

QUANTIFYING SOME OF THE IMPACTS OF ECONOMICS BLOGS*

David McKenzie and Berk Özler

Abstract

Economics blogs represent a significant change in the way research on development economics is discussed and disseminated; yet little is known about the impact of this new medium. Using surveys of development researchers and practitioners, along with experimental and non-experimental techniques, we try to quantify some of their effects. We find that links from blogs cause a striking increase in the number of abstract views and downloads of economics papers. Furthermore, blogging raises the profile of the blogger and changes readers' perceptions about his or her institution. Finally, we find some suggestive evidence that a blog can increase knowledge of the topics it covers for the average, but not the marginal, reader.

Keywords: Blog; Dissemination; Influence; Impact Evaluation.

JEL codes: A11, A23, O12, C93.

* The authors thank Irem Yoruk, Matthew Groh and Gustavo Andrade for excellent research assistance; Jed Friedman and Markus Goldstein for their collaboration on the Development Impact blog; Innovations for Poverty Action for their assistance in conducting the survey; Jean Roth of the National Bureau of Economic Research for providing us with data on downloads of NBER working papers; the editor and the associate editor at EDCC, two anonymous referees, Richard Baldwin, Austin Frakt, Steven Levitt, Greg Mankiw, Dani Rodrik, and Justin Wolfers for comments and suggestions; all the individuals who took time to participate in this survey; and all the individuals who commented on a series of posts on this paper on "Development Impact." The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors. They do not necessarily represent the views of the International Bank for Reconstruction and Development or the World Bank. David McKenzie, The World Bank, 1818 H St., NW, Washington, DC, 20433. Phone: 202-458-9332; email: dmckenzie@worldbank.org. Berk Özler, The World Bank, 1818 H St., NW, Washington, DC, 20433. Phone: 202-458-5861; email: bozler@worldbank.org.

1. Introduction

Practically nonexistent a decade ago, economics blogs are now part of a cultural change in the way research is discussed, disseminated, and even produced.¹ General economics blogs, such as Freakonomics, Marginal Revolution, Paul Krugman and Greg Mankiw, have built large followings – whether measured by subscriptions in Google Reader or by average daily page views.² Blogs have also become increasingly prevalent in the field of development economics, with several academics writing blogs and development organizations like the World Bank, Inter-American Development Bank, Center for Global Development, Oxfam, and Innovations for Poverty Action all maintaining one or more active blogs.

Cowen (2008) argues that blogs are the main way that the general public consumes economics in a given day and suggests that “...about 400,000 people are reading economics blogs and digesting them” on a daily basis. These blogs not only give their creators an outlet to disseminate their ideas and work immediately in a format that is more accessible, but also enable instant feedback, are easy to share online, and allow the bloggers a personal style rather than the inaccessible format of academic journals (Glenn, 2003; Dunleavy and Gilson 2011).

Our motivation in examining the impact of economics blogs stems from two observations about blogs and the questions that arise from these. First, it seems fair to state that “...informing is the core business of blogging.” (McKenna and Pole 2008, p. 102) This leads to the question of whether blogs improve the dissemination of research findings and whether their readers are indeed more informed. On the one hand, the large following scholar bloggers have garnered offers the promise that economics blogs may have sizeable effects on the dissemination of economic research and on the knowledge and attitudes of their readers.³ On the other hand, blogs may fail to cause a significant change in knowledge and attitudes if, as Sunstein (2008) argues, they cause “group polarization” and create “information cocoons.”⁴ Combined with the

¹ On how debates with other economists through blogs can lead to the production of new research, see Gans (2012).

² As of July 2011, Paul Krugman had more than 56,000 subscribers in Google Reader, Mankiw more than 40,000, Freakonomics almost 7,000, and Marginal Revolution more than 4,000. Average daily page views for Marginal Revolution was 36,000 and for Mankiw 13,500 in June 2011 according to SiteMeter.

³ In this sense, economics blogs can serve a similar function to traditional media. For example, Phillips et al. (1991), taking advantage of a natural experiment, shows that articles in the New England Journal of Medicine that were covered by the New York Times received 73% more citations than control articles during the first year after publication.

⁴ Although there is some evidence that polarizing traditional media can affect voter behavior. For example, DellaVigna and Kaplan (2007) find that the introduction of Fox News had an impact on voter turnout, as well as the vote share in Presidential elections between 1996 and 2000.

possibility that blogging gives scholars the freedom to write about topics outside their area of expertise, this would suggest that impacts of blogs might be negligible.⁵

Second, while blogging (and reading blogs) exacts a non-negligible opportunity cost, revealed preference suggests that there is value in blogs to both the scholars and institutions who maintain them (and to the large number of individuals who read them): blogs provide an outlet for ideas and observations not suitable for other media; they allow the scholar blogger to sharpen her arguments by having to defend them publicly; they may lead to regular writing gigs or other professional opportunities, etc. Furthermore, blogging by individual scholars may have positive spillover benefits for their institutions. However, there is, to date, no quantitative evidence of such impacts of economics blogs.⁶

This paper aims to answer three questions regarding the impacts of economics blogs. First, do blogs improve dissemination of working papers or journal articles? Second, do they raise the profile of their creators and affiliated institutions? Third, do they cause changes in attitudes among their readers and/or lead to increased knowledge? We conduct event study analysis using download data from the Research Papers in Economics (RePEc) database to answer the first question. To tackle the remaining questions, we use evidence from a recent survey we conducted among researchers and practitioners of economic development on the role of blogs, and take advantage of a randomized experiment in which a random sample of the respondents of this survey were encouraged to read a new development economics blog.

We find that links to scholarly articles (either publications or working papers) in blogs lead to substantial jumps in their likelihood of being downloaded, with the impact increasing with the popularity of the blog providing the link. Our experiment provided random encouragement to read a new blog produced by four researchers from the World Bank.⁷ Using this experiment, we find that the encouragement increased recognition of the bloggers behind this new blog; led to an improvement in the perceived quality of research produced at the World Bank; caused a reduction in perceived censorship; and sparked an increased interest by the

⁵ Blogs can also potentially leave their readership less informed. Gelman and Fung (2012) discuss the presence of avoidable errors in the case of Freakonomics and suggest that mistakes can arise for a variety of reasons and defeat the information purpose of blogs if the “Readers must work to discern which conclusions are fully quantitative, which are somewhat data driven and which are purely speculative.” Ecker et al. (2011) discuss the relative ineffectiveness of retractions in eliminating misinformation.

⁶ For a nice theoretical exposition of the economics of blogging, particularly the costs and benefits from the perspectives of both the individual blogger and the society see Ribstein (2006).

⁷ McKenzie and Özler are two of the four researchers who produce the “Development Impact” blog.

survey respondents in working at the World Bank as a researcher. Furthermore, combining experimental and non-experimental evidence for this sample, we find some evidence to suggest that reading the new blog improved awareness of findings from recent studies and made them more likely to change how they feel about the effectiveness of a particular intervention for the average reader, but not the marginal reader who only read the blog due to encouragement. Hence, our paper provides the first quantitative evidence on some effects of economics blogs – both for the scholar bloggers and their readers – and presents evidence of spillovers for institutions affiliated with these blogs.

The rest of the paper is as follows. Section 2 presents the impacts of blogs on dissemination of economics research using event study analysis. Section 3 uses our new survey data to describe how blogs are used by graduate students in economics, junior faculty, and field workers in development. Section 4 presents impacts of reading a new blog, exploiting the randomized encouragement design. Section 5 concludes.

2. Event Study Analysis of Dissemination Impact

One of the main purposes of economics blogs is to help better disseminate economic ideas and research: both to other economists and to the broader public. The impact of some of this effort is very hard to measure – for example, many general economics blogs have devoted considerable space to discussing public policy issues in the U.S., while a number of development blogs have discussed issues relating to the debate on aid effectiveness and institutional governance. It is difficult to assess how much any one particular blog may have influenced these debates. However, one area of research dissemination that is potentially important and can be measured is whether blogging about a research paper leads to more people viewing and downloading that paper.

2.1 Descriptive Figures

Several economics blogs regularly link to working papers. However, two issues arise in trying to measure the impact of these links on downloads. First, many of these links are to the web pages of the individual authors or to working paper series for which download statistics are not publicly available. To resolve this, we focus our analysis on blog posts which link directly to papers in RePEc. RePEc is the largest database of economics papers, containing over 430,000 working papers and 675,000 journal articles. In 2011 there were over 28 million abstract views

and over 8 million downloads of papers from this site. Monthly abstract views and download statistics are publicly available. A second issue arises for links provided from blogs to academic papers when they are first released as working papers. It is harder to form a proper counterfactual in these cases with respect to impact on dissemination, since there are often several avenues of dissemination when papers are first released which might also drive download statistics, and heterogeneity in topics amongst papers would rule out comparison to other papers in the same series or papers by the same author as appropriate counterfactuals.

We therefore focus on blog links to papers in RePEc, which have been out at least several months at the time of a blog posting. Figure 1 provides a particularly striking illustration of the phenomenon we wish to measure. Irwin (1997) received an average of 3.4 abstract views and 0.8 downloads per month from the NBER working paper series during 2009. Then on February 16, 2010, Paul Krugman blogged about the paper on his “Conscience of a Liberal” blog, resulting in 940 abstract views and 151 downloads in February 2010.⁸ The paper then went back to averaging 5.3 abstract views and 0.8 downloads per month from April 2010 through March 2011.⁹

Similar patterns occur for other blogs. Figure 2 gives the example of Landry et al. (2006), which was averaging 14.4 abstract views and 5.2 downloads per month in the year before Freakonomics blogged about the paper, and then had 1,478 abstract views and 144 downloads in the month the link was provided from Freakonomics. Figure 3 shows abstract views and download statistics for Arai and Thoursie (2006), which averaged only 1.5 abstract views and 0.67 downloads per month in the year before Chris Blattman blogged about the paper, then had 57 abstract views and 11 downloads in the month the paper was linked from his blog.

