

Being in the Right Place: A Natural Field Experiment on List Position and Consumer Choice

Marco Novarese and Chris M. Wilson*

February 2012

Abstract

By manipulating the order in which new economics research papers are presented in email alerts and tracking economists' subsequent download activity, this paper uses a natural field experiment to explore the relationship between list position and consumer choice. Using a novel method, the paper tests and rejects three common explanations for why individuals show a disproportionate tendency to select items listed in top position. Instead, after demonstrating how the nature of such top position effects changes with list length, it suggests that future research should better consider some less-studied and more complex explanations such as bounded rationality and memory limitations.

Keywords: Position Effects, Order Effects, Primacy Effects, Recency Effects, Choice Fatigue, Prominence

JEL Codes: D01, D03, C93, L00

*Novarese: Center for Cognitive Economics, Universit del Piemonte Orientale, Italy. Novarese is the also the director and general editor of New Economics Papers (NEP). Wilson: School of Business and Economics, Loughborough University, UK. We are very grateful to RePEc, NEP and the participating editors. We thank Cesare Tibaldeschi for his data assistance and Andrew Rhodes for his valuable suggestions. We also thank Bernardo Batiz-Lazo, Luke Garrod, Pim Heijnen, Daniel Houser, Matthew Olczak, Claudio Piga, Yuji Tamura, Jidong Zhou and the seminar participants at UEA, Groningen, EARIE 2010 and Jornadas de Economía Industrial 2011 for their helpful comments. Novarese acknowledges the financial support of Regione Piemonte, 'Bando per le Scienze Umane e Sociali', under the project 'Istituzioni, comportamento e mercati in contesti locali e globali'. Corresponding author contact details: c.m.wilson@lboro.ac.uk.

1 Introduction

Whether selecting a meal at a restaurant, searching for a supplier in a directory, choosing a product from a website or comparing firms within a set of search engine results, consumers frequently make choices from lists of options. Previous research has demonstrated that when faced with such lists, individuals often show a disproportionate tendency to select options that are placed at the top. For example, studies show that demand increases markedly for firms that are positioned at the top of search engine results (e.g. Smith and Brynjolfsson 2001, Xu and Kim 2008, Ghose and Yang 2009, Baye et al 2009, Ellison and Ellison 2009), consumers are more likely to click on the links at the top of a list of email advertisements (Ansari and Mela 2003), voters are more inclined to choose candidates placed at the top of ballot papers (e.g. Miller and Krosnick 1998, Koppell and Steen 2004, Ho and Imai 2008, Blocksom 2008, King and Leigh 2009 and Meredith and Salant 2011) and academics tend to download or cite papers that are listed first in journals (e.g. Pinkowitz 2002 and Coupe et al 2010). However, previous research remains far less clear about the explanations for such ‘top position effects’ or ‘primacy effects’. Are top-placed options more likely to be selected simply because the higher quality options have been placed in top positions and if not, why might individuals show a systematic tendency to select options in top position?

Insights into these questions would help understand a variety of important issues across many areas of economics. For example, within a market context, such insights would increase our understanding of the incentives for suppliers to compete for the top positions within search engines (e.g. Armstrong et al 2009 and Athey and Ellison 2010) and the ability of firms to manipulate consumers’ choices through the selection and presentation of their product ranges. Further, and more generally, such insights would also inform governments of how to prompt their citizens to select socially beneficial options, such as healthier food items (Dayan and Bar-Hillel 2011).

This paper investigates the causes of top position effects by using a natural field experiment with a group of subjects that should be the least likely to depart from the predictions of standard theory - economists. Economists often make their research papers available on a well-known online database, Research Papers in Economics (RePEc). Many economists also choose to be kept informed of recent additions to the database by subscribing to a free email alert service conducted by New Economic Papers (NEP), which regularly compiles lists of new papers. By manipulating the order in which items are presented within such lists and measuring subscribers’ subsequent download activity, this paper improves our understanding of the causes of top position effects in two main regards. First, by using a novel method based on the measurement of other position effects, it is able to test and reject three explanatory hypotheses that are commonly referred to within the literature. Second, it demonstrates how the causes of top position effects are inherently

linked to list length. In short lists, top position effects are shown to have a trivial explanation but as list length increases, top position effects become larger and start to exhibit features that indicate a more fundamental source. In view of these results, we then discuss how top position effects may be more consistent with some less-studied and more complex explanations, such as bounded rationality or memory limitations, and urge the literature to better consider such hypotheses in the future.

The first part of the empirical analysis assesses an explanation based on the specific order in which the items have been presented. Under H1 (specific item order), top position effects arise only because an item with a relatively high value happens to have been placed in top position. Unsurprisingly, as NEP usually tries to sort items into some descending order of value, we find that items in top position receive significantly more downloads than items in other positions. However, inconsistent with a pure explanation of specific item order (H1), significant top position effects remain, albeit at a smaller magnitude, even when the order of items is deliberately randomised as part of the experiment.

Using a second, original method, we then test the explanations of choice fatigue and value signals. Under H2 (choice fatigue), top position effects occur only because the costs of evaluating or selecting an item are increasing from top position downwards, as consistent with individuals who consider the items from the top downwards and have total costs of time that are convex. Alternatively, under H3 (value signals), individuals cannot fully assess an item's quality but are more likely to select the item in top position only because they expect, perhaps incorrectly, that a better informed list-maker has deliberately arranged the items in descending order of value. As these hypotheses imply that individuals either face evaluation/selection costs that are increasing from top position downwards or expect items to be arranged in descending order of value, download activity in both cases should be weakly decreasing from the top to the bottom of the list. Yet, contrary to these explanations, we show that items in bottom position receive significantly more downloads than the items in the position immediately above them. In fact, items in bottom position receive significantly more downloads than average, such that the data is characterised by both top position effects and some relatively smaller, 'bottom position effects' or 'recency effects'.

Consequently, the paper rejects the possibility that the common hypotheses H1-H3 can offer a full explanation for top position effects and indicates that their causes remain more complex. To help provide a better explanation, the paper then explores some empirical regularities relating to the role played by the number of items within a list. The size of top position effects is shown to increase with list length. When lists are longer, individuals focus their download activity towards top position. More substantially, by dividing the data into subsamples of short and long lists, the paper then demonstrates how list length has an important impact on the causes of top position

effects. When lists are short, top position effects disappear once the order of items is randomised such that the observations can be fully explained by specific item order (H1). However, when lists are long, top position effects become robust to randomisation and significant bottom position effects begin to emerge in ways that reject all three hypotheses, H1-H3. Increases in list length appear to prompt qualitative changes in individual behaviour which generate larger top position effects that have a more fundamental source. In line with this possibility, the data also offers a suggestion that, in longer lists, download activity may be less sensitive to aspects of each paper’s listed summary information, including the paper’s number of authors, abstract and number of keywords, as consistent with individuals making relatively less use of the listed summary information to guide their decision-making. The final section of the paper then discusses how top position effects may be more consistent with some less-studied and more complex explanations, such as bounded rationality or memory limitations.

The paper continues as follows. Section 2 now provides a detailed comparison of the paper to the existing literature. Before Section 5 reviews the hypotheses and discusses the two empirical tests, Sections 3 and 4 discuss the NEP email alert service, the experimental procedures and the data. The main empirical analysis is conducted in Sections 6 and 7, before Section 8 explores the role of list length. Section 9 discusses some other possible explanations and Section 10 concludes.

2 Previous Literature

The existence of top and/or bottom position effects has been previously well documented in a variety of contexts but our paper contributes to the conceptually distinct and less-studied issue of position effects within individual choice when options are visually presented in a list.¹ Moreover, our paper differs to much past research by focusing on better understanding the *causes* of top position effects, including the role played by the number of listed items. This section now provides a relatively detailed review because many existing studies come from outside economics and have rarely been connected. Previous studies are classified into three settings, which we refer to as ‘limited selection’, ‘unlimited selection’ and ‘repeated choice’.

2.1 Limited Selection Setting. In this setting, individuals face a one-off decision to select one item (or some other fixed number of items) from the list of alternatives. Examples include the

¹Other contexts include how individuals i) use lists of evidence to form impressions or judgments (e.g. Asch 1946, Hogarth and Einhorn 1992), ii) evaluate between alternatives in performance contests or product sampling tests (e.g. Ginsburgh and van Ours 2003, Biswas et al 2010), iii) choose which response to endorse in surveys (e.g. Schwarz et al 1992) and iv) recall items in memory tasks (e.g. Tan and Ward 2000). We return to discuss the effects of memory in Section 9.

decision of which site to buy a DVD from after a price comparison search, which main course to choose at a restaurant, or which candidate(s) to vote for. This is the most common setting for studying top position effects within individual choice but differs from our own setting of ‘unlimited selection’, where individuals are not inherently constrained in the number of items they are willing or able to select.

The market studies listed within the introduction (e.g. Smith and Brynjolfsson 2001, Xu and Kim 2008, Ghose and Yang 2009, Baye et al 2009, Ellison and Ellison 2009; Ansari and Mela 2003) offer some useful evidence of top position effects but they do not focus on understanding their causes. Notably, because the lists were always arranged by a private firm, one cannot rule out an explanation of specific item order (H1). To avoid this possible explanation, Murphy et al (2006) and Dayan and Bar-Hillel (2011) present field experiments where the order of items on a restaurant’s website or menu was randomised. They both still report the existence of positive top position effects, while also documenting some positive bottom position effects. Our paper has some related findings but we use our results to test between competing explanations, while also exploiting variation in the length of lists to further understand the causes of top position effects.

In a related literature, voters are shown to often exhibit a tendency to vote for the candidate placed at the top of the ballot list (e.g. Miller and Krosnick 1998, Koppell and Steen 2004, Ho and Imai 2008, Blocksom 2008, King and Leigh 2009 and Meredith and Salant 2011). As legislation often requires the order of ballots to be determined (quasi-) randomly, these results cannot be explained by specific item order (H1). Instead, most papers jump to an explanation of satisficing (Simon 1955), where individuals start from the top and face marginal inspection costs for each item, which in our different setting, is analogous to choice fatigue (H2). Further results within this literature suggest that top position effects are i) accompanied by lower votes for the candidate in the median ballot position (Meredith and Salant 2011) and ii) larger in ballots with a higher number of candidates (Ho and Imai 2008 and Blocksom 2008). We find some related results but go beyond these papers by using the existence of other position effects to test between a variety of explanations and by showing how list length can have an important effect on the cause, as well as the size, of top position effects.

2.2 Unlimited Selection Setting. This is the setting considered in the current paper, where there is no inherent constraint on the number of items an individual is able or willing to select. Here, individuals are free to select any number of items from a list and the items are sufficiently non-substitutable that it is quite reasonable for individuals to select multiple items. In addition to our download environment, other examples include browsing amongst different items on a website or considering which books to buy from a list of bestsellers.

