefs $\mathbf{p}_t(\alpha)$ and the counter $\mathbf{k}_t(\alpha)$ as follows. Let $k_0^\pi(\alpha) = 0$ for all π and $\mathbf{p}_0(\alpha) = \mathbf{p}_0$. In each period t , the focus of attention ι_t is chosen according to the attention strategy $\alpha(\mathbf{p}_t, t)$. Recursively define

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

and

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

By construction, for each strategy α , the beliefs $\mathbf{p}_t(\alpha)$ follow the law P_α^{ea} .

We make use of an auxiliary process of the beliefs that would result from learning about a single item in all periods starting from a prior belief π . This process consists of beliefs $a_t(\pi)$ and a counter $h_t^{\tilde{\pi}}$ for each $\pi, \tilde{\pi} \in \Pi$. Formally, let $a_0(\pi) = \pi$ and $h_0^{\tilde{\pi}} = 0$ for all $\pi, \tilde{\pi} \in \Pi$, and

recursively define $a_{t+1}(\pi) = \ell\left(a_t(\pi), h_t^{a_t(\pi)}\right)$, and, for each $\tilde{\pi} \in \Pi$,

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

That is, $a_t(\pi)$ is the belief in period t and $h_t^{\tilde{\pi}}$ counts how many times belief $\tilde{\pi}$ has been updated up to period t in this auxiliary process.

It is convenient to reformulate the stopping rule as follows. For each strategy α , let $M_t(\alpha) = \max_{j=1,2} p_t^j(\alpha)$, $\overline{M}_t(\alpha) = \max_{s=0,\dots,t} M_s(\alpha)$, and $\underline{M}_t(\alpha) = \min_{s=0,\dots,t} M_s(\alpha)$. The strategy α stops by period t , i.e., $\tau^{\text{ea}}(\alpha) \leq t$, if

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

We now 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 . First, consider $t = 1$. The statement is trivial if $p_0^1 = p_0^2$; accordingly, assume $p_0^1 \neq p_0^2$. Again, the statement is trivial if α focuses on the item with the higher belief in period 0; accordingly, assume that α focuses on the item with the lower belief in period 0. 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 $M_1(\alpha) = \max\{p_0^1, p_0^2\}$. 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 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^* . Therefore, let $\tilde{\alpha}$ focus on item 1 in the first period and follow α^* thereafter.

To compare $\tilde{\alpha}$ and α^* , 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 $s = 1, 2, \dots, t$,

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

and

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

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 $\tilde{\alpha}$ 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,

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

and the induction step holds.

The second set is the complementary set to the first; that is, it is the set of updating draws for which there exists a period $s \in \{1, 2, \dots, t-1\}$ in which $a_s(p_0^2) = p_0^1$. 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^* under the strategy α^* . We also have

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

since the belief p_0^1 is updated once in period 0 and in each period $s = 1, \dots, s^*$, $\langle \mathbf{p}_s(\tilde{\alpha}) \rangle = \{a_1(p_0^1), a_{s-1}(p_0^2)\}$ and $a_{s-1}(p_0^2) \geq a_1(p_0^1)$. Thus, the strategy $\tilde{\alpha}$ 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

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

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

Proof of Proposition 7. Without loss of generality, we prove the statements for $i = 1$. We claim that

$$D_{\text{ea}}^1(\mathbf{p}) = \begin{cases} \frac{\bar{p}-p^1}{2(\bar{p}-p^2)} \left(1 - \frac{(\bar{p}-p^2)^2}{(\bar{p}-p)^2}\right) + \frac{p^1-p^2}{\bar{p}-p^2} & \text{if } p^1 \geq p^2, \\ \frac{\bar{p}-p^2}{2(\bar{p}-p^1)} \left(1 - \frac{(\bar{p}-p^1)^2}{(\bar{p}-p)^2}\right) & \text{if } p^1 \leq p^2. \end{cases} \quad (16)$$

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, for $p_0^1 = p_0^2$, conditional on not choosing

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

$$D_{\text{ea}}^1(p, p) = \frac{1}{2} \left(1 - \left(\frac{\bar{p} - p}{\bar{p} - \underline{p}} \right)^2 \right). \quad (17)$$

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 $p^1 > p^2$,

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

which, together with (17), yields the expression in (16) for this case.