2.2 Formal Estimation

We systematically searched the 50 most read finance and economics blogs (defined according to one list based on blogs with public traffic logs¹⁰) for links to research papers. Out of this list, only 12 were blogs written by academic economists – the remainder consisted of mostly financial or macro blogs, typically without reference to academic papers. To these 12 blogs, we

⁸ The paper is in the NBER working paper series, which is gated (requiring payment) to readers from institutions which do not subscribe to the NBER series, limiting downloads.

⁹ This pattern is similar to what is known as the “Slashdot effect” (http://en.wikipedia.org/wiki/Slashdot_effect) which is used to describe the spike in web traffic that occurs when a popular website links to a smaller site. However, our work is the first we are aware of to show that this also occurs for downloads of research papers following links from blogs.

¹⁰ <http://www.gongol.com/lists/bizeconsites/> [March 2011 data, accessed August 2011].

added the six other blogs that our survey respondents interested in development economics (described below) were most likely to report reading (Aid Watch, Chris Blattman, Economix, Freakonomics, Paul Krugman, and Dani Rodrik).¹¹

Examining these 18 academic blogs for a period of four months in 2011, we find that the mean (median) number of links to research papers is 8.3 (6) per month. The distribution is highly skewed, with Marginal Revolution and Freakonomics both linking to more than 100 papers over this four month period, whereas seven other blogs (including Greg Mankiw and Dani Rodrik) linked to 10 or fewer papers. Just over half (53%) of the links are to newly released working papers or journal articles (within three months of being released), typically going directly to the working paper series, journal website, or the author's personal webpage. As mentioned above, we exclude links to recently released papers from our analysis due to the fact that identification of the impact of blogs on their dissemination is problematic.

Out of remaining links to papers that were released at least three months before the date of the link, half of them are directly to author's web pages, with another 22% to JSTOR or to journal websites. The NBER working paper series (10%) and RePEc (6%) together account for another 16% with the remainder distributed across a range of sources, including SSRN and other working paper series. Out of these sources, RePEc is the only one with publicly available abstract view and download data, so we use these data for our main analysis. Nevertheless, as we describe below, we also carry out robustness analysis using download data obtained privately from the NBER.

In our main analysis, we use data on abstract views and downloads for papers for which there was a direct link from one of these 18 blogs to the RePEc database, and which had been published at least three months before the link appeared. We excluded blogs which only linked to one or two working papers at most.¹² This resulted in a database of 107 research papers linked from one of eight blogs during a period of four and a half years spanning January 2007 through June 2011: Aid Watch, Baseline Scenario, Chris Blattman, Economix, Freakonomics, Marginal

¹¹ Using other sources of data on the daily page views for these six additional blogs, we find that each of them would appear in the list of top 50 business and economics websites if their traffic levels were also publicly available in the same format as used by gongol.com.

¹² There also seem to be spikes for these other blogs. For example, the blog Angry Bear linked to a paper which had abstract views jump from 3 to 150, and downloads from 0 to 40 comparing the month before to month of blogging.

Revolution, Overcoming Bias, and Paul Krugman.¹³ We use this database to formally test for whether blogging about a paper increases its abstract views and downloads through event study analysis.

The RePEc statistics are available at a monthly frequency, and so for each paper i that is blogged about, we define $t=0$ in the month in which the blog entry occurred, $t=-1$ in the month before, $t=+1$ in the month after, etc. We then estimate the impact of blog s linking to a paper using the following regression:

$$Abstract\ Views_{i,t} = \alpha_i + \sum_{s=1}^8 \beta_s Blog_{i,t} + \sum_{s=1}^8 \gamma_s Blog_{i,t-1} + \sum_{s=1}^8 \delta_s Blog_{i,t+1} + \varepsilon_{i,t} \quad (1)$$

$Blog_{i,t}$ is a dummy variable which takes value one if the paper is blogged about in time period t , β_s is our coefficient of interest, measuring the increase in abstract views in the month of blog s blogging compared to the paper-specific average, and γ_s allows for a one-month lagged effect which may arise particularly for blog posts towards the end of the month. We include paper-specific fixed effects (α_i), Error terms are clustered at the paper level, which allows for serial dependence and assumes independence across papers. The corresponding equation is likewise estimated for paper downloads.

Equation (1) is known as the constant mean model in event study analysis (Campbell et al. 1997). A first threat to this assumption is if abstract views or downloads are trending over time. Paper view statistics appear to trend downwards over the first couple of months of release of the typical paper, but otherwise seem reasonably stable. Excluding access statistics for papers which are blogged about during the first two months after publication should therefore alleviate this concern. Nevertheless, for robustness we also re-estimate equation (1) after adding paper-specific linear time trends.

A second concern is that of reverse causation, with bloggers blogging about a paper because people have suddenly started downloading it and talking about it. The inclusion of the lead term $Blog_{i,t+1}$ allows us to test whether $\delta = 0$, and thereby rule out the case that a paper which attracts a lot of attention in month t gets blogged about in month $t+1$. A related concern is

¹³ There is a perception that papers with flashy findings or that are “Freakonomics-like” in applying economic analysis to issues that appeal to a broader audience are more likely to be discussed in blogs. While this might be the case for papers that are discussed and linked immediately upon release, the papers in our sample are blogged about several months or years after their initial release dates and are comprised of a broad range of macro, labor and development papers. It also seems clear that pure theory or econometrics papers are much less likely to be blogged about.

that a particular paper attracts a lot of attention for some unrelated reason in month t , resulting in a simultaneous increase in interest in the paper and in blog entries about the paper. If this were the case, we should see the same paper being blogged about on multiple blogs. This is likely to be an issue when looking at papers which are just released, but is much less of a concern for older papers. There are only two occasions where this occurred in our sample. The first is multiple blog entries pointing to Mortensen and Pissarides (1994) in October 2010, when they were awarded the Nobel Prize. The second case is Rockey (2009), which was blogged about by Marginal Revolution on June 26, 2010, and then picked up (with acknowledgement to Marginal Revolution) in a Freakonomics blog post on July 8, 2010.¹⁴ We exclude the first case, and code the second case as having been blogged about in both June and July 2010.

In our baseline specification we estimate equation (1) using monthly data up to 2 years on either side of the month of the link from a blog. We then examine robustness by narrowing the window to ± 1 year and to ± 6 months.

2.3. Results

Table 1 shows the results of estimating equation (1). We see large and significant impacts of blogging on both paper abstract views and paper downloads in the month in which the paper is blogged about. There are also some significant, but smaller, impacts on these access statistics in the month after the paper is blogged about. The lead terms are all small, and in all but one case, insignificant.¹⁵ These results are consistent with the graphical illustrations seen in Figures 1-3, and show a big spike in abstract views and download of papers in the month they are discussed in one of these blogs.

To place the impacts in perspective, it is useful to first compare them to the download and abstract numbers for an average NBER working paper from RePEc: 10.3 abstract views per month and 4.2 downloads per month in months 3-14 after release. A blog post on Chris Blattman or Aid Watch is thus equivalent to an extra 7-9 months of abstract views, and 4-6 months of

¹⁴ This custom in blogging to acknowledge when the blogger has come across a paper while reading another blog (via the syntax “hat tip”, “heard through”, or H/T in short for either) also makes us confident that these papers weren’t linked from our sample of blogs as a result of the bloggers reading about these papers on other blogs not included in our sample (such as blogs by non-scholars that occasionally reference research papers, like Ezra Klein, Andrew Sullivan, or and Megan McArdle). Again this would be more of an issue for recently released papers. In addition, we found no cases of the New York Times, Washington Post or CNN linking to the studies included in our paper, suggesting that coverage in the mainstream media is unlikely to be driving our results.

¹⁵ The exception is on downloads for Freakonomics, and arises from the case mentioned above, in which it blogged about a paper the month after Marginal Revolution had. Excluding this paper reduces the Freakonomics lead term download coefficient to 2.5 with $p=0.17$ in column 5.

downloads. The impacts of Freakonomics, Marginal Revolution and Paul Krugman are even larger – with the abstract view impact of 300-470 equivalent to 3 or more years of regular views, and the download impact of 33-100 downloads equivalent to 8 months to 2 years of regular downloads.

Exact and consistent data across all the blogs in our list are not available, but the data that are available suggest that the most-read blogs have significantly lower click-through rates than the more research-focused niche blogs. Marginal Revolution and Freakonomics are both estimated to have approximately 35,000-40,000 page views and 25,000 unique visits per day. This suggests a click-through rate of only 1-2 percent for abstract reads and 0.1-0.4 percent for downloads. Baseline Scenario has 6,800 page views and 5,000 visits per day, and Overcoming Bias 4,000 page views and 2,700 visits. This suggests a click-through rate of 3-4 percent for abstract views and 0.7 percent for downloads. Chris Blattman's blog is estimated to have approximately 2,200 page views per day, suggesting a click-through rate of 4.3 percent for abstract reads and 1.1 percent for downloads.¹⁶ This seems consistent with the intuition that as an academic's blog starts reaching a larger audience, the additional readers are less likely to be interested in reading academic papers.