Previous work in this setting has been minimal and has only considered individuals’ decisions

to download (or cite) academic articles from journal issues. Pinkowitz (2002) uses data from the Journal of Finance website where individuals are able to download two types of papers: those already published and those accepted for publication but not yet assigned a position within an issue. The fact that top-placed papers receive significantly more downloads than other papers even before being assigned to a top position supports an explanation of specific item order (H1). However, the fact that such papers receive an additional download effect after being assigned also points to some other explanation. In a related study, Coupe et al (2010) consider how future citations vary across papers placed in the top five positions within issues of the European Economic Review. Between 1975 and 1997, the journal switched repeatedly between presenting its papers in author-alphabetical order and editor-chosen order. Consistent with an explanation of specific item order (H1), top position effects exist for editor-ordered issues. However, contrary to H1, smaller top position effects also exist within alphabetically-ordered issues. Without any further testing and without an acknowledgment of choice fatigue (H2), both papers explain their results in terms of value signals (H3). In contrast, we show how the measurement of other positions, especially bottom position, can be used to test both H2 and H3, and investigate the important role of list length in understanding top position effects. As explained in Section 6, we also use a less restrictive estimation methodology.

2.3 Repeated Choice Setting. Finally, some related, but somewhat different studies focus on a setting where individuals face a sequence of decisions. Individuals are shown to be more likely to abstain or resort to heuristics after having made a larger number of previous decisions. For example, Levav et al (2010) use a set of field experiments to examine how consumers behave when faced with a series of product customization decisions. They show that consumers are more likely to pick the default option in later decisions and that this effect is stronger when the preceding decisions have involved a relatively large set of alternatives. Augenblick and Nicholson (2011) provide a detailed study of some similar effects in voting behaviour when ballot papers contain *multiple* contests. Using some exogenous variation in the order in which the contests are presented, they show that as contests move down the ballot paper, individuals are more likely i) to abstain from voting, ii) to vote for the default option and iii) to display a bias towards candidates listed first within a contest. Both papers explain these effects by suggesting that the act of making difficult previous decisions reduces the availability of an individual's cognitive resources and Augenblick and Nicholson refer to this as 'choice fatigue'. Although in a different setting, our paper also tests a form of choice fatigue (under H2) but rejects it as an explanation of top position effects due to the presence of positive bottom position effects. This indicates that future studies within this literature may benefit from specifically exploring individuals' decisions at the very end of a series.

3 RePEc and NEP

Before considering the hypotheses and empirical analysis, Sections 3 and 4 now discuss the context of the experiment and the data. Research Papers in Economics (RePEc) is a popular (decentralized) online database of both published and unpublished economics research papers. As part of RePEc, New Economics Papers (NEP) offers a free email alert service. It provides regular email announcements of new papers that have been recently added to the RePEc database. Such announcements are often provided on a weekly basis and are generated for 80 separate research subfields, such as monetary economics or time series. Subscribers can select which subfields they wish to subscribe to and by 2011, NEP had over 64000 total subscriptions.²

Each email announcement has two sections of text. An extract from an example announcement is provided in Appendix 1. The top section informs the reader of how many papers are included within the announcement and includes a list of papers with their titles and authors. If a reader clicks on the title of any paper within the list, or scrolls down, she is taken to the bottom section of the announcement. The bottom section repeats the same list with additional summary information for each paper, including the paper's abstract, keywords, JEL classification codes, date (if these are available) and most importantly, a link to a full text version of the paper. By clicking on a paper's link, a new window is opened and the paper is downloaded. There are no restrictions on the number of different papers an individual can download. For brevity, from this point forward, the terms 'announcement' and 'list' will be used interchangeably.

The announcements for each subfield are managed by an editor, who is a volunteer from academia or the public sector, and are compiled as follows. First, NEP gathers a list of all new papers that have been recently added to the RePEc database. An algorithm then uses past data with information from each paper's title and abstract to arrange these papers into some descending order of estimated value. The list is then passed on to the subfield editors so that they can pick out the papers that are relevant for their own subfield announcement (although some editors may prefer to use a non-sorted version of the list). After selecting their relevant papers, each editor is free to amend the order in which the papers are presented within their subfield announcement or leave them in the order suggested by the algorithm. (Within our sample) about two-thirds of editors usually amend the order of their lists, most with the intention of improving upon the algorithm's attempts to put the more interesting and relevant papers towards the top. Neither the use of the algorithm nor an editor's decision rule is ever made explicit to subscribers.

As later discussed in more detail, it can be the case that a paper is selected to be in the list of more than one subfield. Therefore, to avoid confusion, we will now make a distinction between

²For more, see <http://repec.org/>, <http://nep.repec.org/> or Batiz-Lazo and Krichel (2010) for a brief history.

‘papers’ and ‘items’. An item will refer to an entry on a specific list, whereas a paper will refer to the underlying piece of research that can appear as an item in the lists of multiple subfields.

RePEc measures the download activity for each item in an extremely precise manner. First, it measures the number of downloads that occur specifically via the links contained within NEP announcements, not just those that occur through RePEc more generally. Second, in cases where a paper appears in the list of multiple subfields, RePEc offers a measurement of the number of downloads received for each separate list appearance. It is important to understand this point. The measurement of downloads is item-specific, not paper-specific, and it is this measurement which allows us to analyse the relationship between list position and subsequent download activity in a meaningful manner.

4 Experimental Procedures and Data

After requesting permission from the NEP editors, we were granted the opportunity to measure the number of downloads received by items in lists released between 31st August 2008 and 24th January 2009 across 29 subfields.³ The average number of subscribers per subfield at the start (end) of the sample period was 501 (554).

As the more relevant and popular papers are likely to be placed near the top of lists, one would automatically expect the existence of top position effects. However, to explore the possibility of top position effects in more detail, we were granted the further opportunity to manipulate the order in which items were presented for a small proportion of lists within our sample released after September 20th.

To conduct the experiment, we asked NEP and the relevant editors to continue collecting and ordering their lists as they would do under normal circumstances. However, before the release of any given list, we intervened and randomly allocated the list into one of two groups. Within each subfield, around two-thirds of the lists were allocated to the control group and the remaining lists were allocated to the treatment group. Any list within the control group was sent to subscribers with no alterations. The order of items was left completely unchanged. However, any list within the treatment group had its items re-arranged into a random order with the use of a random number generator. Subscribers were then emailed an identical announcement containing the list in

³The 29 subfields appear representative and cover a wide range of different areas of economics. Their titles are Africa, Ageing, Agricultural, Cognitive and Behavioural, Collective Decision Making, Computational Economics, Dynamic General Equilibrium, Education, Efficiency and Productivity, Time Series, Experimental, Forecasting, Happiness, Health, History and Philosophy, Human Capital, International Trade, Intellectual Property, Knowledge Management, Microfinance, Microeconomics, Migration, Marketing, Monetary, Post Keynesian, Project and Portfolio Management, Risk Management, Sports and Transition.

its new order. Lists were always sent in the common format with no indication of whether the list was randomised or otherwise. No changes were made to the content of the lists and at no point were the subscribers made aware of the experiment.

Using NEP’s item-specific download measurements, the accumulated download activity was then measured for each item on each list from the release of its announcement until a final cut-off date, almost two years later, on the 20th December 2010. The use of a cut-off date could be inappropriate if it prevented the full measurement of downloads, but here the measurement period is extremely long and should capture all related download activity. However, the use of a single cut-off date does imply that lists with different release dates are monitored for different lengths of time. While this should not be problematic given the long measurement period, explicit account for this is made later in the testing and estimation methodologies.

Table 1 provides some summary statistics in regard to the final sample and breaks the data down across the control and treatment groups. (All tables and figures are included at the end of the paper.) After discarding a few announcements that contained only a single item, the final sample includes 577 lists with 6740 listed items which, in aggregate, received 36276 downloads.⁴ The average list contains 11.68 items and the average item received 5.38 downloads. 6% of the items received zero downloads. The 6740 listed items stem from 4985 different papers. Within our sample, an average paper therefore appears on 1.35 subfield lists (or 3.88 subfield lists if one also considers lists outside our sample of 29 subfields). We later address this feature of the data within our estimation procedure.

5 Explanatory Hypotheses and Empirical Tests

As a platform for the empirical analysis, we now review three explanatory hypotheses for top position effects (H1-H3) that are commonly referred to within the literature and outline two empirical tests. As the existing literature contains no theoretical model of top position effects within our unlimited selection setting, we use the following simplified framework to help clarify our discussion.⁵

Consider a list with $n \geq 2$ items. Define the position of item j as p_j , where $p_j = 1$ if item j is in top position and $p_j = n$ if item j is in bottom position. Let individual i ’s true value of downloading and reading item j , V_{ij} , derive from two components, such that $V_{ij} = v_i(s_j) + u_i(\varepsilon_j)$. The first

⁴Two further lists were discarded as outliers. Within these lists, two items (placed last in a list of 10 items and 13th in a list of 18 items) had an excessively large interest with 244 and 222 downloads, relative to the maximum number of downloads in the remaining sample, 69. In a rare occurrence, it is suspected that the NEP list links for these two items were copied onto some other website, with the result of artificially increasing their prominence.

⁵A small literature of theoretical work exists within the limited selection setting where individuals must directly choose between alternatives (e.g. Rubinstein and Salant 2006 and Salant 2011) but this is not easily transferred to our setting. We further discuss this work in Section 9.

component, s_j , can be observed by individual i through inspection of item j 's listed summary information (title, authors, abstract, keywords, JEL codes and date). Individual i can then assign item j an 'observable value', $v_i(s_j)$. However, the second component, ε_j , remains unobservable to individual i until after she has read the item. It could reflect the underlying quality of the paper not captured by, or related to, s_j . The level of ε_j provides individual i with an 'unobservable value', $u_i(\varepsilon_j)$, and we assume individuals prefer higher levels of ε_j , with $u'_i(\varepsilon_j) > 0$. Before downloading item j , individual i can only form an estimate of this second component, $\hat{\varepsilon}_j$.

To explain why individuals might download some items and not others, one must assume the existence of some form of costs. We could introduce a cost of inspecting each item's summary information such that individuals would have to search amongst the items. However, this only adds unnecessary complexity, especially within our non-standard setting of unlimited selection. Instead, and without loss for our illustrative purposes, suppose that individual i can freely inspect each item's summary information and position, but faces a cost of time (or effort) to actually download any given item, c_i . As there is no inherent constraint on the number of items that individuals are able to download, the following simple decision rule is optimal - download any item j if its expected value, $\hat{V}_{ij} = v_i(s_j) + u_i(\hat{\varepsilon})$, is greater than or equal to its associated download cost, c_i .

Top position effects will be defined to exist when the item in top position receives significantly more downloads than the average number of downloads received by items in all other positions. The literature often refers to the following three simple hypotheses. These explanations are not exhaustive, as further discussed in Section 9.