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

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

which, together with (17), yields the expression in (16) for this case.

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

$$e^1(\mathbf{p}; \bar{p}) = \frac{(1 - 2\lambda)^2 (1 - p^1)^2 (p^1)^2 (\bar{p} - p^2)}{(\bar{p} - p^1)(\bar{p}p^1 - \bar{p} + \lambda p^1 - \bar{p}\lambda(2p^1 - 1))(p^1 - \bar{p}p^1 - \lambda p^1 + \bar{p}\lambda(2p^1 - 1))}. \quad (18)$$

Statement 4 of the proposition follows immediately from this expression. Statements 2 and 3 follow from the monotonicity of the right-hand side of (18) with respect to p^2 and p^1 , respectively. Monotonicity with respect to p^2 is straightforward. For p^1 , one may verify that $\frac{\partial^2}{\partial p^1 \partial \bar{p}} \log e^1(\mathbf{p}; \bar{p}) < 0$ for $0 < p^1 < \bar{p} < 1$. Thus, it suffices to show that $\log e^1(\mathbf{p}; \bar{p})$ for $\bar{p} = 1$ increases in p^1 . At $\bar{p} = 1$, (18) 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 8. We use the same coupling construction as in the proof of Proposition 4, again denoting by \mathbf{p}_t and $\hat{\mathbf{p}}_t$ the assessment sequences generated by strategies β_l and μ_l , respectively. Similarly, we denote variables generated with strategy β_l by τ , $k(j, t)$, etc., and variables generated by μ_l by $\hat{\tau}$, $\hat{k}(j, t)$, etc.

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

The second set, R_2 , consists of all \mathbf{r} such that $r^i = \bar{p}_l^i$ for the target item i . Consider $\mathbf{r} \in R_2$ and any draw $\pi \in \Pi(\mathbf{r})$. We claim that if β_l leads to the choice of the target item i in draw π , then so does μ_l . The claim is immediate if π^j reaches \underline{p}_l^j for each $j \neq i$; accordingly, suppose that π^j reaches \bar{p}_l^j for some $j \neq i$. If, under β_l , the DM chooses i in period τ , then $p_\tau^i \geq \bar{p}_l^i$ (since it is not the case that the belief for each other item reaches its lower threshold). Assume for contradiction that, under μ_l , the DM chooses $j \neq i$ in some period $\hat{\tau} \leq \tau$, and hence $\hat{p}_{\hat{\tau}}^j \geq \bar{p}_l^j$. By the Attention Theorem, $k(j, \hat{\tau}) \geq \hat{k}(j, \hat{\tau})$ and thus $\hat{p}_{\hat{\tau}}^j \geq \bar{p}_l^j$ for some $\tau' \leq \hat{\tau} \leq \tau$. Therefore, β_l leads to the choice of $j \neq i$ in period $\tau' \leq \tau$, contradicting that β_l leads to the choice of i in period τ . Again by the Attention Theorem, $k(i, \tau) \leq \hat{k}(i, \tau)$, hence $\hat{p}_\tau^i \geq \bar{p}_l^i$, and μ_l leads to the choice of i in period τ or earlier.

The third set, R_3 , is the singleton $\{(\underline{p}_l^1, \dots, \underline{p}_l^L)\}$. Consider $\mathbf{r} = (\underline{p}_l^1, \dots, \underline{p}_l^L)$ and a draw $\pi \in \Pi(\mathbf{r})$. We claim that if μ_l leads to the choice of the target item i in draw π , then so does β_l . If, under μ_l , the DM chooses i in period $\hat{\tau}$, then $\hat{p}_{\hat{\tau}}^j \leq \underline{p}_l^j$ for all $j \neq i$. Assume for contradiction that, under β_l , the DM chooses $j^* \neq i$ in some period $\tau \leq \hat{\tau}$, and hence $p_\tau^{j^*} \leq \underline{p}_l^{j^*}$ for all $j \neq j^*$. By the Attention Theorem, $k(i, \tau) \leq \hat{k}(i, \tau)$ and thus $\hat{p}_\tau^i \leq \underline{p}_l^i$ for some $\hat{\tau}' \leq \tau \leq \hat{\tau}$. Non-wastefulness of μ_l together with the Coupling Lemma imply that $\mathbf{p}_\tau = \hat{\mathbf{p}}_\tau$, and thus μ_l leads to the choice of j^* in period $\tau \leq \hat{\tau}$, contradicting that μ_l leads to the choice of i in period $\hat{\tau}$. Again by the Attention Theorem, $k(i, \hat{\tau}) \leq \hat{k}(i, \hat{\tau})$, which, together with the Coupling Lemma, implies that $\hat{p}_{\hat{\tau}}^j \leq \underline{p}_l^j$ for all $j \neq i$ and hence β_l leads to the choice of i in period $\hat{\tau}$ or earlier.