We note that the estimates in Table 1 show the average impacts of links from these blogs. In practice, there appears to be considerable heterogeneity in the spike in blog traffic caused by a particular blog. For example, just taking the difference in abstract views in the month of the link compared to the mean abstract views over the months before the paper was discussed gives a range of +33 to +2908 over the 31 papers linked by Marginal Revolution in our sample (25th percentile to 75th percentile range is 69 to 314). It is likely that the size of the increase reflects a combination of the interest in the topic to the general blog reader, and the manner in which the blog links to the paper (e.g. full post about a paper vs. a single line link; positive, neutral or negative link, etc.). Unfortunately the data in our sample are not suitable to explore this phenomenon systematically.

Finally, we do not believe there is any reason to think that the impacts of linking to papers in RePEc are likely to be different than linking to papers in a working paper series or in an author's own webpage. Nonetheless, as a robustness check, we requested access statistics

¹⁶ Blog traffic statistics from <http://www.gongol.com/lists/bizeconsites/> [accessed March 2011 rankings on July 28, 2011]; Marginal Revolution and Freakonomics traffic data from SiteMeter; and www.websitevalue.us. Chris Blattman traffic statistics based on a blog posting in which he said he had 800,000 page views in 2010.

from the NBER and SSRN. SSRN does not provide monthly data for abstract views or downloads. The NBER noted their data are not normally set up to be reported in this manner and don't capture abstract views, but, in response to our request, agreed to extract paper download data for a small number of papers. We requested monthly download data for each NBER paper that had been published at least six months before a link was provided from a blog post in Marginal Revolution, Paul Krugman, Baseline Scenario or Freakonomics between July, 2009 and April, 2011.¹⁷ Appendix Table 1 estimates equation (1) on these data and shows similar impacts to those in Table 1 – a link by Paul Krugman leads to an increase of 235 downloads, while links from Marginal Revolution and Baseline Scenario cause an increase of 43-46 downloads.¹⁸ The confidence intervals for these point estimates overlap with those for the impacts of these same blogs in Table 1.

3. Survey Data from Development Economists

In order to explore in more detail how potential readers use blogs, in February 2011 we conducted surveys of several groups who would be potential readers of development economics blogs. Development economics is a useful area to examine for several reasons. First, it is a field in which a mixture of academic and less academic potential audiences can be identified. Secondly, since only about the half the readers of development and aid blogs live in the United States, it opens up the possibility of considering readers in a range of different countries.¹⁹ Finally, from a practical standpoint, our own contacts and work in this field made it easy to identify potential survey participants and is likely to have increased response rates.

3.1 Sample frames and response rates

The first group we identified consists of students in Ph.D. and Master's Programs in economics who are studying development economics. We contacted development economics faculty at 48 institutions in the U.S. and abroad and asked them to forward an invitation to their graduate students to participate in our online survey. This faculty list was comprised of members of the organization BREAD who teach development at a school with a Ph.D. or Master's

¹⁷ We thank Jean Roth of the NBER for her help in providing these data. We chose these blogs (and Chris Blattman's, which didn't link to an NBER paper published at least six months earlier during this period) to examine because they are the ones with the highest impact in our RePEc sample.

¹⁸ Freakonomics is omitted since there is only one NBER paper linked to during this time.

¹⁹ This is according to the Smart Aid survey (<http://findwhatworks.wordpress.com/2011/09/19/blog-survey-findings-2-demographics-and-professional-status/>).

Program, as well as additional faculty whom, through personal contacts, we knew to teach development at this level. Students were told that the purpose of the survey was to assess how the next generation of development economists finds out about new studies in development economics, and about the role of social media such as blogs. They were told that the first 100 individuals to reply would receive a copy of one of two new popular press books on development, as would a random drawing of other respondents. A total of 405 Ph.D. students and 181 Masters students not in Ph.D. programs completed the survey. Faculty were asked to tell us how many students they had sent the invitation to, and based on these responses, we estimate that the survey response rate was at least 60 percent of those who received invitations, which is high for an online survey. Although we cannot say whether those who didn't participate in the survey are less likely to read economics blogs, a comparison of students who responded fastest to those who took more time to reply does not show any significant difference in likelihood of reading economics blogs, suggesting that marginal respondents are not those who are less interested in reading economics blogs.²⁰

The second group surveyed were field staff for Innovations for Poverty Action (IPA), an NGO which implements randomized experiments in a number of countries around the world; and fellows of the Overseas Development Institute (ODI). These ODI fellows are young postgraduate economists who are sent to work on two-year contracts in the public sectors in selected developing countries. This group therefore represents individuals who are more involved in the practice of development work. They were given the same incentives to respond to the survey as the student group. A total of 150 field staff replied to the survey, representing a response rate of approximately 60-70 percent.

The third group surveyed was assistant professors in development economics. These were identified through their membership in the organization BREAD or through their participation in the NEUDC development economics conference. This group comprises of potential readers most engaged in research among our survey populations. Invitations were sent to 120 individuals, with 76 taking part in the survey (63 percent).

The final group surveyed was individuals with the job title "Economist" at the World Bank who were not in the research department. New Ph.D.s are hired as Economists and typically

²⁰ 83 percent of the fastest quintile of responders read economics blogs, versus 77 percent of the slowest quintile (p-value for a test of equality is 0.34). 16 percent of the fastest quintile of responders read them daily, versus 14 percent for the slowest quintile (p=0.71).

spend 6-8 years in this position before getting promoted to a different job title. This provides a group of economists engaged in operational development work without a research focus. Survey invitations were sent to 170 staff, but responses were only received from 43 individuals (25 percent).²¹

The top panel of Table 2 provides some basic summary information for these different groups. Average ages range from 27 to 34, and women range from 42 to 54 percent of the different groups. Approximately two-thirds of the assistant professors and Ph.D. students are located in the United States, as are half the Masters students. Most of the field staff is located in developing countries, but 20 percent list their location as the U.S., either because they are U.S.-based staff for IPA, or because they are temporarily in the U.S. before heading to a field office.

This age range overlaps well with the most common ages for blog readers according to other recent surveys. The SmartAid blog carried out a survey of readers of different development and aid blogs in 2011 and found 63 percent of survey respondents were aged 20 to 34 and 51 percent were female²², while Davis et al. (2011) find in their survey of professors aged 59 on average that only 44 percent do read blogs. Despite concerns with response rates and unknown representativeness with both of these other surveys, it appears that blog readership is much more common among younger professionals. This pattern also appears to be true of the general public, with a 2009 consumer survey finding that 16.1 percent of 25-34 year olds had read a blog in the past 30 days, compared to 8.6 percent of 45-54 year olds and 6.4 percent of 55-64 year olds.²³

The baseline survey asked about 12 working papers released in 2010 on the BREAD working papers website, a leading source of working papers in development economics. Even with self-reports of having read the paper, and counting having seen a seminar on the paper as having read it, the majority of survey respondents have not read most papers. The assistant professors in development have on average only read 2.2 out of the 12 papers, and 22 percent have read none. PhD students have read only 1.4 of these papers on average, and field staff and World Bank economists outside the research department less than 1 of the 12 papers. Given these

²¹ We readily admit that this sample is not a random sample of the entire “audience of interest,” which would be also comprised of policymakers, staff of aid organizations and other bilateral and multilateral donor agencies, as well as a random sample of students, faculty, and practitioners of development economics. To the extent that our sample differs from this population, the external validity of our findings would be compromised.

²² Source: <http://findwhatworks.wordpress.com/2011/09/19/blog-survey-findings-2-demographics-and-professional-status/> [accessed 16 January 2012].

²³ Source: <http://www.marketingcharts.com/interactive/bloggers-few-in-number-wield-disproportionate-influence-10799/mediamark-research-mri-visited-wrote-blog-30-days-october-2009jpg/> [accessed 3 February, 2012].

low levels of reading working papers, there certainly seems to be a role for other forms of dissemination about new results.

3.2 Self-reported uses of economics blogs in survey data

The second panel of Table 2 provides more detailed information on how individuals read and use blogs. Readership of economics blogs is high among all five groups surveyed, with between 76 and 84 percent of those surveyed having read an economics blog in the past six months. Female graduate students are significantly less likely to read blogs than males ($p < 0.02$), although there is no significant difference in blog readership by gender among assistant professors, field staff, and World Bank economists. Among those who don't read economics blogs, the most common reasons given as the most important for not reading them were lack of time, and lack of knowledge about what economics blogs are out there.

Many of those who read blogs do so only sporadically – only 40 percent of graduate students and 34 percent of assistant professors who read blogs do so at least a few times a week. The majority read blog posts by entering the URL address in their browser, rather than email subscriptions, an RSS feed such as Google Reader, or another social media channel like Facebook. Perhaps as a result, the average reader does not read very many blogs – the median and mean number of economics blogs read in the past month is about 3. Table 2 shows the most read blogs among this sample at the time of the survey are a couple of mainstream blogs – Marginal Revolution and Freakonomics – and four development-focused blogs – Chris Blattman, Aid Watch (now defunct), Dani Rodrik and the IPA blog.

The last part of Table 2 shows what actions blog readers say they have taken in the past month as a result of reading economics blogs. Consistent with the evidence from Section 2, a majority (between 50 and 72 percent depending on subgroup) of readers say they have read a new economics paper as a result of a blog posting. Although one-third of graduate students and one half of assistant professors are currently involved in designing and implementing a survey, only 2 percent of students and 4 percent of assistant professors say they have added a survey question in the past month as a result of a blog posting. Likewise, the percentage of blog readers who have changed how they plan to analyze data as a result of a blog posting is also low, ranging from 6 to 11 percent.