H1: Specific Item Order. Top position effects exist only because individuals can observe that the item in top position happens to have a relatively large value.

Suppose that individual i believes that each item's position reveals no information about its unobservable component, $\hat{\varepsilon}_j = \varepsilon \forall j$. She therefore expects any variation in the value of the listed items to be captured solely by the items' summary information, with $\hat{V}_{ij} = v_i(s_j) + u_i(\varepsilon)$. H1 then suggests that top position effects exist only because of the rather trivial reason that the items happen to be ordered in such a way that the item in top position has a relatively large observable value. For example, given constant download costs, $c_i(p_j) = c \forall j$, individual i could display top position effects when $v_i(s_j) + u_i(\varepsilon) \geq c$ for $p_j = 1$, and $v_i(s_k) + u_i(\varepsilon) < c$ for some $p_k > 1$.

H2: Choice Fatigue. Top position effects exist only because individuals have download costs that are increasing from top position downwards.

Suppose that download costs now vary with item position, $c_i(p_j)$, and that individuals have

download costs such that $c_i(p_j) \leq c_i(p_k)$ for any $p_j < p_k$. H2 then suggests that top placed-items are more likely to be downloaded only because individuals find lower-placed items increasingly costly to download. This would be consistent with the possibility where an individual i) exhibits total costs of time that are convex in the number of downloads she completes and ii) considers whether to download each item sequentially in a strict linear order from top position downwards.⁶

H3: Value Signals. Top position effects exist only because individuals believe (perhaps incorrectly) that the items have been arranged in descending order of value by a better informed agent.

Suppose that individual i believes (perhaps incorrectly) that the list has been arranged by an expert agent m who has the ability to observe each item's observable and unobservable components, s_j and ε_j . Moreover, suppose that individual i believes that agent m has placed the items in descending order according to agent m 's assessment of item values, such that $p_j < p_k$ if $V_{mj} = v_m(s_j) + u_m(\varepsilon_j) \geq V_{mk} = v_m(s_k) + u_m(\varepsilon_k)$. In our context, this would correspond to a situation where an individual has, perhaps through experience, formed the belief that NEP sorts its items by value. H3 then suggests that top-positioned items are more likely to be downloaded only because, on average, individuals expect items with larger unobservable components to be placed in higher positions. For example, consider an extreme case where each item's summary information reveals nothing about its value, $v(s_j) = v(s) \forall j$. As items then vary only in unobservable components, individual i believes that the items have been ordered in line with agent m 's assessment of unobservable value, such that $p_j < p_k$ if $u_m(\varepsilon_j) \geq u_m(\varepsilon_k)$. Individual i then rationally infers that higher placed items must have higher unobservable components, $\varepsilon_j \geq \varepsilon_k$ for $p_j < p_k$, as $u'_m(\cdot) > 0$, and consequently expects the items to be weakly ordered in value from top to bottom, with $V_{ij} = v_i(s) + u_i(\varepsilon_j) \geq V_{ik} = v_i(s) + u_i(\varepsilon_k)$ for any $p_j < p_k$.

To assess the validity of these hypotheses, we propose the following two empirical tests. While Empirical Test I has been used within existing studies, Empirical Test II is new to the literature.

Empirical Test I: Analysis of the Control and Treatment Groups. The explanation of specific item order (H1) can be tested by comparing the control and treatment groups. Under H1, top position effects exist only because an item with a relatively large observable value has been placed in top position. Consequently, under H1, any possible top position effects should *only* arise within the control group, where the items are likely to have been deliberately ordered by some form of value. They should definitely not exist within the treatment group, where the item order

⁶This explanation is analogous to satisficing (Simon 1955) in a limited selection setting, where individuals start from the top and face marginal inspection costs for each item.

has been disrupted through randomisation. Hence, H1 can be rejected as a full explanation of top position effects if such effects are documented within the treatment group.

In contrast, the explanations of choice fatigue (H2) and value signals (H3) are not dependent upon the specific order of items and so it is possible that top position effects could exist in both the control and the treatment groups. Under H2, lower positioned items may be too costly to download regardless of the actual order of items, and under H3, despite the experiment's randomisation of item order, individuals may continue to hold, a now incorrect, belief that top-placed items have the highest value. In order to test these hypotheses, we introduce a new test.

Empirical Test II: Analysis of Other Position Effects. The explanations of choice fatigue (H2) and value signals (H3) can be tested by analysing the possible significance of a broader set of position effects, beyond just top position. Under H2, individuals have download costs that are weakly increasing from top position downwards. Under H3, individuals expect, or have grown accustomed, to item values that are weakly decreasing from top position downwards. In both cases, average download activity should then be weakly decreasing in position. Hence, H2 and H3 can be rejected as pure explanations of top position effects if we observe that the average number of downloads received by items in a given position, p , is significantly larger than the average number of downloads received by items in a preceding position, $p' < p$.

6 Empirical Test I: Comparing the Treatment and Control Groups

This section now uses Empirical Test I to assess the explanatory power of H1 by examining how any top position effects differ between the control and treatment groups. For an initial view of the data, Section 6.1 employs a simple descriptive test, before Section 6.2 uses a more comprehensive estimation methodology.

6.1 Descriptive Analysis. If top position effects do not exist within the data then the download activity for items in top position should be similar to items in other positions. In particular, if one defines the number of downloads received by the item in position p of list l as d_{pl} , then we would expect the number of downloads received by an item in top position averaged across all L lists, $(1/L) \sum_{l=1}^L d_{p=1,l}$, to tend towards the average number of downloads received per item across all positions and all lists, $(1/L) \sum_{l=1}^L (1/n_l) \sum_{p=1}^{n_l} d_{pl}$. As Appendix 2 demonstrates, this statement can be tested by using a simple test statistic, z , based on the central limit theorem. Alternatively, one can also use a non-parametric test, such as a Wilcoxon test. In either case, note that the tests

remain robust to the fact that the lists were monitored for different lengths of time because the tests compare levels of downloads averaged across all lists, rather than between different lists.

Table 2 presents the results for the control and treatment groups separately. It displays the average number of actual downloads received per item in top position, the number of downloads we would expect each top-placed item to receive absent position effects, the percentage difference between the two, and the associated test statistics.

The results suggest that an item in top position within the control group receives 44% more downloads than expected. However, quite contrary to an explanation based purely on specific item order (H1), significant top position effects are also documented within the treatment group. Despite the randomisation of item order, top-placed items are still downloaded 28% more than expected, although this effect appears to be smaller and less significant to that in the control group. At best, H1 therefore seems to offer only a partial explanation for top position effects.

6.2 Further Analysis. These results are now examined in more detail by estimating the relationship between the number of downloads received by each item, d_{pt} , and a top position dummy variable, $T_{p=1}$, while controlling for a range of other variables.

Any estimation strategy must account for two important features of the data. First, item downloads can only take the form of a non-negative integer. Previous research using download (or citation) data has not fully addressed this issue. Some studies have ignored it (e.g. Pinkowitz 2002), while others have treated it in a restrictive way. For example, Coupe et al (2010) employ a Poisson model which tackles the problem by specifying that the dependent variable is drawn from a Poisson distribution with conditional density, $f(y|x) = (e^{-\lambda})(\lambda^y/y!)$ for $y = 0, 1, 2, \dots$, where the explanatory variables enter via $\lambda = e^{x\beta}$. However, the Poisson model makes the strong assumption of equi-dispersion, requiring the conditional mean to be equal to the conditional variance, $E(y|x) = Var(y|x)$. As later confirmed in our results, this property is known to be rarely observed in practice, and consequently, the Poisson model typically under-estimates standard errors. To resolve this issue, we use a quasi-maximum likelihood estimator based on the Poisson distribution (Poisson QMLE). The Poisson QMLE fully recognises that the Poisson distribution may be inappropriate, but persists in using it in the knowledge that i) the coefficient estimates are still consistent, ii) one can correct for the biased standard errors by using a robust estimator for the variance-covariance matrix under a weaker assumption, $E(y|x) = \sigma^2 Var(y|x)$, (where σ^2 is a parameter to be estimated), and usefully, iii) the model is also robust to any further forms of mis-specification.⁷

⁷See Wooldridge (1999) for more details on the QMLE Poisson and its relative advantages. An alternative approach could use a negative binomial model, but this is argued to be less robust. Unless otherwise stated, all our main results can be replicated using the negative binomial model.

Second, as discussed previously, some papers within the sample are included as items on the lists of more than one subfield. This implies that some item-level observations that share the same underlying paper are unlikely to be independent, resulting in invalid estimates of the standard errors. To resolve this issue, we cluster the standard errors by paper. This approach allows the error terms of a group of observations with the same underlying paper to have a correlated error structure, while maintaining the assumption of independent errors for groups of observations with different underlying papers.

Consequently, we proceed to model the downloads received by each item, d_{pl} , using a QMLE Poisson model with clustered standard errors. In particular, we first estimate Model 1 on the control and treatment groups separately. Model 1 specifies the explanatory variables, $x'_{pl}\beta$, as

$$\beta_0 + \beta_1 T_{p=1} + z'_l \beta_z + q'_{pl} \beta_q \quad (1)$$

where $T_{p=1}$ is a dummy variable which equals one only if an item is in a top position ($p = 1$), and where z_l and q_{pl} are vectors of list-specific and item-specific control variables, respectively. Any top position effects will be captured by the estimated value of β_1 .

The variables included in the list-specific controls, z_l , are summarised in Table 3. These include the total number of items within the item's list, the number of subscribers registered on the list's subfield email alert at the date of the list's release, and a set of subfield fixed effects. As the lists vary in the length of time for which the downloads were monitored, we also include a measure of each list's 'availability' by calculating the number of days between the list's release date and the final observation cut-off date. In the vast majority of results, this variable proves insignificant due to our long measurement period.

The item-specific control variables, q_{pl} , are summarised in Table 4 and are constructed using each item's summary information. In particular, they refer to the item's title language, length of title, number of authors, length of abstract, length of keywords and number of JEL classification codes.⁸ Finally, we also include a measure of the total number of lists (within the entire population of NEP), in which each item's underlying paper appeared. This variable could be negatively correlated with our item-specific measure of downloads if it captures the possibility that a paper's total downloads are dispersed over a larger number of lists. Alternatively, it could be positively correlated if it is related to the item's general appeal.