Observe that as $l \rightarrow \infty$, for any \mathbf{v} such that $v^i = 1$,

$$\Pr \left(\Pi \left(\underline{p}_l^1, \dots, \underline{p}_l^I \right) \mid \mathbf{v}; \mathbf{p}_l \right) \leq \Pr \left(r^i = \underline{p}_l^i \mid v^i = 1 \right) \rightarrow 0,$$

and for any \mathbf{v} such that $v^i = 0$,

$$\Pr \left(\bigcup_{\mathbf{r} \in R_2} \Pi_l(\mathbf{r}) \mid \mathbf{v}; \mathbf{p}_l \right) \leq \Pr \left(r^i = \bar{p}_l^i \mid v^i = 0 \right) \rightarrow 0.$$

Therefore, for \mathbf{v} such that $v^i = 1$, $\liminf_{l \rightarrow \infty} \left(\tilde{D}^i(\mathbf{v}; \mu_l, \mathbf{p}_l) - \tilde{D}^i(\mathbf{v}; \beta_l, \mathbf{p}_l) \right)$ is nonnegative since it has the same sign as

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

which is nonnegative for all l since for any draw in $\Pi_l(\mathbf{r})$ with $\mathbf{r} \in R_2$, $\mathbf{p}_{\tau} \in F_{\text{no}}^i(\mathbf{p}_l)$ implies $\hat{\mathbf{p}}_{\hat{\tau}} \in F_{\text{no}}^i(\mathbf{p}_l)$.

Similarly, for \mathbf{v} such that $v^i = 0$, the sign of $\limsup_{l \rightarrow \infty} \left(\tilde{D}^i(\mathbf{v}; \mu_l, \mathbf{p}_l) - \tilde{D}^i(\mathbf{v}; \beta_l, \mathbf{p}_l) \right)$ is the same as that of

$$\Pr \left(\hat{\mathbf{p}}_{\hat{\tau}} \in F_{\text{no}}^i(\mathbf{p}_l) \mid \Pi \left(\underline{p}_l^1, \dots, \underline{p}_l^I \right), \mathbf{v} \right) - \Pr \left(\mathbf{p}_{\tau} \in F_{\text{no}}^i(\mathbf{p}_l) \mid \Pi \left(\underline{p}_l^1, \dots, \underline{p}_l^I \right), \mathbf{v} \right).$$

which is nonpositive since for any draw in $\Pi \left(\underline{p}_l^1, \dots, \underline{p}_l^I \right)$, $\hat{\mathbf{p}}_{\hat{\tau}} \in F_{\text{no}}^i(\mathbf{p}_l)$ implies $\mathbf{p}_{\tau} \in F_{\text{no}}^i(\mathbf{p}_l)$. \square

B Counterexample: failure of stationarity

We show here that the conclusions of Proposition 4 can fail if the attention strategy satisfies IIA but is non-stationary. Suppose there are three items and the DM is Bayesian with prior beliefs $p_0^j = 1/2$ for each $j \in \{1, 2, 3\}$. The stopping thresholds \underline{p} and \bar{p} and the signal structure for items 1 and 2 are the same as in the counterexample from Section 5.1. The value of item 3 is perfectly revealed as soon as the DM focuses on it for a single period.

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

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

Under β , 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 $1/2$. Let μ be the manipulated strategy that focuses on item 1 in the first period and follows β thereafter. Under μ , the DM first focuses on item 1, and then focuses on item 3—thereby learning its value—in the second period. Thus, under μ , the DM chooses item 1 only if item 3 has value 0 *and* p^1 reaches \bar{p} , which occurs with probability $1/4$.

Proposition 4 does not apply for a reason akin to that in the counterexample from Section 5.1. The baseline and manipulated processes “meet” after two periods in that 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, this cannot happen since, 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 .