Finally, the survey results suggest that blogs are having an influence on how people feel about the effectiveness of particular policies, particularly among more policy-oriented

respondents. Thus 44 percent of field staff who read blogs and 34 percent of World Bank operational economists say that, in the past month, they have changed their views about the effectiveness of some policy intervention as a result of a blog post. This is also the case for 29 percent of Ph.D. student readers, but only 10 percent of assistant professors who read blogs.

4. A Randomized Experiment

In order to measure the causal impact of blogging, we would ideally randomly assign some people to read a blog, and others to not. This is difficult to do for many blogs in existence, since most potential readers would have heard of the blog and potentially sampled it to see whether they are interested or not. We therefore took advantage of the launch of a new blog, Development Impact²⁴, which went online on April 1 2011, and provided randomized encouragement to read this blog.

Development Impact is a blog initiated by the authors and two other World Bank researchers (Jed Friedman and Markus Goldstein), and is hosted by the World Bank. It covers a range of issues related to impact evaluations and evaluative research, including discussions of new research papers, reviews of new books about impact evaluations, methodological issues, and experiences from evaluations around the world. During the months of May and June of 2011, it received a total of just under 50,000 page views, or just over 800 page views per day. While small relative to the most-read mainstream economics blogs like Freakonomics and Marginal Revolution, this traffic level would place it in the top 50 of economics and business blogs according to the ranking used in section 2. Moreover, this makes it the most read World Bank blog during this period, and perhaps a reasonable representative of blogs in economic development: it has about 40 percent of the traffic of Chris Blattman's blog, one of the most popular development blogs, exceeds that of Innovation for Poverty Action's blog, and has more Google Reader subscribers than the Center for Global Development's Views from the Center blog.²⁵ As such, while the findings of the experiment are internally valid only for the Development Impact blog, they seem likely to extend more generally to other research-oriented development blogs.

²⁴ <http://blogs.worldbank.org/impactevaluations/>

²⁵ Chris Blattman traffic statistics based on a blog posting in which he said he had 800,000 page views in 2010.

4.1 Randomized Encouragement

We used the 619 graduate student and field staff respondents to our baseline survey that had provided us with email addresses as the sample for this experiment.²⁶ We stratified the data by three variables that could potentially influence attitudes towards research methods and interest in the World Bank: whether the respondent was a Ph.D. student, field staff, or Masters student; whether or not the respondent indicated reading Chris Blattman’s blog at baseline; and whether or not the respondent reported being involved in a randomized experiment at baseline. Individuals were then randomly assigned to treatment and control within each stratum, with Appendix Table 2 showing that this succeeded in balancing baseline observable characteristics.

The encouragement then consisted of two emails to those in the treatment arm. The first was sent on April 6th, five days after the blog was launched, thanking this group for participating in the survey and alerting them to the new blog. They were told about the purpose of the blog and some of the topics that would be discussed, as well as saying “We consider you our most important audience for such a blog, and so want to make it something interactive and useful. We therefore very much would like if you check out the new blog, comment on things, and raise any issues or thoughts on things you would like the blog to discuss.” We then sent a second, shorter, email to the same group three weeks after the blog had been launched, which asked how the blog was doing so far, and asked whether there were particular topics people would like to see covered, again encouraging people to check out the blog. In this sense, the encouragement was stronger than a simple statement informing people of the existence of a new blog, but did not involve any financial incentives to take-up the treatment offered.

4.2 Follow-up Survey

A follow-up survey was then sent at the start of June 2011 to both treatment and control groups, i.e. two months after the blog was launched. There were two main reasons for this time frame. First, we were concerned that some of the sample would graduate or change email addresses over the summer, making it hard to track them over a longer period. Second, given the quite rapid increase in readership experienced over the first two months of the blog and links to Development Impact from other blogs such as Marginal Revolution, IPA, and Chris Blattman,

²⁶ We did not use the World Bank or Assistant Professor samples because of the smaller size of these groups, and because the World Bank economists would be automatically notified about a new World Bank blog.

we were concerned that the control group might rapidly become readers of the blog, thereby reducing the encouragement effect.

The follow-up survey was answered by 445 of the 619 initial respondents (72 percent), which is a high response rate for an online survey. The response rate was slightly higher in the control group than in the treatment group (74.9 percent vs. 68.8 percent, $p=0.092$). A comparison of those who responded quickly to those who responded after multiple attempts shows no significant differences in age, gender, location, or interest in becoming an academic researcher. However, those who required more time and effort to reply were less likely to be frequent blog readers at baseline. Appendix Table 2 however shows that the treatment and control groups are still balanced on baseline characteristics among the follow-up sample, so that any selective attrition according to observable characteristics is not significantly related to treatment status.

Neither the treatment nor the control group was told that they were in an experiment, and both had been surveyed previously on similar topics. It therefore seems unlikely that any results obtained are the result of Hawthorne effects.

4.3 Did the Encouragement Work?

We estimate the following linear regression equation to test whether the randomized encouragement succeeded in increasing readership of the Development Impact blog among the treated group:

$$\text{Read Development Impact}_i = \alpha_s + \beta \text{Treat}_i + \varepsilon_i \quad (2)$$

where the α_s are controls for randomization strata (Bruhn and McKenzie, 2009) and the coefficient of interest is β .

The first column of Table 3 shows that the encouragement succeeded in increasing the proportion of respondents who read Development Impact by 9.9 percentage points, a 55 percent increase relative to the 18 percent of the control group who had read Development Impact in the last month before the follow-up survey. Columns 2 and 3 then re-estimate equation (1) by gender, and columns 4 and 5 by whether or not the individual stated at baseline that they wished to become a researcher in an academic institution. We see that the encouragement worked for males but not females, and for individuals who wish to become academic researchers but not others. It seems reasonable that encouragement to read a research-oriented blog is likely to work

better for individuals who are more interested in pursuing a research career. Some of the gender difference is due to females being less likely to say they want an academic research career, but the encouragement treatment also has no effect on females who want research careers. The lack of effect for females could potentially also be related to some of the reasons hypothesized for why few female economists blog (Kahn, 2011).

4.4 Impacts of the “Development Impact” Blog

We now use our follow-up survey data to estimate the impact of reading Development Impact on various outcomes. We can do this by using the randomized encouragement as an instrument for reading Development Impact in the following regression:

$$Outcome_i = \alpha_s + \gamma Read\ Development\ Impact_i + \omega_i \quad (3)$$

Randomized encouragement designs have a long history of being used to assess the impacts of different media, ranging from early evaluations of Sesame Street (Ball and Bogatz, 1970) to more recent evaluations of radio programs (Paluck, 2009). The parameter γ that is identified through this design is the local average treatment effect or LATE (Angrist et al, 1996), which in our case, is the impact of reading Development Impact for individuals who read it when encouraged and do not read it otherwise. This group consists of about half the male and just over half of the research-focused individuals in our sample, so it is a non-trivial group (Table 3). Moreover, this is the parameter of interest for answering questions like should blogs attempt outreach exercises to have more readers.

Nevertheless, if the marginal readers, who only read the blog because of encouragement, are those who find it less interesting or read it less intensively than those read it of their own accord, the average impact of reading the blog may differ from the LATE. We therefore also employ the bias-adjusted nearest-neighbor matching estimator of Abadie et al. (2004) to estimate the average treatment effect on the treated (ATT). We match on a rich set of baseline characteristics: age, gender, whether or not the individuals attend an elite (top 5) economics department²⁷, whether they live in the U.S., whether they are Ph.D. students, Master’s students or field staff, whether they plan on careers as academic researchers, whether they are currently

²⁷ We define the “top 5” as Harvard, MIT, Chicago, Princeton, Stanford, Berkeley and Yale – i.e. 7 schools who all are sometimes considered as “top 5” depending on which ranking criteria is used. In addition, Berkeley and Yale, which are tied for 6th in the U.S. News and World Report rankings, have very strong programs in development economics.

involved in conducting surveys, whether they read economic blogs, the frequency of reading, dummies for readership of the four most commonly read blogs in our survey (Chris Blattman, Aid Watch, Marginal Revolution, and Freakonomics), and the number of current research papers (out of 12) they had read at baseline. The identifying assumption is that, conditional on this large number of measures of interest in research, blog reading habits, and other characteristics, there is no selection into reading Development Impact on the basis of unobserved characteristics. This assumption may be more likely to hold in our context, whereby potential readers are coming across the blog by chance during the first two months of the blog, than might be the case for well established blogs. Nevertheless, identification remains more of a concern for these estimates than for those based on the experimental encouragement design.

Blogs may affect the reputation of the individuals and institutions producing them. In Table 4, we examine whether reading the Development Impact blog changes readers opinions about the World Bank and their awareness of its producers. Survey respondents were asked to rate on a 10 point scale (10 being the highest) their interest in working as a researcher at the World Bank, and at other institutions. Columns 1 and 2 show the ITT and IV-LATE using the encouragement experiment on the full sample, while column 3 shows the ATT estimated using matching. Columns 4 through 7 then examine the experimental effects for the two groups which responded most to the encouragement: males and individuals who say they would like to become academic researchers (research-focused for short). Point estimates are positive for all specifications, and are significant at the 10 percent level for the experimental estimates on the research-focused group. For this research-focused group, there is also a marginally significant impact on interest in working in a top-10 research university, and a strong negative effect on working in a liberal arts university.