In addition to Model 1, we also estimate Model 2 on the entire sample in order to formally test how the size of any top position effects differs between the control and treatment groups. As summarised in (2), Model 2 simply includes two additional variables: $treat_l$ is a dummy variable

⁸Announcements also include information about the year in which each item was written but almost all items were written at a time very close to the list release date and this proved insignificant in estimations.

that equals one only if list l is in the treatment group, while $T_{p=1} * treat_l$ is an interaction term which will capture any difference in the size of top position effects between the two groups.

$$\beta_0 + \beta_1 T_{p=1} + z'_l \beta_z + q'_{pl} \beta_q + \beta_T treat_l + \beta_{1T} T_{p=1} * treat_l \quad (2)$$

It is important to remember that NEP uses data from each item's summary information to assign it a list position and that list position may, in turn, affect an item's downloads. Consequently, within the control group, we should be wary of interpreting the item-specific control variables which derive from each item's summary information because they may affect downloads both directly and indirectly, through their effect on item position. As the item-specific control variables may also be correlated with the top position dummy, there may be a further problem of multicollinearity creating inefficient coefficient estimates. To improve robustness in this regard and more generally, all later results are presented not just in the main specification with both the list-specific and item specific controls (specification i), but also in specifications that omit the item-specific controls (ii) and omit the list-specific controls (iii). In addition, we should remember that these problems cannot be present in the treatment group where item order is unrelated to the items' summary information.

Table 5 presents the results for Model 1 across the control and treatment groups separately and for Model 2 across the entire sample. After considering several possible quadratic terms, we include them where significant. For brevity, only the calculated marginal effects are presented with their associated z-statistics. In addition, each estimation displays the results of a Wald test for the joint significance of all coefficients, χ^2 , the Akaike Information Criterion, AIC , and an estimate of the mean-variance ratio, $\hat{\sigma}^2$, as described above. In support for our chosen methodology and in rejection of the basic Poisson model, the estimates suggest a mean-variance ratio that is always substantially greater than one.

The results offer some robust evidence for the existence and causes of top position effects. In line with the descriptive findings, Table 5 suggests that items in top position within the control group receive an average of 2.6-3.2 more downloads than items in other positions. When compared to the average number of downloads per item, 5.46, this implies a large effect of around 50-60%. However, top position effects are also significant within the treatment group. Items in top position within the treatment group receive an average of 1.6-2.2 (30-40%) more downloads than items in other positions. The fact that top position effects remain significant despite the randomisation of item order rules out an explanation based purely on the specific order of items (H1). Nevertheless, while it is clear that specific item order cannot be the main cause of top position effects, the randomisation of item order does significantly weaken the size of the estimated top position effects. By considering the effect of the interaction term, $T_{p=1} * Treat_l$, in Model 2, one can see that top

position effects are significantly smaller in the treatment group by a factor of about 20%. This suggests that H1 has a minor explanatory role.⁹

Before further investigating the other possible causes of top position effects, we briefly comment on some interesting side results in Table 5 concerning the effects of the control variables on item downloads. To minimise any of the interpretation problems discussed above, these effects are best viewed within the treatment group. First, items with an English title receive more downloads. Second, downloads are decreasing in an item's length of title (at an increasing rate). This fits with Emerald Journals' advice to keep titles relatively short to attract more attention.¹⁰ Third, perhaps surprisingly, items with a higher number of authors receive fewer downloads. Fourth, the length of an abstract provides no effect, but items with no abstract receive higher downloads. This is consistent with the possibility that subscribers download such papers in order to find out whether they wish to read them. Fifth, the number of JEL codes has no effect on downloads, but items with zero JEL codes are associated with fewer downloads. Sixth, downloads are increasing in the length of an item's selection of keywords. Seventh, downloads are also increasing in the number of lists in which the item's underlying paper appears, perhaps reflecting the paper's general appeal. Eighth, as expected, a higher number of subscribers is associated with higher item downloads. Finally, as further discussed in Section 8, an item's downloads are decreasing in the number of items included within its list.

For brevity, we do not always continue to report the full results of the control variables in later tables, even when they are included in the estimations. However, we will return to further discuss their effects in Section 8.

7 Empirical Test II: Analysing Other Position Effects

Having found that H1 can only offer a minor explanation for top position effects, this section now assesses the explanations of choice fatigue (H2) and value signals (H3) by appealing to Empirical Test II which involves the measurement of a broader set of position effects. As the two hypotheses imply that either download costs are weakly increasing, or that expected item values are weakly decreasing, from top position downwards, Empirical Test II suggests that H2 and H3 can both be

⁹Within the presented analysis, we have ignored the fact that some of the lists belong to a subfield where the editor usually sorts the lists beyond that determined by the algorithm. Breaking down the data further in this regard does not offer any major insights and adds substantial complication to the results. Within the control group, top position effects are observed to be larger in editor-ordered lists as consistent with the intuitive possibility that additional editor ordering improves the accuracy of the algorithm's ordering. However, within the treatment group, there are no differences between lists in subfields that usually receive additional editor sorting and those that do not.

¹⁰<http://www.emeraldinsight.com/authors/guides/promote/optimize1.htm?PHPSESSID=7602t7ftgvi5r6g7shm06gg6c6>, 01/04/2011.

rejected as pure explanations of top position effects if the average number of downloads received in a given position, p , is significantly larger than the average number of downloads received in a preceding position, $p' < p$.

However, any study of position effects (beyond top and bottom position), can force a necessary reduction in sample size. This problem arises because positions can become ill-defined in lists that contain too few items. For example, in a study of three positions; top, second and bottom, one would have to drop all lists that contain only two items because the second and bottom position are not uniquely defined. Consequently, there is a trade-off - the inclusion of more positions may bring additional insights but it may also generate a reduction in sample size.

To help decide how many positions to focus on, Figure 1 displays the average number of downloads received per-item across 8 positions (top, second, third, fourth, fourth-from-bottom, third-from-bottom, second-from-bottom and bottom) and compares this to the average number of downloads received across all positions, using only the 391 lists within the sample that contain 8 or more items. This offers two suggestions. First, one should definitely include bottom position as it appears that, contrary to H2 and H3, items in bottom position may receive significantly more downloads than items in the position immediately above it. Second, the four interior positions (third, fourth, fourth-from-bottom and third-from-bottom) appear to be relatively unimportant as their effects are roughly decreasing in size and quite weak when compared to the sample average. Consequently, we focus on the four ‘exterior’ positions (top, second, second-from-bottom and bottom) using a sample of all lists with four or more items. However, our conclusions remain robust under a variety of different approaches.¹¹

To estimate such position effects, we extend the estimation methodology used in section 6.2 by estimating Model 3 on the control and treatment groups separately. As summarised in (3), Model 3 estimates item downloads as a function of four position dummies, $\{T_{p=1}, T_{p=2}, T_{p=n_i-1}, T_{p=n_i}\}$, and the original list-specific and item-specific control variables. The position effects will be captured by the estimated coefficients, β_1 to β_4 .

$$\beta_0 + \beta_1 T_{p=1} + \beta_2 T_{p=2} + \beta_3 T_{p=n_i-1} + \beta_4 T_{p=n_i} + z'_i \beta_z + q'_{pl} \beta_q \quad (3)$$

Empirical Test II suggests that H2 and H3 can be rejected if there exists a position effect that is significantly larger than an effect observed in a preceding position. Without loss for our purposes, we perform three Wald tests to assess the equality of position effects in ‘adjacent’ positions; $\beta_1 -$

¹¹In particular, we have tested the following alternative approaches, i) (top, bottom) with all lists, ii) (top, second, third, third-from-bottom, second-from-bottom, bottom) with all lists of 6 or more, and iii) (top, second, third, fourth, fourth-from-bottom, third-from-bottom, second-from-bottom, bottom) with all lists of 8 or more. The results from ii) and iii) also confirm the insignificance of the four ‘interior’ positions, in line with Figure 1.

$\beta_2 = 0$, $\beta_2 - \beta_3 = 0$ and $\beta_3 - \beta_4 = 0$. Any rejection of these tests with $\beta_{j-1} - \beta_j < 0$ will provide evidence against H2 and H3.

While the effect of item randomisation on the position effects is no longer the main focus, it will still be useful to assess how randomisation impacts upon the estimated coefficients. Therefore, in a similar vein to Section 6.2, we estimate a further model with some additional terms using the whole sample. In particular, Model 4 comprises of Model 3 plus $treat_i$ (which equals one if the list is within the treatment group) and a vector of terms that interact each of the position dummies with $treat_i$.

The results are presented in Table 6. As before, it contains the marginal effects and z-statistics, but in addition, the bottom of Table 6 also includes the results of the Wald tests and uses a minus sign, $(-)$, to highlight any cases where, contrary to H2 and H3, $\beta_{j-1} - \beta_j < 0$. For ease of reference, Figure 2 plots the estimated position effects for the control and treatment groups using the results from Table 6 under the main specification, 3(i).

First, if, like many previous studies, we ignore the importance of bottom position effects, our results appear initially consistent with choice fatigue (H2) and value signals (H3). In the control group, the top position effect is accompanied by a positive second position effect and overall, downloads are strictly decreasing in position. Similarly, in the treatment group, while item randomisation eliminates the second position effect, downloads are weakly decreasing in position as consistent with H2 and H3. However, once we consider bottom position, our results provide a very different conclusion. Contrary to a pattern of weakly-decreasing downloads, items in bottom position are reported to receive significantly more downloads than the items placed in the position immediately above them. Moreover, items in bottom position actually receive significantly more downloads than average. While weaker than the documented top position effects, these ‘bottom position effects’ are extremely robust, as item randomisation has no significant effect on their magnitude within the treatment group. The presence of such bottom position effects rules out the common hypotheses of choice fatigue (H2) and value signals (H3) as pure explanations for top position effects. It cannot be that top position effects exist just because individuals find it increasingly costly to make downloads from lower positions (H2). Nor can it be that top position effects exist just because individuals expect, or have grown accustomed to, the possibility that NEP arranges its items in descending order of value (H3). Such hypotheses cannot explain the disproportionate download activity for items in bottom position. However, in Section 9 where we discuss some further explanations for our results, we do note that these findings could still be consistent with some more complex versions of choice fatigue and value signals.

Finally, we must stress that the existence of bottom position effects is not artificially generated by the two-section design of the NEP announcements. As explained in Section 3 and illustrated in

Appendix 1, by clicking on the top item within the upper section of the announcement an individual is indeed taken to the summary information of the top item in the lower section. However, when inspecting this summary information, it is not the case that the individual’s attention is artificially drawn upwards to the bottom item within the upper section of the announcement. Instead, the bottom item is off-screen and the individual would have to deliberately scroll upwards in order to see this item.

8 Further Empirical Analysis - The Role of List Length

The results offer no easy explanation of top position effects. Specific item order (H1), choice fatigue (H2) and value signals (H3) have been ruled out as pure explanations for top position effects due to the existence of bottom position effects and due to the fact that top position effects remain, even when the order of items has been randomised. Consequently, the documented top (and bottom) position effects must have some other more complex cause. To understand more, this section now explores some further features of the data. In particular, it considers the role of list length. In addition to showing how list length can affect the size top position effects, we step further beyond the existing literature by demonstrating how list length has an important influence on the cause of top position effects.

Across the sample, the number of items included within most lists varies from 2-54.¹² This variation arises, not only from differences in the supply of academic papers over time within subfields, but also from differences between subfields, with some subfields typically having longer or shorter lists. Hence, we should be clear that the variation in list length may not be exogenous. For example, as subfield editors may vary in their selectivity, items from relatively long lists may be more likely to belong to subfields with less selective editors with lists that typically contain items of relatively lower (expected) value or quality. While the subfield fixed effects may offer some control for this possibility, all results should be viewed with this caveat in mind.

Within the limited selection setting, Ho and Imai (2008) and Blocksom (2008) show that top position effects can be increasing in list length. In a related search context, Meyer (1997) also shows that laboratory subjects conduct fewer searches to select an item when faced with longer lists. Within our different, unlimited selection setting, we now extend this evidence by examining how our four position effects vary with list length. In particular, we estimate Model 5, summarised by (4). Like Model 3, it contains the four position dummies and the control variables but it also includes a set of interaction terms, $T_p * n_l$, to capture how each position effect varies with list length, n_l , (and specifies list length as a separate variable, rather than as just a member of the

¹²One outlying list contains 118 items but its exclusion does not affect the results.

control variables).¹³

$$\begin{aligned}
& \beta_0 + \beta_1 T_{p=1} + \beta_2 T_{p=2} + \beta_3 T_{p=n_l-1} + \beta_4 T_{p=n_l} \\
& + \beta_{1n} T_{p=1} * n_l + \beta_{2n} T_{p=2} * n_l + \beta_{3n} T_{p=n_l-1} * n_l + \beta_{4n} T_{p=n_l} * n_l \\
& + \beta_n n_l + z'_l \beta_z + q'_{pl} \beta_q
\end{aligned} \tag{4}$$

The results are presented in Table 7. While the effects of second, second-from-bottom and bottom position remain independent of list length, top position effects become significantly larger in lists that contain more items, especially in the treatment group. As list length increases, individuals appear to focus their download activity more towards items in top position.¹⁴

We now take a more original step beyond the existing literature by demonstrating how list length can have a fundamental role in determining, not only the *size* of top position effects, but also their *cause*.

The estimations of the four position effects from Section 7 (Models 3 and 4) are now repeated on two subsamples of ‘short’ and ‘long’ lists separately. The two subsamples are created by dividing the sample of lists with 4 or more items around its median list length of 11. The subsample of short lists then includes 282 lists that contain between 4-11 items, while the subsample of long lists includes the 250 lists that contain 12 or more items. The results for the estimations and the corresponding Wald tests are presented in Table 8 for short lists, and in Table 9 for long lists. For ease of reference, the estimated position effects from the main model specification, 3(i), are also displayed graphically for short and long lists in Figure 3.

First, consider the results for short lists. The findings are now wholly consistent with the rather trivial explanation of specific item order (H1). Within the control group, significant position effects are documented in top and second position. However, both of these effects are completely eliminated once item order has been randomised within the treatment group. Indeed, after randomisation, there is little evidence of any position effects whatsoever within short lists. This also adds further evidence against the possibility that position effects are driven by some form of habit.

Next, consider the results for long lists. The contrast is dramatic. Again, top position and second position effects are documented within the control group but now there are two big differences. First, randomisation of item order has *no* significant effect on the size of the top position

¹³Once again, the results remain robust using a different number of positions. The addition of further interaction terms to allow for the effects of list length squared proved insignificant.

¹⁴In addition, as also previously reported in Table 5, it is interesting to note how list length itself exerts a strong negative effect on the number of downloads received by each item regardless of item position (via the variable, n_l). While one can check that the aggregate number of downloads received by all items within each list is increasing in list length, this result suggests that, on average, an individual item within a list that includes relatively more items, receives significantly fewer downloads. An explanation for this finding involving the caveat above cannot be ruled out - longer lists could contain items of lower (expected) value.

effects, suggesting that top position effects are more robust in longer lists and cannot be explained by specific item order (H1). Second, in both the control and the treatment groups, items in bottom position now receive significantly more downloads than items in the penultimate position and significantly more downloads than average, ruling out choice fatigue (H2) and value signals (H3) as pure explanations. Consequently, in contrast to short lists, top position effects in long lists are now inconsistent with all three explanations, H1-H3.¹⁵

In summary, increases in list length generate i) larger top position effects, ii) more robust top position effects and iii) significant bottom position effects. This indicates that the source of top position effects is inherently linked to list length. In longer lists, top position effects have a cause that is deeper and more complex than the simple hypotheses commonly referred to within the literature, H1-H3.

These findings also suggest that increases in list length may trigger qualitative changes in individuals' behaviour. In line with this possibility, we note one final effect of list length with respect to the item-specific control variables. By viewing the estimation results within the treatment group to minimise any interpretation concerns (as discussed in Section 6), one can see that i) fewer item-specific control variables are significant within the long list sample relative to the short list sample, and ii) within the long list sample, the only significant item-specific control variables relate to an item's title. This suggests that, when faced with longer lists, individuals may reduce their use of the listed summary information to consider only the item title as a guide for their decision-making. However, because the variation in list length is not exogenous, we cannot rule out the alternative possibility that individuals make less use of the provided information in longer lists because items in longer lists have less informative summary information.¹⁶

9 Possible Explanations

In light of these results, rather than the common hypotheses H1-H3, this section now discusses how top position effects may be more consistent with some less-studied and more complex hypotheses

¹⁵The results here differ very slightly if one uses the alternative estimation procedure involving a negative binomial model or if one divides the sample in two at the median observation which gives subsamples with 4-14 and 15+ items rather than dividing the sample by the median list length. In these cases, top position effects within the treatment group of the short lists subsample are marginally significant implying that top position effects in the short sample are only largely, rather than wholly, consistent with H1. However, the differences between short and long lists remains clear.

¹⁶By observing the variable, *treat*, in Model 4, one can also note that randomisation appears to actually increase item downloads within short lists. One explanation may suggest that individuals realise that the items are not ordered as usual and decide to inspect the lower positions more thoroughly than they would do normally. However, this effect is not present in longer lists.

and urges further research to better consider such explanations in the future.

First, bounded rationality offers an excellent description of our results. When faced with cognitively demanding tasks, it is well known that individuals often replace fully rational decision strategies with biased heuristics to economise on cognitive resources (e.g. Payne et al 1993). By explicitly analysing the complexity of alternative choice rules, Salant (2011) demonstrates this within the context of choice within lists. In longer lists, when the decision is more difficult, he shows how it may be optimal for individuals to switch to a decision rule that is simpler than the rational decision rule and that simpler decision rules necessarily generate order effects, such as top and bottom position effects. This explanation seems highly consistent with our findings and also with the suggestion in the data that individuals may become less dependent upon using each item's full summary information to guide their download decisions in longer lists. However, Salant's model is constructed under the assumption that individuals select one item from a list and it remains unclear how his results extend to a setting such as ours with unlimited selection.

Second, we suggest that our findings are also consistent with an explanation of memory limitations. Rather than making all their download decisions in sequence as they progressively inspect each item, some individuals may move back and return to download a previously inspected item. In this scenario, individuals may be more likely to return to items that they can better remember. This could then lead to top and bottom position effects because individuals are well known to exhibit biases in memory recall towards items listed first and last (e.g. Tan and Ward 2000). Moreover, as consistent with our results, evidence suggests that these effects become more pronounced in longer lists (e.g. Ward et al 2010).

Finally, while we have rejected the common simple explanations of choice fatigue (H2) and value signals (H3), some more complex versions of these hypotheses may offer an account for our findings. For example, consider a modified version of choice fatigue where a fraction of individuals find it increasingly costly to download items but inspect the list from bottom to top rather than from top to bottom. As consistent with the results, in aggregate, items in top and bottom position would then be more likely to receive downloads in lists that are sufficiently long to induce fatigue effects. Alternatively, in a setting where individuals have positive inspection costs, it is possible that bottom position effects may also arise when individuals inspect items from top to bottom and either i) have additional recall costs of returning to a previously searched item or ii) progressively learn more about the quality distribution of items as they search. However, these latter explanations appear less able to account for the documented patterns regarding list length.¹⁷

¹⁷Although not an explanation in itself, we also note that top (and bottom) position effects could be enhanced by a further factor. If there is an advantage from selecting the same items as other individuals and if individuals know that other individuals exhibit top (and bottom) position effects, then it may be optimal to also select items in top (and bottom) position.

10 Conclusion

This paper uses a natural field experiment to analyse the causes of top position effects. When lists are short, the paper shows that any tendency for individuals to select items located at the top of a list is totally eliminated once the order of items has been randomised. This suggests that top position effects existed only because of the specific order in which the items were presented. However, as lists grow longer, our results indicate that individuals display several qualitative changes in behaviour. Individuals start to concentrate their download activity towards items in top position, generating top position effects which are larger and more robust even when the order of items has been randomised, and individuals begin to exhibit positive bottom position effects. The paper then uses these findings to reject three explanations for top position effects that are commonly cited within the literature and suggests, instead, that they may be more consistent with some less-studied and more complex explanations, such as bounded rationality and memory limitations.

While we cannot formally test between these further hypotheses using our current data, we urge the literature to better consider such explanations in the future and suggest that the development of eye-tracking software may offer much hope in this regard. Expanding on some early work on advertising by Lohse (1997) and others, Reutskaja et al (2011) use such software to analyse subjects' choices of snacks from a grid of options. Among many other results, they find that subjects look more frequently at, and are more likely to choose, items located in the top left-hand corner, or the middle, of the grid. We imagine that the full application of such techniques to the study of position effects in lists is likely to be very fruitful.