The second panel of Table 4 examines the perceptions of the quality of research produced by different institutions, again rated on a 10-point scale with 10 being the highest. Reading Development Impact is found to have large and statistically significant impacts on the perceptions of quality of research produced at the World Bank in all specifications. Blog readership seems to have spillover results on the reputation of the IMF's research, despite no posts directly relating to work by IMF researchers (perhaps because readers revise upwards their opinion of the quality of work at international institutions) and also on the quality of work from Harvard, Yale, and MIT – schools strongly associated with rigorous impact evaluation work in

development economics and whose work featured strongly in the first two months of posts.²⁸ In contrast, there is no significant impact on the perceived quality of research at a range of other good, but not at the very top, economics programs whose work was not blogged about very much during this first two months.²⁹ Taken together these results therefore show that, over a very short period, reading the Development Impact blog has made readers view the World Bank more favorably both as a place to work and as a producer of good research.

The four core bloggers on Development Impact post their blogs without going through any approval process, and the blogs are written in a conversational tone – perhaps more so than other World Bank blogs at the time. There is an impression that World Bank researchers face some restrictions and censorship on what they can write - as evidenced by the control group mean of 3.4 out of 5 on a scale of 1 = high degree of censorship and 5 = complete freedom. Both the ATT from matching and the experimental results for the research-focused group show a positive impact of blog readership on this score, indicating readers of the blog are less likely to view researchers at the World Bank as censored in terms of what they can write in a blog post.

Finally, we asked the survey respondents about their awareness of a list of nine development researchers, all approximately 5-10 years post-Ph.D. Included in this list were the two other regular bloggers on the Development Impact blog (since the survey was administered by the authors of this manuscript, we didn't ask whether people had heard of us or our work). The final panel of Table 4 shows some evidence for greater awareness of the Development Impact bloggers. The experimental impact is positive and significant at the 10 percent level for males, but insignificant for the research-focused sub-sample. The ATT is also significant at the 10 percent level. In contrast, we do not find any greater awareness of seven reasonably

²⁸ Development Impact had 16 posts linking to research from someone from Harvard, MIT or Yale in its first two months, including feature reviews on books by Dean Karlan and by Esther Duflo and Abhijit Banerjee.

²⁹ We pool these schools together in these two groups for ease of presentation, and because Harvard/Yale/MIT are typically the top ranked three programs in development economics among the schools we asked about. They are also home to JPAL and IPA, which are strongly associated with randomized experiments. If we run the analysis separately for each school, the ITT effect is positive and significant at the 5% level for Harvard and Yale, positive but insignificant ($p=0.17$) for MIT, and insignificant at standard levels ($p>0.27$ in each case) for all other schools. The choice of other schools was made to give a range of other types of programs in development, without requiring participants to answer survey questions about every top-25 program.

comparable development economists, none of whose research was discussed in our blog during this period.³⁰

Table 4 therefore shows a number of significant and positive changes in attitudes and general impressions towards the World Bank and its researchers as a result of the Development Impact blog. The mere existence of the blog and a casual reading of articles to get a sense of the issues being discussed may have been sufficient to generate these findings. In Table 5, we look for changes in knowledge and attitudes which might only be expected to occur from more in-depth reading of blog posts. To measure knowledge, we asked detailed questions related to six blog posts that had appeared on the Development Impact blog (Appendix 2 provides the exact questions). These questions were quite specific and proved difficult for the respondents, with the mean individual in the control group only getting 0.91 out of 6 correct. The experimental impacts estimated on the full sample and on the sub-groups vary in sign and are not significant. However, the matching estimate is positive, large relative to the mean, and significant at the 1 percent level.³¹ Our point estimate of a 0.8 question increase in knowledge is equivalent to 0.7 standard deviations for the control group. This is large relative to increases in knowledge found in many educational interventions, which often find increases of 0.1-0.3 standard deviations in learning (e.g. Angrist et al, 2002; Kremer, 2003).³²

When we examine the survey respondents' answers to the question of whether they changed their mind about the effectiveness of an intervention during the past month, we get results that are qualitatively identical: the average readers of Development Impact, but not the marginal ones, are substantially more likely to have changed their feelings about the effectiveness of an intervention (a 16 percentage point increase over a control mean of 25%).

Three possible interpretations for this difference between the ATT and IV-LATE suggest themselves. The first is that the power for detecting experimental effects is less than for the matching estimator due to incomplete take-up, so these differences just arise due to statistical variability. Indeed the IV-LATE for the impact on knowledge for the full sample in Table 5,

³⁰ The seven other researchers were Kathleen Beegle, Jishnu Das, Pascaline Dupas, Eric Edmonds, Erica Field, Joseph Kaboski, and Asim Khwaja. They were chosen as people at approximately the same career stage as our co-bloggers and to cover people both within and outside of the World Bank.

³¹ We do not know the effect of reading our blog on time use – it could be crowding out other productive learning activities. Hence, we cannot claim welfare gains from the improved test scores for the average reader.

³² A school-building program that was exceptionally successful in increasing school enrollment (by 42 percentage points) in Afghanistan led to a 1.2 SD increase in test scores for those enrolled in school (Burde and Linden, 2012).

while negative, has a wide confidence interval that would include the positive and significant point estimate found with the matching estimator.

A second explanation is that reading the blog impacts knowledge for the average reader, but not for the marginal reader who only reads because of encouragement. The IV-LATE estimate will in general only be equal to the ATT estimate if either the treatment effect is constant, or if no one in the control group reads the blog (Angrist, 2004). The latter does not hold in our case, and it is plausible that there is treatment heterogeneity since the readers who would read the blog regardless of whether they are encouraged or not might be the ones most likely to read closely, learn from it, and change their minds about interventions.

A third possible reason for this difference is that the matching estimate might just show there is positive selection on knowledge into blog readership. However, recall that among the variables used for matching are an indicator of attending an elite economics PhD program, interest in being an academic researcher, the number of recent papers out of 12 read at baseline, and indicators of baseline blog reading habits. Therefore we are controlling for a large number of characteristics that might well proxy for knowledge.

To lend additional support to interpreting the matching estimates presented in Tables 4 and 5 as causal rather than due to selection, we conduct two robustness exercises which are reported in Appendix Tables 3 and 4. The first of these excludes from our analysis baseline readers of Chris Blattman and AidWatch – presumably the two most similar audiences to that of Development Impact – and then conduct our matching analysis on the impact of reading Development Impact for individuals who did not report reading these other two development blogs at baseline. The results are qualitatively similar to those for our full sample of matching, although there is some inevitable loss of statistical power due to the smaller sample size.

Second, we conduct a falsification test, in which we match on whether or not the survey respondents read Chris Blattman’s blog at baseline instead of Development Impact, again chosen because it is likely the closest substitute. Reading this alternative blog has no significant impact on interest in working at the World Bank or perceptions of quality of research produced by the World Bank. In fact, it has exactly the opposite effect on a number of outcome variables, including perceptions of censorship at the World Bank. While it does not improve awareness of our co-bloggers at Development Impact, it does improve awareness of other researchers included in our survey (many of whose research has been mentioned on Chris Blattman’s blog). Appendix

Table 4 shows that reading Chris Blattman’s blog neither increases knowledge of the six papers covered in *Development Impact* between the baseline and follow-up surveys nor changes feelings about effectiveness of an intervention during the past month. We view these robustness checks as suggesting the matching estimates may indeed be picking up the causal impact of reading the *Development Impact* blog for the average reader.

We also examined whether blog readership is affecting attitudes towards different methodologies. There has been a lot of recent debate about the role of experiments in development economics, with some critics such as Deaton (2010) claiming that experiments have no special role to produce more credible knowledge than any other method, and others such as Ravallion (2009) worrying that development researchers are letting methodology drive the questions they answer. Our survey results among the control group find little agreement with Deaton, but that many share this particular concern of Ravallion. We find no consistent impacts on a number of questions concerning attitudes towards methodological issues, which may reflect lack of power or the short-time frame of our follow-up survey.³³

Before concluding the paper, we present some descriptive evidence on what development economics blogs discuss. In the analysis above, we have shown that (a) economics blogs lead to a higher number of abstract views and downloads of papers to which they provide links; (b) they can change readers’ perceptions about the bloggers and the institutions with which they are affiliated; and (c) they can lead to modest increases in knowledge and beliefs about effectiveness of interventions. However, if bloggers mainly discuss their own work, or the work of others from their own institutions, then the effects could be merely self-promotional rather than advancing knowledge by disseminating new research and ideas in development economics. If the benefits of blogging accrue mainly to self (or the sponsoring institution), then it is harder to claim positive spillover effects for the rest.

To shed some light on this question, we have examined every post from 11 development blogs for the six-week period between September 1 and October 15, 2012. These included popular blogs from larger development institutions, such as Inter-American Development Bank’s *Development that Works*, Center for Global Development’s *Views from the Center*, World Bank’s *Development Impact* and *Let’s Talk Development*, as well as Oxford University’s *CSAE Blog*, *Chris Blattman.com*, Acemoglu and Robinson’s *Why Nations Fail*, among others. We then

³³ Results, not shown, are available from the authors upon request.

categorized each post into one or more of the following four categories according to the contents of the entry: (C1) blog post discusses blogger's own research paper; (C2) blog post discusses blogger's own ideas on research, policy, impact evaluation, etc.; (C3) blog post discusses a research paper by someone else from the same institution; and (C4) blog post discusses a research paper by someone else from another institution.³⁴

Our findings are reported in Table 6. We find that, of the 255 posts in these 11 blogs during this six-week window, only 36 entries discuss the blogger's own research paper (14%), and another 29 mention research papers by someone else from the same institution (11%) – indicating that approximately three quarters of all blog entries contained no mention of research by the blogger or someone else from the blogger's institution.³⁵ In contrast, 171 blog posts (67%) in this sample contain some discussion of a research paper by someone from another institution, including a non-negligible number of posts that contain a substantive discussion of a research paper rather than a simple link to the paper or a 'cut and paste' of an abstract.