References

- [1] Ansari A. and Mela C.F. (2003) "E-customization" *Journal of Marketing Research* vol.40 p.131-145
- [2] Armstrong M., Vickers J. and Zhou J. (2009) "Prominence and Consumer Search" *RAND Journal of Economics* vol.40 p.209-233
- [3] Asch S.E. (1946) "Forming Impressions of Personality" *Journal of Abnormal and Social Psychology* vol.41 p.258-290
- [4] Athey S. and Ellison G. (2010) "Position Auctions with Consumer Search" *Quarterly Journal of Economics*, forthcoming

- [5] Augenblick N. and Nicholson S. (2011) "Ballot Position, Choice Fatigue, and Voter Behavior" Working Paper, July
- [6] Batiz-Lazo B. and Krichel M. (2010) "The Creation of Internet Communities: A Brief History of On-line Distribution of Working Papers through NEP, 1998-2010" Working Paper, November
- [7] Baye M., Gatti J.R.J., Kattuman P. and Morgan J. (2009) "Clicks, Discontinuities and Firm Demand Online" *Journal of Economics and Management Strategy* vol.18 p.935-975
- [8] Biswas D., Grewal D. and Roggeveen A. (2010) "How the Order of Sampled Experiential Products Affects Choice" *Journal of Marketing Research* vol.47 p.508-519
- [9] Blocksom D. (2008) "Moderators of the Name-Order Effect: The 2004 Presidential Election in State Ohio" Thesis, Stanford University, available at http://economics.stanford.edu/files/Honors_Theses/Theses.2008/Blocksom,%20D.%202008.pdf
- [10] Coupe T., Ginsburgh V. and Noury A. (2010) "Are Leading Papers of Better Quality? Evidence from a Natural Field Experiment" *Oxford Economic Papers* vol.61 p.1-11
- [11] Dayan E. and Bar-Hillel M. (2011) "Nudge to Nobesity II: Menu Positions Influence Food Orders" *Judgment and Decision Making* vol.6 p.333-342
- [12] Ellison, G. and Ellison, S.F. (2009) "Search, Obfuscation and Price Elasticities on the Internet" *Econometrica* vol.77 p.427-52
- [13] Ghose A. and Yang S. (2009) "An Empirical Analysis of Search Engine Advertising: Sponsored Search in Electronic Markets" *Management Science* vol.55 p.1605-1622
- [14] Ginsburgh V.A. and van Ours J.C. (2003) "Expert Opinion and Compensation: Evidence from a Musical Competition" *American Economic Review* vol.93 p.289-296
- [15] Ho D.E. and Imai K. (2008) "Estimating Causal Effects of Ballot Order from a Randomized Natural Experiment: California Alphabet Lottery 1978-2002" *Public Opinion Quarterly* vol.72 p.216-240
- [16] Hogarth R.M. and Einhorn H.J. (1992) "Order Effects in Belief Updating: The Belief-Adjustment Model" *Cognitive Psychology* vol.24 p.1-55
- [17] King A. and Leigh A. (2009) "Are Ballot Order Effects Heterogeneous?" *Social Science Quarterly* vol.90 p.71-87

- [18] Koppell K. G. S. and Steen J.A. (2004) "The Effects of Ballot Position on Election Outcomes" *Journal of Politics* vol.66 p.267-281
- [19] Levav J., Heitmann M., Hermann A. and Iyengar S. (2010) "Order in Product Customization Decisions: Evidence from Field Experiments" *Journal of Political Economy* vol. 118 p.274-299
- [20] Lohse G.L. (1997) "Consumer Eye Movement Patterns on Yellow Pages Advertising" *Journal of Advertising* vol.9 p93-123
- [21] Meredith M. and Salant Y. (2011) "On the Causes and Consequences of Ballot Order Effects" forthcoming, *Political Behavior*
- [22] Meyer R.J. (1997) "The Effect of Set Composition on Stopping Behavior in a Finite Search Among Assortments" *Marketing Letters* vol.8 p.131-143
- [23] Miller J.M. and Krosnick J.A. (1998) "The Impact of Candidate Name Order on Election Outcomes" *Public Opinion Quarterly* vol.62 p.291-330
- [24] Murphy J., Hofacker C. and Mizerski R. (2006) "Primacy and Recency Effects on Clicking Behavior" *Journal of Computer-Mediated Communication* vol.11 p.522-535
- [25] Payne J.W., Bettman J.R. and Johnson E.J. (1993) "The Adaptive Decision-Maker" New York: Cambridge University Press
- [26] Pinkowitz L. (2002) "Research Dissemination and Impact: Evidence from Website Downloads" *Journal of Finance* vol.57 p.485-499
- [27] Reutskaja E., Nagel R., Camerer C.F. and Rangel A. (2011) "Search Dynamics in Consumer Choice under Time Pressure: An Eye-Tracking Study" *American Economic Review* vol.101 p.900-926
- [28] Rubinstein A. and Salant Y. (2006) "A Model of Choice from Lists" *Theoretical Economics* vol.1 p.3-17
- [29] Salant Y. (2011) "Procedural Analysis of Choice Rules with Applications to Bounded Rationality" *American Economic Review* vol.101 p.724-48
- [30] Schwarz N., Hippler H.-J. and Noelle-Neumann E. (1992) "A Cognitive Model of Response Order Effects in Survey Measurement" in "Context Effects in Social and Psychological Research" (Eds) Schwarz N. and Sudman S., New York: Springer
- [31] Simon H. A. (1955) "A Behavioral Model of Rational Choice" *Quarterly Journal of Economics* vol.69 p.99-118

- [32] Smith M. and Brynjolfsson E. (2001) “Consumer Decision-Making at an Internet Shopbot: Brand Still Matters” *Journal of Industrial Economics* vol.49 p.541-558
- [33] Tan L. and Ward G. (2000) “A Recency-Based Account of Primacy Effects in Free Recall” *Journal of Experimental Psychology: Learning, Memory and Cognition* vol.26 p.1589-1625
- [34] Ward G., Tan L. and Grenfell-Essam R. (2010) “Examining the Relationship Between Free Recall and Immediate Serial Recall: The Effects of List Length and Output Order” *Journal of Experimental Psychology: Learning, Memory and Cognition* vol.35? p.1207-1241
- [35] Wooldridge J.M. (1999) “Quasi-Likelihood Methods for Count Data” Ch.8 in ‘Handbook of Applied Econometrics Volume II: Microeconomics’, Pesaran M.H and Schmidt P. (Eds), Blackwell
- [36] Xu Y. and Kim H.-W. (2008) “Order Effect and Vendor Inspection in Online Comparison Shopping” *Journal of Retailing* vol.84 p.477-486

Appendix:

Appendix 1: An Example Announcement

NEP: New Economics Papers
Cognitive and Behavioural Economics

Edited by: Marco Novarese
<http://econpapers.repec.org/RAS/pno2.htm>
 University of the Piemonte Orientale
 Date: 2008-11-18
 Papers: 7

This document is in the public domain, feel free to circulate it.

+++++

+ Note: Access to full contents may be restricted +

+ NEP is sponsored by SUNY Oswego +

+++++

In this issue we have:

1. We-thinking and 'double-crossing': frames, reasoning and equilibria
Smerilli, Alessandra
2. Is Libertarian Paternalism an Oxymoron?
Christophe Salvat
3. Follow the Leader: Simulations on a Dynamic Social Network
David Goldbaum
4. In what sense do firms evolve?
Bart Nooteboom
5. The Dynamics of Learning: An Economic Model of Student Motivation and Achievement
Barkley, Andrew
6. Heterogeneity, Bounded Rationality and Market Dysfunctionalit
Xue-Zhong He; Lei Shi
7. Knowledge Creation and Sharing in Organisational Contexts: A Motivation-Based Perspective
Lam, Alice; Lambermont-Ford, Jean-Paul

1. We-thinking and 'double-crossing': frames, reasoning and equilibria

Smerilli, Alessandra

The idea of we-thinking, or we-reasoning, is increasingly drawing the attention of more and more economists. The two main contributors are Bacharach and Sugden, and they approach the topic in two different ways. Sugden's aim is to show that we-reasoning is a consistent and logical way of thinking, but he does not face the problem of how we-reasoning can arise. Bacharach's theory is based on frames and his never reached aim (because of his death) was to explain we-thinking in terms of Variable Frame Theory. But some of his intuitions conflict with the logical analysis he proposes. In the present paper, I take a different approach to the way in which we-thinking works. Based on a not fully developed intuition of Bacharach's, i.e. the 'double-crossing'

problem in Prisoners' Dilemma (PD) game, I propose a framework in which a person is allowed to have both I-thoughts, when she is we-reasoning, and we-concepts, when she is I-reasoning, and develop my analysis in terms of equilibrium concepts.

Keywords: we-thinking; frames; we-equilibria

JEL: D70 Z10 C70

Date: 2008-11-02

URL: <http://d.repec.org/n?u=RePEc:pra:mprapa:11545&r=cbe>

...The remaining items in positions 2 to 7 are then presented sequentially in a similar format.

Appendix 2: Construction of the Descriptive Test Statistic

The test is related to that used by Koppell and Steen (2004) and Meredith and Salant (2011) in the context of ballot ordering. Let d_{pl} derive from some unspecified list-specific random variable with finite mean, μ_l , and variance, σ_l^2 . Under a null hypothesis of no position effects, the following two statements should be true. First, the average number of downloads per item in list l , $(1/n_l) \sum_{p=1}^{n_l} d_{pl}$, should tend towards the mean for list l , μ_l . Second, the average number of downloads per-item observed in position p across all L lists, $(1/L) \sum_{l=1}^L d_{pl}$, should tend towards the average mean across all L lists, $(1/L) \sum_{l=1}^L \mu_l$. We wish to test if these relationships hold for $p=1$. For any set of parent distributions, the Lindberg-Fuller version of the central limit theorem suggests that the test statistic, $\sqrt{L}[(1/L) \sum_{l=1}^L d_{p=1,l} - (1/L) \sum_{l=1}^L (1/n_l) \sum_{p=1}^{n_l} d_{pl}]$, is asymptotically distributed by $N(0, \sigma^2)$ where σ is the average standard deviation in the limit. By rearranging and replacing σ with its sample variant, $s = \sqrt{(1/L) \sum_{l=1}^L s_l^2}$ (where s_l^2 is the sample variance of list l), the test statistic, $z = [\sum_{l=1}^L d_{p=1,l} - \sum_{l=1}^L (1/n_l) \sum_{p=1}^{n_l} d_{pl}] / \sqrt{\sum_{l=1}^L s_l^2}$, is distributed by $N(0, 1)$.

Tables and Figures

Throughout the tables, we use two-tailed tests and employ *, ** and *** to denote significance at 5%, 1% and 0.1%. The reported Wilcoxon test statistic is in the form of its equivalent z-value.

Table 1: Sample Summary Statistics

	All	Control	Treatment
Number of lists	577	389	188
Total number of items	6740	4363	2377
Average number of items per list	11.68	11.21	12.64
Total number of downloads across items	36276	23829	12447
Average number of downloads per item	5.38	5.46	5.24
Total number of papers	4985	-	-
Average number of list appearances per paper (within sample)	1.35	-	-
Average number of list appearances per paper (within NEP)	3.88	-	-

By definition, a paper may appear in both the control and treatment groups as a different item.