These results suggest that while there is a significant element of self-promotion and the promotion of affiliated individuals and institutions that is present in blogging, bloggers also disseminate information about new research, ideas, and issues in their fields of expertise. They cause increases in downloads of research papers of other researchers in their field, and may cause modest changes in the knowledge and attitudes of their readers on issues that are unrelated to the blogger or her institution.

5. Conclusions

Economic blogs are doing more than providing a new source of procrastination for writers and readers. Using experimental and non-experimental approaches, we have provided the first quantitative evidence that they are having some effects. We observe large impacts on dissemination of research – a link on a popular blog results in a substantial increase in abstract views and downloads. In future work, it would be interesting to see whether blogging about research papers increases citations. Testing this would require more time to pass given the long

³⁴ Please note that because a blog post could be categorized into multiple categories, the percentages of posts in each category can add up to more than 100%.

³⁵ These figures increase to 22% and 18%, respectively if we exclude *Udadisi* from the list, which accounts for a large number of posts in our sample (102 out of 255). Similarly, 78 out of 153 total posts (or 51%) contain a discussion of research papers by someone from another institution when *Udadisi* is excluded from our analysis.

lag from reading about a paper to citing it, as well as different empirical methods than those used to test dissemination here as event study analysis is not suited for this type of analysis.³⁶

We also have evidence from our experiment that blogs influence perceptions and attitudes towards development policy: readers of the new Development Impact blog think more highly of World Bank research and are more interested in working in this institution, while the average reader may have gained more knowledge about the contents of recent research papers and changed her feelings about the effectiveness of interventions as a result of reading this blog.

When it comes to making a judgment about the overall effects of economics blogs in general or development blogs in particular, we advise caution. In our study we did not measure effects of reading blogs on time use, therefore we cannot make statements regarding overall increases in knowledge. Furthermore, blogs can also spread misinformation fast and leave their readers less, rather than more, informed, especially if their readership trusts the bloggers to sift through important details of research and summarize it for them (Gelman and Fung, 2012). And, while we have shown some positive externalities of blogs, they can also have negative ones. A poorly written blog can have implications for the day-to-day operations of organizations and cause reputational harm before they can be taken down. Given that mainly successful blogs survive and unsuccessful ones likely disappear, our evidence, with the exception of the experiment with Development Impact during its first two months, comes from a sample of blogs that are in equilibrium. This makes it difficult for us to estimate what would happen if new blogs were encouraged to start.

Nonetheless, our results show that there are a number of positive externalities from economics blogs that are unlikely to be captured by the blogger herself. The presence of these externalities, coupled with positive costs of blogging, suggests that there may be an undersupply of good economics blogs. Organizations and academic institutions, which are curious but unsure about the benefits of having more of their staff blog in their field of expertise, may provide some incentives for blogging, closely track the impacts on institution-specific outcomes of interest, and keep the subsidies for successful blogs and discontinue them for others.

³⁶ Ideally one would have several blogs randomly choose among a set of papers to blog about, and then follow this up several years later to see whether this generates more citations of these papers. We have suggested this idea to a couple of the most popular bloggers who did not wish to conduct such an experiment.

References

- Abadie, Alberto, David Drukker, Jane Herr, and Guido Imbens. "Implementing matching estimators for average treatment effects in Stata" *Stata Journal* 4(3): 290-311, 2004.
- Arai, Mahmood and Peter Thoursie, "Giving up foreign names: An empirical examination of surname change and earnings," *Stockholm University Department of Economics Working Paper*, 2006.
- Angrist, Joshua "Treatment Effect Heterogeneity In Theory and Practice", *Economic Journal* 502: C52-C83, 2004.
- Angrist, Joshua, Guido Imbens and Donald Rubin, "Identification of causal effects using instrumental variables," *Journal of the American Statistical Association* 91 (1996): 444-455.
- Angrist, Joshua, Eric Bettinger, Erik Bloom, Elizabeth King and Michael Kremer "Vouchers for private schooling in Colombia: Evidence from a randomized natural experiment", *American Economic Review* 92(5): 1535-58, 2002.
- Ball, Samuel and Gerry Bogatz, *The First Year of Sesame Street: An Evaluation*, 1970. Educational Testing Service, Princeton, NJ.
- Bruhn, Miriam and David McKenzie, "In pursuit of balance: Randomization in practice in development field experiments," *American Economic Journal: Applied Economics* 1 (2009): 200-32.
- Burde, Dana, and Leigh L. Linden. 'The Effect of Village-Based Schools: Evidence from a Randomized Controlled Trial in Afghanistan,' *NBER Working Paper* No. 18039, 2012.
- Campbell, John, Andrew Lo and A. Craig MacKinlay, *The Econometrics of Financial Markets*, 1997. Princeton University Press. Princeton, NJ.
- Cowen, Tyler, "Why everything has changed: the recent revolution in cultural economics," *Journal of Cultural Economics* 32 (2008): 261-73.
- Davis, William, Bob Figgins, David Hedengren, and Daniel Klein, "Economics professors' favorite economic thinkers, journals, and blogs (along with party and policy views)," *Econ Journal Watch* 8 (2011): 126-46.
- Deaton, Angus, "Instruments, randomization, and learning about development," *Journal of Economic Literature*, 48 (2010): 424-55
- Dellavigna, Stefano and Ethan Kaplan, "The Fox news effect: Media bias and voting," *Quarterly Journal of Economics* 122 (2007): 1187-1234.

- Dunleavy, Patrick and Chris Gilson, "Improving Academic Communication", presentation at the Investigating Academic Impact conference, 2011:
<http://blogs.lse.ac.uk/impactofsocialsciences/presentations/> [accessed July 31, 2011].
- Ecker, Ullrich K. H., Stephan Lewandowsky, Briony Swier, and Darren Chang, "Correcting false information in memory: Manipulating the strength of misinformation encoding and its retraction," *Psychonomic Bulletin and Review*, 18 (2011): 570-78.
- Gans, Joshua, "Blogs and academic research: A timely story," January 18, 2012.
<http://www.digitopoly.org/2012/01/18/blogs-and-academic-research-a-timely-story/>
- Gelman, Andrew, and Kaiser Fung, "Freakonomics: What Went Wrong?" *American Scientist*, 100 (2012): 6.
- Glenn, Owen, "Scholars who blog," *Chronicle of Higher Education*, June 6, 2003.
- Irwin, Douglas, "Higher tariffs, lower revenues? Analyzing the fiscal aspects of the "Great tariff debate of 1888," *NBER Working Paper no. 6239*, 1997.
- Kahn, Matthew, "Where are the female economics bloggers?" *Christian Science Monitor*, March 26, 2011. <http://www.csmonitor.com/Business/Green-Economics/2011/0326/Where-are-the-female-economics-bloggers> [accessed August 2, 2011].
- Kremer, Michael "Randomized evaluations of educational programs in developing countries: Some lessons", *American Economic Review Papers and Proceedings* 93(2): 102-106, 2003
- Landry, Craig, Andreas Lange, John List, Michael Price, and Nicholas Rupp, "Toward an understanding of the economics of charity: Evidence from a field experiment," *Quarterly Journal of Economics* 121 (2006): 747-82.
- McKenna, Laura and Antoinette Pole, "What do bloggers do: an average day on a political blog," *Public Choice* 134 (2008): 97-108.
- Mortensen, Dale and Christopher Pissarides, "Job creation and job destruction in the theory of unemployment," *Review of Economic Studies* 61(1994): 397-415.
- Paluck, Elizabeth, "Reducing Intergroup Prejudice and Conflict Using the Media: A Field Experiment in Rwanda," *Journal of Personality and Social Psychology* 96 (2009): 574-87.
- Ravallion, Martin, "Should the Randomistas rule?" *The Economists Voice* 6 (2009).
- Phillips, David P., Elliot J. Kanter, Bridget Bednarczyk, and Patricia L. Tastad, "Importance of the Lay Press in the Transmission of Medical Knowledge to the Scientific Community," *The New England Journal of Medicine* 325 (1991): 1180-1183.

Ribstein, Larry, "From Bricks to Pajamas: The Law and Economics of Amateur Journalism,"
William & Mary Law Review 48 (2006): 185-249.

<http://scholarship.law.wm.edu/wmlr/vol48/iss1/3>

Rockey, James, "Who is left-wing, and who just thinks they are?" University of Leicester
Working Paper, 2009.

Sunstein, Cass, "Neither Hayek nor Habermas", *Public Choice* 134 (2008): 87-9.

Table 1: Do blogs increase abstract views and downloads of papers blogged about?

	Abstract Views				Paper Downloads			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>Month of blog post effects</i>								
Aid Watch	67.9*** (14.6)	66.8*** (14.5)	66.1*** (14.4)	65.0*** (14.3)	17.1*** (6.2)	16.5*** (6.2)	15.9** (6.1)	15.7** (6.4)
Baseline Scenario	150.1*** (31.9)	150.1*** (32.3)	149.4*** (32.3)	148.7*** (33.3)	35.1*** (7.2)	35.1*** (7.2)	35.2*** (7.3)	35.0*** (7.6)
Chris Blattman	94.7*** (28.5)	88.5*** (27.8)	94.4*** (28.4)	94.4*** (28.6)	25.3*** (8.1)	23.6*** (8.1)	24.9*** (8.0)	24.7*** (8.1)
Economix	134.0*** (37.9)	134.7*** (38.6)	138.2*** (38.5)	140.1*** (40.5)	20.0*** (5.8)	20.1*** (5.9)	20.5*** (6.0)	20.6*** (6.3)
Freakonomics	466.4** (231.5)	397.1** (180.3)	473.3* (240.3)	450.9* (230.0)	100.3 (68.8)	82.9 (53.7)	102.1 (71.3)	96.3 (68.5)
Marginal Revolution	295.2*** (83.3)	258.6*** (62.0)	296.2*** (86.4)	286.7*** (84.1)	38.8*** (12.5)	29.8*** (6.7)	39.3*** (13.2)	36.6*** (12.2)
Overcoming Bias	102.9*** (34.5)	102.9*** (35.0)	101.6*** (35.0)	101.5*** (35.9)	18.8** (7.9)	18.8** (8.0)	18.5** (7.8)	18.2** (7.9)
Paul Krugman	446.5*** (160.7)	448.6*** (163.0)	437.9*** (160.4)	425.9*** (160.5)	83.3*** (31.0)	83.9*** (31.4)	80.6*** (30.2)	76.3** (29.2)
<i>Month after blog post effects</i>								
Aid Watch	-2.5 (5.4)	-3.8 (6.4)	-4.3 (6.4)	-5.3 (6.4)	-1.6 (2.7)	-2.4 (3.2)	-2.7 (3.2)	-2.9 (2.9)
Baseline Scenario	16.8*** (3.0)	16.8*** (3.0)	16.2*** (3.1)	15.4*** (3.5)	5.6*** (1.3)	5.6*** (1.3)	5.7*** (1.4)	5.5*** (1.3)
Chris Blattman	11.2*** (3.2)	9.8*** (1.9)	10.8*** (2.7)	10.9*** (2.4)	2.6* (1.4)	2.0** (1.0)	2.1 (1.4)	2.0 (1.6)
Economix	20.3** (8.6)	20.5** (8.8)	18.9** (8.1)	17.2** (8.6)	2.8** (1.4)	2.8* (1.5)	2.5* (1.4)	2.3 (1.7)
Freakonomics	152.6 (131.7)	24.9 (22.2)	159.1 (139.3)	111.0 (103.8)	23.8 (22.1)	-8.4 (6.5)	25.5 (24.0)	13.0 (14.4)
Marginal Revolution	138.2 (92.4)	105.3 (68.7)	139.2 (96.1)	128.8 (91.9)	45.8 (31.6)	37.8 (25.8)	46.2 (32.5)	43.3 (31.5)
Overcoming Bias	11.2*** (2.5)	11.2*** (2.5)	9.9*** (1.9)	9.8*** (1.6)	2.8** (1.1)	2.8** (1.1)	2.5** (1.1)	2.2* (1.1)
Paul Krugman	111.9* (66.7)	114.2* (67.9)	103.3* (61.9)	91.3* (55.0)	28.7 (19.5)	29.3 (19.9)	26.0 (18.0)	21.7 (15.3)
<i>Month before blog post effects</i>								
Aid Watch	1.5 (4.2)	0.8 (3.8)	-0.3 (3.3)	-1.3 (3.3)	5.8 (5.9)	5.3 (5.7)	4.6 (5.4)	4.4 (5.9)
Baseline Scenario	0.7 (1.6)	0.7 (1.7)	0.0 (2.1)	-0.7 (2.5)	1.0 (0.9)	1.0 (0.9)	1.1 (1.3)	0.8 (1.4)
Chris Blattman	-0.5 (5.3)	-1.6 (5.9)	-0.8 (5.0)	-0.8 (4.9)	2.2 (3.5)	1.8 (3.8)	1.7 (3.3)	1.6 (3.1)
Economix	7.3 (8.2)	8.1 (8.6)	6.3 (8.6)	3.7 (9.1)	0.8 (1.7)	0.9 (1.8)	0.6 (1.8)	0.5 (2.0)
Freakonomics	14.6 (9.7)	10.4 (8.0)	21.9 (17.5)	28.3 (27.7)	5.0* (2.9)	4.1 (2.6)	6.9 (4.8)	8.6 (7.4)
Marginal Revolution	8.0 (6.9)	3.7 (8.3)	9.2 (10.0)	11.2 (15.1)	2.9 (2.5)	2.1 (2.9)	3.4 (3.2)	3.8 (4.2)
Overcoming Bias	1.0 (3.5)	1.0 (3.5)	-0.5 (2.7)	-1.1 (2.0)	1.0 (1.6)	1.0 (1.7)	0.6 (1.5)	0.4 (1.2)
Paul Krugman	14.7 (10.7)	17.2 (11.4)	6.1 (5.2)	-5.9* (3.2)	3.9 (4.5)	4.5 (4.7)	1.2 (3.1)	-3.1 (2.2)
Window on either side of blog date	24 months	24 months	12 months	6 months	24 months	24 months	12 months	6 months
Paper-specific linear time trend	No	Yes	No	No	No	Yes	No	No
Observations	3,841	3,841	2,295	1,310	3,841	3,841	2,295	1,310

Notes: Robust standard errors in parentheses clustered at the paper level, *, **, and *** indicate significance at 10, 5 and 1% levels respectively.

Table 2: Basic Characteristics and Blog Reading of Development Economist Survey Respondents

	PhD Students	Masters Students	Field Staff	Assistant Professors	World Bank Economists
<i>Basic Characteristics</i>					
Mean Age	28.4	26.5	27.0	32.7	34.2
Proportion Female	0.45	0.54	0.51	0.46	0.42
Proportion in the U.S.	0.68	0.54	0.20	0.65	0.78
Currently Writing a Research Paper	0.87	0.71	0.25	0.95	0.95
Currently Implementing a Survey	0.31	0.17	0.44	0.53	0.49
Currently Implementing an Experiment	0.20	0.09	0.75	0.49	0.26
Mean number of current research papers (out of 12) read	1.44	0.98	1.21	2.24	0.70
Proportion who have read 0 out of 12 recent papers	0.41	0.54	0.39	0.22	0.63
<i>Economics Blog reading characteristics</i>					
Has read an Economics Blog in last 6 months	0.76	0.76	0.84	0.79	0.78
Males	0.82	0.85	0.84	0.74	0.77
Females	0.69	0.68	0.84	0.84	0.80
Made a comment on an Economics Blog in last 6 months	0.10	0.09	0.17	0.10	0.14
<i>Conditional on reading economics blogs</i>					
Reads blog by going manually to blog webpage	0.69	0.73	0.68	0.74	0.89
Reads blogs daily or several times a week	0.40	0.39	0.55	0.34	0.31
Read Marginal Revolution in last month	0.36	0.20	0.38	0.40	0.14
Read Freakonomics in last month	0.42	0.36	0.34	0.40	0.28
Read Chris Blattman in last month	0.44	0.34	0.64	0.48	0.17
Read Aid Watch in last month	0.24	0.19	0.43	0.08	0.10
Read Dani Rodrik in last month	0.31	0.48	0.42	0.36	0.52
Read IPA blog in last month	0.21	0.36	0.68	0.18	0.07
<i>Actions taken in last month as a result of reading blogs(conditional on reading)</i>					
Read a new economics paper	0.59	0.53	0.57	0.50	0.72
Added a question to a survey questionnaire	0.02	0.02	0.06	0.04	0.07
Changed how they plan on analyzing data	0.08	0.09	0.11	0.06	0.10
Changed feelings about effectiveness of a particular intervention	0.29	0.44	0.44	0.10	0.34
Sample Size	405	181	150	76	43

Table 3: Did the Encouragement Work, and for Whom?