Table 2: Descriptive Tests of Top Position Effects

	Control	Treatment
Number of lists	389	188
Average actual downloads per top-placed item	8.65	7.47
Expected downloads per item	6.01	5.83
% Difference (actual - expected)	43.87	28.12
z test	11.28***	5.22***
Wilcoxon test	5.07***	2.13*

Table 3: Summary of List-Specific Control Variables

Name	Description	(across items)		(across lists)		Min	Max
		Mean	St. Dev.	Mean	St. Dev.		
n	The number (no.) of items included in the list	17.72	16.18	11.68	8.41	2	118
subs	The no. of subscribers to the list's subfield at release date	575	428	526	415	59	2080
av	The no. of days for which the list was available	769	43	770	44	695	841
subfields	A set of fixed effects for the subfields	-	-	-	-	-	-

Table 4: Summary of Item-Specific Control Variables

Name	Description	(across items)		(across lists)		Min	Max
		Mean	St. Dev.	Mean	St. Dev.		
engtitle	Whether the item has an English title (1=yes)	0.99	0.09	0.99	0.10	0	1
title	No. of characters in item's title (divided by 100)	0.76	0.28	0.75	0.28	0.05	2.43
authors	No. of authors	2.16	1.09	2.16	1.10	1	15
zeroab	Whether the item has no abstract (1=no abstract)	0.02	0.15	0.02	0.15	0	1
abstract	No. of characters in item's abstract (divided by 100)	9.82	5.30	9.80	5.52	0	148
zerokey	Whether or not the item has no keywords (1=no keywords)	0.19	0.39	0.20	0.40	0	1
keywords	Total no. of characters in item's keywords (divided by 100)	0.64	0.48	0.63	0.49	0	4.65
zerojel	Whether the item has no JEL codes (1=no JEL codes)	0.40	0.49	0.42	0.49	0	1
jel	No. of JEL codes provided for the item	1.91	1.92	1.85	1.91	0	13
repstotal	No. of lists within NEP in which the paper appears	4.21	1.51	3.88	1.39	2	12

Table 5: Estimation Results of Top Position Effects

Sample:	Control			Treatment			All		
Model Spec:	1(i)	1(ii)	1(iii)	1(i)	1(ii)	1(iii)	2(i)	2(ii)	2(iii)
top	2.593*** (9.99)	2.714*** (10.14)	3.156*** (9.11)	1.566*** (4.76)	1.578*** (4.53)	2.184*** (5.10)	2.497*** (9.77)	2.632*** (9.97)	3.103*** (9.08)
top*treat	-	-	-	-	-	-	-0.562* (2.09)	-0.611* (2.21)	-0.557 (1.57)
treat	-	-	-	-	-	-	0.202 (1.87)	0.127 (1.14)	0.013 (0.10)
engtitle	2.366*** (8.49)	-	2.660*** (8.05)	2.468*** (5.52)	-	2.756*** (4.90)	2.366*** (9.98)	-	2.679*** (9.41)
title	-3.867*** (4.28)	-	-5.617*** (5.02)	-3.851*** (3.39)	-	-5.827*** (3.84)	-3.951*** (5.50)	-	-5.748*** (6.31)
title ²	1.363** (2.76)	-	1.717** (2.82)	1.288* (2.13)	-	2.083* (2.51)	1.384*** (3.58)	-	1.888*** (3.79)
authors	-0.216*** (3.54)	-	-0.251** (3.18)	-0.210** (2.70)	-	-0.222* (2.27)	-0.223*** (4.56)	-	-0.245*** (3.91)
zeroab	1.787** (3.15)	-	0.942 (1.38)	1.898* (2.41)	-	0.942 (0.92)	1.804*** (3.93)	-	0.925 (1.63)
abstract	-0.028* (2.06)	-	-0.044* (2.50)	0.001 (0.10)	-	-0.03 (1.12)	-0.018 (1.59)	-	-0.040** (2.60)
zerokey	0.964** (3.06)	-	1.001* (2.50)	0.232 (0.64)	-	1.034 (1.96)	0.774** (3.19)	-	1.034** (3.21)
keywords	0.404 (1.20)	-	-0.281 (0.64)	-0.929 (1.93)	-	-0.919 (1.51)	0.068 (0.25)	-	-0.487 (1.35)
keywords ²	-0.061 (0.69)	-	0.075 (0.64)	0.474** (2.81)	-	0.273 (1.49)	0.042 (0.52)	-	0.133 (1.35)
zerojel	-0.317 (1.33)	-	-0.649* (2.13)	-0.585* (2.00)	-	-0.69 (1.80)	-0.427* (2.27)	-	-0.693** (2.87)
jel	-0.096 (1.66)	-	-0.093 (1.25)	-0.066 (0.92)	-	-0.029 (0.31)	-0.09 (1.96)	-	-0.076 (1.26)
repstotal	0.003 (0.07)	-	-0.014 (0.25)	0.165** (3.05)	-	0.159* (2.37)	0.056 (1.64)	-	0.045 (1.02)
n	-0.042*** (9.08)	-0.045*** (9.61)	-	-0.044*** (3.43)	-0.052*** (3.89)	-	-0.042*** (9.41)	-0.047*** (10.28)	-
subs	0.007* (2.37)	0.011*** (3.58)	-	0.008** (3.26)	0.010*** (4.21)	-	0.008*** (4.20)	0.011*** (5.52)	-
av	-0.004* (2.18)	-0.002 (0.88)	-	0.002 (0.46)	0.006 (1.41)	-	-0.003 (1.84)	-0.001 (0.49)	-
subfields	Yes	Yes	No	Yes	Yes	No	Yes	Yes	No
Obs	4363	4363	4363	2377	2377	2377	6740	6740	6740
Lists	389	389	389	188	188	188	577	577	577
Clusters	3361	3361	3361	1907	1907	1907	4985	4985	4985
χ^2	2672***	2377***	372***	1404***	1184***	167***	3749***	3331***	512***
AIC	25051	25625	29743	12944	13256	15477	38066	38930	45238
$\hat{\sigma}^2$	3.02	3.30	4.38	2.47	2.70	3.83	2.84	3.10	4.19

Table 6: Other Position Effects - Estimation and Test Results

Sample:	Control			Treatment			All		
Model Spec:	3(i)	3(ii)	3(iii)	3(i)	3(ii)	3(iii)	4(i)	4(ii)	4(iii)
top	2.952*** (10.73)	3.055*** (10.64)	3.393*** (9.64)	1.761*** (4.98)	1.818*** (4.86)	2.371*** (5.21)	2.841*** (10.60)	2.976*** (10.50)	3.329*** (9.56)
top*treat	-	-	-	-	-	-	-0.601* (2.22)	-0.619* (2.18)	-0.571 (1.59)
sec	1.673*** (6.57)	1.651*** (6.28)	2.048*** (6.10)	0.489 (1.69)	0.593 (1.87)	0.899* (2.19)	1.611*** (6.52)	1.590*** (6.18)	2.006*** (6.05)
sec*treat	-	-	-	-	-	-	-0.739** (2.63)	-0.657* (2.14)	-0.785* (2.08)
secbot	-0.239 (1.18)	-0.242 (1.16)	-0.190 (0.71)	-0.169 (0.65)	-0.125 (0.45)	0.114 (0.32)	-0.273 (1.36)	-0.276 (1.33)	-0.201 (0.76)
secbot*treat	-	-	-	-	-	-	0.206 (0.57)	0.279 (0.72)	0.327 (0.68)
bot	0.795** (2.70)	0.804* (2.53)	1.077** (3.01)	0.941*** (3.56)	1.040*** (3.66)	1.447*** (3.81)	0.742* (2.56)	0.755* (2.43)	1.040** (2.93)
bot*treat	-	-	-	-	-	-	0.318 (0.88)	0.394 (1.01)	0.394 (0.84)
treat	-	-	-	-	-	-	0.250* (2.06)	0.150 (1.22)	0.061 (0.40)
Item Controls	Yes	No	Yes	Yes	No	Yes	Yes	No	Yes
List Controls	Yes	Yes	No	Yes	Yes	No	Yes	Yes	No
Obs	4271	4271	4271	2360	2360	2360	6631	6631	6631
Lists	351	351	351	181	181	181	532	532	532
Clusters	3311	3311	3311	1900	1900	1900	4937	4937	4937
χ^2	2707***	2449***	400***	1398***	1195***	187***	3763***	3410***	556***
AIC	24084	24652	28517	12819	13119	15269	36991	37845	43809
$\hat{\sigma}^2$	2.88	3.15	4.14	2.47	2.69	3.81	2.76	3.01	4.02
Wald Tests:									
Top-Sec	14.05***	15.77***	8.63**	9.18**	7.26**	6.26*	-	-	-
Sec-Secbot	42.54***	39.59***	31.89***	3.42	3.34	2.38	-	-	-
Secbot-Bot	(-)9.85**	(-)8.97**	(-)8.85**	(-)10.33**	(-)9.60**	(-)7.05**	-	-	-

Table 7: The Effect of List Length on the Size of Position Effects

Sample:	Control			Treatment		
Model Spec:	5(i)	5(ii)	5(iii)	5(i)	5(ii)	5(iii)
top	2.229*** (5.87)	2.411*** (6.02)	1.874*** (4.02)	0.138 (0.29)	-0.006 (0.01)	0.085 (0.14)
top*n	0.039* (2.51)	0.034* (2.12)	0.060** (3.00)	0.105*** (3.89)	0.120*** (4.00)	0.119** (3.24)
sec	1.103** (2.73)	1.043* (2.57)	0.719 (1.53)	0.473 (0.62)	0.594 (0.70)	0.614 (0.61)
sec*n	0.036 (1.46)	0.038 (1.55)	0.061* (2.13)	-0.001 (0.03)	-0.003 (0.06)	-0.010 (0.16)
secbot	0.182 (0.36)	0.134 (0.26)	-0.334 (0.70)	0.367 (0.53)	0.538 (0.69)	0.391 (0.46)
secbot*n	-0.038 (1.11)	-0.035 (0.95)	-0.016 (0.49)	-0.047 (1.04)	-0.056 (1.17)	-0.054 (0.99)
bot	0.718 (1.77)	0.676 (1.57)	0.347 (0.79)	1.098 (1.89)	1.238* (2.08)	1.148 (1.47)
bot*n	0.004 (0.14)	0.007 (0.26)	0.027 (0.97)	-0.014 (0.39)	-0.017 (0.50)	-0.011 (0.26)
n	-0.038*** (8.28)	-0.040*** (8.74)	-0.047*** (10.76)	-0.045*** (3.39)	-0.052*** (3.84)	-0.077*** (5.79)
Item Controls	Yes	No	Yes	Yes	No	Yes
List Controls	Yes	Yes	No	Yes	Yes	No
Obs	4271	4271	4271	2360	2360	2360
Lists	351	351	351	181	181	181
Clusters	3311	3311	3311	1900	1900	1900
χ^2	2721***	2460***	576***	1428***	1215***	230***
AIC	24067	24639	28215	12789	13077	15087
$\hat{\sigma}^2$	2.87	3.15	4.04	2.45	2.66	3.67

Table 8: Estimated Position Effects for Short Lists (with 4-11 items)

Sample:	Control			Treatment			All		
Model Spec:	3(i)	3(ii)	3(iii)	3(i)	3(ii)	3(iii)	4(i)	4(ii)	4(iii)
top	2.445*** (6.28)	2.633*** (6.27)	2.921*** (5.52)	0.841 (1.61)	0.702 (1.28)	1.223 (1.67)	2.418*** (6.14)	2.658*** (6.26)	3.026*** (5.48)
top*treat	-	-	-	-	-	-	-1.147** (2.77)	-1.323** (3.16)	-1.453** (2.61)
sec	1.336*** (4.05)	1.266*** (3.69)	1.620*** (3.34)	0.156 (0.32)	0.312 (0.54)	0.512 (0.63)	1.307*** (3.85)	1.244*** (3.51)	1.701*** (3.37)
sec*treat	-	-	-	-	-	-	-0.883 (1.92)	-0.700 (1.31)	-1.076 (1.59)
secbot	-0.173 (0.55)	-0.210 (0.65)	-0.201 (0.46)	-0.424 (0.95)	-0.179 (0.34)	-0.265 (0.40)	-0.259 (0.80)	-0.282 (0.85)	-0.202 (0.45)
secbot*treat	-	-	-	-	-	-	-0.055 (0.10)	0.240 (0.37)	-0.104 (0.13)
bot	0.384 (1.26)	0.497 (1.44)	0.630 (1.41)	0.568 (1.48)	0.985* (2.10)	0.956 (1.51)	0.338 (1.05)	0.450 (1.24)	0.674 (1.46)
bot*treat	-	-	-	-	-	-	0.360 (0.74)	0.554 (0.99)	0.192 (0.28)
treat	-	-	-	-	-	-	0.937** (3.02)	0.800* (2.47)	1.224** (2.87)
engttitle	2.425*** (5.76)	-	3.042*** (4.95)	2.374** (2.68)	-	3.339** (3.24)	2.488*** (6.32)	-	3.185*** (6.20)
title	-0.348** (3.10)	-	-0.211 (1.31)	-4.855* (2.16)	-	-7.182* (2.19)	-3.816** (2.89)	-	-4.529* (2.30)
title ²	-1.864 (1.09)	-	-1.828 (0.78)	1.645 (1.38)	-	2.573 (1.42)	1.166 (1.55)	-	1.034 (0.86)
authors	-0.080 (0.08)	-	-0.717 (0.49)	-0.462* (2.54)	-	-0.299 (1.12)	-0.395*** (4.01)	-	-0.253 (1.86)
zeroab	0.804 (0.90)	-	0.267 (0.24)	4.424* (2.37)	-	2.333 (0.88)	1.748* (2.07)	-	0.794 (0.72)
abstract	-0.053* (2.06)	-	-0.068 (1.86)	-0.011 (0.30)	-	-0.065 (1.17)	-0.037 (1.77)	-	-0.069* (2.22)
zerokey	0.973 (1.90)	-	1.028 (1.54)	0.348 (0.49)	-	1.601 (1.36)	0.816 (1.91)	-	1.214* (2.02)
keywords	0.226 (0.41)	-	-0.077 (0.10)	-1.960* (2.29)	-	-0.498 (0.39)	-0.288 (0.60)	-	-0.210 (0.32)
keywords ²	0.015 (0.11)	-	0.051 (0.26)	0.842*** (4.26)	-	0.226 (0.78)	0.163 (1.50)	-	0.097 (0.61)
zerojel	-0.478 (1.28)	-	-0.543 (1.08)	-1.365** (2.60)	-	-1.543* (2.08)	-0.742* (2.35)	-	-0.824 (1.91)
jel	-0.238** (2.58)	-	-0.289* (2.31)	-0.177 (1.40)	-	-0.231 (1.22)	-0.220** (2.88)	-	-0.262* (2.43)
repstotal	0.050 (0.74)	-	-0.041 (0.43)	0.221 (1.95)	-	0.115 (0.77)	0.086 (1.46)	-	-0.001 (0.02)
n	-0.068 (1.12)	-0.077 (1.20)	-	-0.222* (2.35)	-0.186 (1.83)	-	-0.153** (3.14)	-0.147** (2.82)	-
subs	0.020* (2.04)	0.024* (2.36)	-	0.004 (0.48)	0.014 (1.68)	-	0.013 (1.92)	0.017** (2.59)	-
av	0.001 (0.34)	0.004 (1.05)	-	-0.006 (0.63)	0.002 (0.23)	-	-0.002 (0.51)	0.002 (0.55)	-
subfields	Yes	No	Yes	Yes	No	Yes	Yes	No	Yes
Obs	1471	1471	1471	633	633	633	2104	2104	2104
Lists	199	199	199	83	83	83	282	282	282
Clusters	1295	1295	1295	572	572	572	1814	1814	1814
χ^2	1407***	1236***	126***	695***	523***	51***	1846***	1582***	167***
AIC	8320	8588	10540	3728	3939	4814	12175	12613	15375
$\hat{\sigma}^2$	2.68	2.97	4.60	2.62	3.14	4.76	2.74	3.07	4.66
Wald Tests									
Top-Sec	6.76**	9.52**	4.59*	1.21	0.30	0.53	-	-	-
Sec-Secbot	14.14***	12.92***	9.98**	1.08	0.50	0.77	-	-	-
Secbot-Bot	(-)2.06	(-)2.85	(-)2.24	(-)3.84*	(-)3.39	(-)2.20	-	-	-

Table 9: Estimated Position Effects for Long Lists (with 12+ items)

Sample:	Control			Treatment			All		
Model Spec:	3(i)	3(ii)	3(iii)	3(i)	3(ii)	3(iii)	4(i)	4(ii)	4(iii)
top	3.338*** (8.67)	3.396*** (8.24)	3.776*** (7.96)	2.433*** (5.33)	2.575*** (5.21)	2.739*** (4.45)	3.249*** (8.58)	3.321*** (8.16)	3.666*** (7.94)
top*treat	-	-	-	-	-	-	-0.464 (1.32)	-0.379 (0.99)	-0.492 (1.08)
sec	1.946*** (5.07)	1.979*** (4.83)	2.372*** (4.90)	0.485 (1.55)	0.555 (1.67)	0.541 (1.43)	1.909*** (5.13)	1.941*** (4.90)	2.287*** (4.86)
sec*treat	-	-	-	-	-	-	-1.016** (3.26)	-0.986** (2.98)	-1.188** (3.18)
secbot	-0.460 (1.85)	-0.435 (1.65)	-0.439 (1.32)	-0.368 (1.39)	-0.400 (1.50)	-0.378 (1.07)	-0.436 (1.79)	-0.413 (1.61)	-0.424 (1.31)
secbot*treat	-	-	-	-	-	-	0.042 (0.10)	-0.035 (0.09)	0.025 (0.05)
bot	1.088* (2.06)	1.039 (1.86)	1.395* (2.22)	0.872* (2.41)	0.820* (2.25)	1.052* (2.20)	1.067* (2.06)	1.025 (1.90)	1.325* (2.16)
bot*treat	-	-	-	-	-	-	-0.118 (0.23)	-0.155 (0.29)	-0.151 (0.24)
treat	-	-	-	-	-	-	0.074 (0.58)	-0.002 (0.01)	-0.178 (1.19)
engtitle	2.161*** (5.61)	-	2.304*** (5.43)	3.005*** (10.46)	-	3.266*** (11.03)	2.305*** (7.38)	-	2.489*** (7.43)
title	-4.232*** (4.09)	-	-6.420*** (5.32)	-3.217** (2.66)	-	-4.545** (2.98)	-3.819*** (4.74)	-	-5.710*** (6.01)
title ²	1.485** (2.77)	-	2.065*** (3.32)	1.073 (1.64)	-	1.589 (1.87)	1.307** (3.11)	-	1.868*** (3.74)
authors	-0.151* (2.24)	-	-0.231** (2.79)	-0.130 (1.77)	-	-0.117 (1.31)	-0.145** (2.82)	-	-0.192** (3.12)
zeroab	2.278*** (3.42)	-	1.436 (1.79)	0.682 (1.51)	-	0.099 (0.14)	1.706*** (3.71)	-	0.962 (1.70)
abstract	-0.011 (0.73)	-	-0.021 (1.22)	0.005 (0.42)	-	-0.017 (0.66)	-0.004 (0.42)	-	-0.02 (1.31)
zerokey	0.789* (2.15)	-	0.701 (1.57)	0.288 (0.67)	-	0.833 (1.41)	0.605* (2.22)	-	0.755* (2.21)
keywords	0.223 (0.53)	-	-0.668 (1.29)	-0.059 (0.09)	-	-0.759 (0.85)	0.095 (0.28)	-	-0.713 (1.71)
keywords ²	0.005 (0.04)	-	0.211 (1.35)	0.066 (0.22)	-	0.154 (0.41)	0.029 (0.23)	-	0.204 (1.47)
zerojel	-0.136 (0.43)	-	-0.438 (1.18)	-0.063 (0.18)	-	-0.282 (0.69)	-0.092 (0.39)	-	-0.392 (1.41)
jel	-0.001 (0.02)	-	0.017 (0.18)	0.022 (0.26)	-	0.053 (0.53)	0.009 (0.15)	-	0.026 (0.37)
repstotal	-0.027 (0.51)	-	-0.001 (0.01)	0.139* (2.40)	-	0.137 (1.94)	0.031 (0.79)	-	0.052 (1.11)
n	-0.034*** (7.63)	-0.036*** (7.96)	-	-0.014 (1.10)	-0.018 (1.39)	-	-0.029*** (7.11)	-0.032*** (7.65)	-
subs	0.008* (2.48)	0.011*** (3.49)	-	0.009*** (3.48)	0.010*** (3.76)	-	0.009*** (4.34)	0.011*** (5.14)	-
av	-0.002 (0.79)	0.000 (0.01)	-	0.000 (0.08)	0.004 (0.91)	-	-0.001 (0.59)	0.000 (0.12)	-
subfields	Yes	No	Yes	Yes	No	Yes	Yes	No	Yes
Obs	2800	2800	2800	1727	1727	1727	4527	4527	4527
Lists	152	152	152	98	98	98	250	250	250
Clusters	2362	2362	2362	1484	1484	1484	3692	3692	3692
χ^2	1580***	1412***	297***	879***	768***	154***	2270***	2035***	413***
AIC	15606	15945	17901	8932	9039	10313	24620	25045	28230
$\hat{\sigma}^2$	2.90	3.18	3.85	2.30	2.41	3.28	2.71	2.90	3.64
Wald Tests:									
Top-Sec	6.87**	6.34*	4.43*	14.58***	13.43***	11.10***	-	-	-
Sec-Secbot	33.56***	30.03***	26.46***	4.88*	5.68*	3.45	-	-	-
Secbot-Bot	(-)8.59**	(-)6.98**	(-)7.86**	(-)8.86**	(-)8.53**	(-)6.52*	-	-	-

Figure 1: Average Per-Item Downloads by Position (Compared to Average Across All Positions)




Figure 2: Estimated Position Effects (Marginal Effects from Main Specification)




Figure 3: Estimated Position Effects in Short and Long Lists Respectively