Dependent Variable: Read Development Impact Blog in last month

	(1)	(2)	(3)	(4)	(5)
	Full sample	Males	Females	Research Focused	Not Research Focused
Treatment	0.099*** (0.036)	0.137*** (0.048)	0.038 (0.057)	0.195*** (0.066)	0.054 (0.043)
Proportion of Control Group who read Development Impact	0.18	0.15	0.21	0.14	0.19
Observations	445	239	202	135	310

Notes: Robust standard errors in parentheses, *, **, and *** indicate significance at 10, 5 and 1% levels respectively. Research-focused denotes individuals who say at baseline they wish to become an academic researcher

Table 4: Impact of Reading Development Impact Blog on Perceptions of Institutions

	Control Group Mean	Full sample			Males		Research-focused	
		(1) ITT	(2) IV-LATE	(3) Matching	(4) ITT	(5) IV-LATE	(6) ITT	(7) IV-LATE
<i>Interest in Working as a researcher:</i>								
at World Bank	7.87	0.122 (0.213)	1.243 (2.151)	0.143 (0.316)	0.102 (0.311)	0.760 (2.236)	0.748* (0.388)	3.777* (2.087)
at IMF	5.18	0.221 (0.272)	2.008 (2.566)	-0.545 (0.471)	-0.107 (0.396)	-0.770 (2.761)	0.468 (0.508)	2.366 (2.522)
at top-10 research university	7.63	0.230 (0.233)	2.163 (2.223)	0.215 (0.402)	0.177 (0.305)	1.258 (2.101)	0.512* (0.304)	2.587 (1.671)
at Liberal Arts University	5.55	0.0169 (0.249)	0.149 (2.162)	-1.110*** (0.414)	-0.278 (0.343)	-1.810 (2.264)	-0.893** (0.420)	-4.464* (2.423)
<i>Perception of Quality of Research Produced</i>								
at World Bank	7.73	0.309** (0.156)	2.968* (1.681)	0.483* (0.261)	0.575** (0.222)	4.298** (2.043)	0.739** (0.285)	3.465** (1.487)
at IMF	6.39	0.431** (0.199)	3.987* (2.249)	0.288 (0.334)	0.537* (0.285)	3.530* (2.083)	0.737** (0.370)	3.453* (1.834)
at Harvard, Yale and MIT	8.70	0.354*** (0.124)	3.374** (1.580)	0.177 (0.154)	0.544*** (0.186)	3.867** (1.726)	0.195 (0.224)	0.930 (1.046)
at selection of other schools	6.85	0.111 (0.131)	1.087 (1.311)	0.005 (0.207)	0.258 (0.184)	1.825 (1.360)	0.113 (0.244)	0.524 (1.071)
<i>Perception of Extent to which World Bank staff face</i>								
Censorship over blog posts (1=high, 5 = low)	3.41	0.130 (0.107)	1.147 (0.921)	0.597*** (0.168)	0.208 (0.144)	1.296 (0.860)	0.537*** (0.188)	2.465*** (0.931)
<i>Awareness of Individuals</i>								
Proportion aware of 2 Development Impact bloggers	0.61	0.012 (0.041)	0.114 (0.379)	0.119* (0.069)	0.107* (0.0575)	0.837* (0.491)	-0.021 (0.074)	-0.102 (0.358)
Proportion aware of 7 comparable researchers not mentioned on Development Impact blog	0.63	-0.019 (0.033)	-0.167 (0.294)	-0.038 (0.056)	0.025 (0.046)	0.182 (0.330)	-0.052 (0.057)	-0.255 (0.282)
Sample Size		439	439	433	235	235	134	134

Notes: Robust standard errors in parentheses, *, ** and *** indicate significance at the 10%, 5% and 1% levels respectively.

Selection of other schools is average over Oxford, Paris School of Economics, Williams, Cornell, Michigan, British Columbia, and Duke.

Table 5: Impact of Reading Development Impact Blog on Reader Knowledge and Attitudes Regarding the Effectiveness of Interventions

	Control Group	Full sample			Males		Research-focused	
	Mean	(1)	(2)	(3)	(4)	(5)	(6)	(7)
		ITT	IV-LATE	Matching	ITT	IV-LATE	ITT	IV-LATE
Knowledge								
Number of questions correct about 6 papers	0.91	-0.103 (0.0982)	-1.038 (1.159)	0.815*** (0.167)	-0.0367 (0.133)	-0.267 (0.975)	0.0273 (0.183)	0.140 (0.889)
Changed how they feel about effectiveness of an intervention in the past month	0.25	-0.0515 (0.0325)	-0.417 (0.474)	0.162** (0.082)	-0.0626 (0.0466)	-0.388 (0.456)	0.0154 (0.0625)	0.0272 (0.362)
Sample size		445	445	439	239	239	135	135

Notes: Robust standard errors in parentheses, *, ** and *** indicate significance at the 10%, 5% and 1% levels respectively.

Table 6 - What Do Bloggers Discuss?

Blog Name	Number of posts	Numbers of Papers in Each Category			
		Category 1	Category 2	Category 3	Category 4
Chris Blattman (Columbia University)	16	2	7	4	8
CSAE Blog (Oxford University)	3	2	1	0	0
Development Impact (World Bank)	29	3	10	10	16
Development that Works (Inter-American Development Bank)	12	0	2	1	10
Financial Access Initiative (New York University)	11	2	4	0	9
Innovations for Poverty Action Blog	9	2	2	4	3
Let's Talk Development (World Bank)	16	6	4	5	4
Marc Bellemare (Duke University)	17	1	6	1	11
Udadisi (Andres Marroquin's blog)	102	2	6	2	93
Views from the Center (Center for Global Development)	28	7	23	2	6
Why nations fail (Daron Acemoglu & James Robinson)	12	9	6	0	11
TOTAL	255	36 (14%)	71 (28%)	29 (11%)	171 (67%)

Notes: All blog entries between September 1 and October 15, 2012 were included in the sample, excluding entries that were of a personal nature or unrelated to research. We excluded entries titled 'Links I Liked' from Chris Blattman's blog. Blog entries were categorized as follows: Category 1 - blog post discusses blogger's own research paper; Category 2 - blog post discusses blogger's own ideas on research, policy, impact evaluation, etc.; Category 3 - blog post discusses a research paper by someone else from the same institution; and Category 4 - blog post discusses a research paper by someone else from another institution. Each entry could be categorized as fitting into multiple categories, meaning that the percentage of posts in each category add up to greater than 100%. Percentages of each category is presented in parentheses in the final

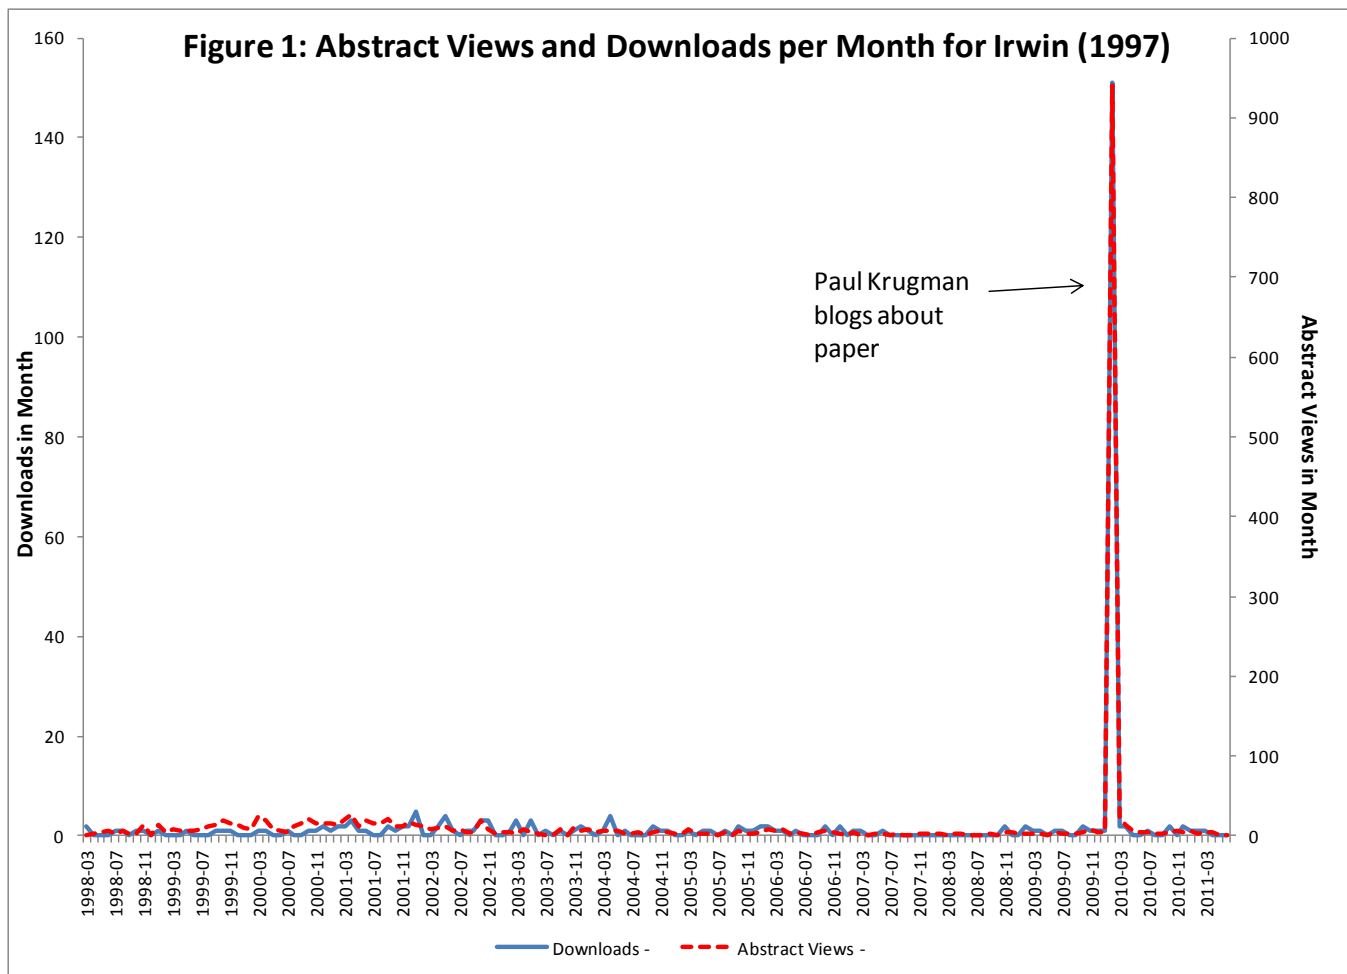


Figure 2: Abstract Views and Downloads per Month for Landry et al. (2006)


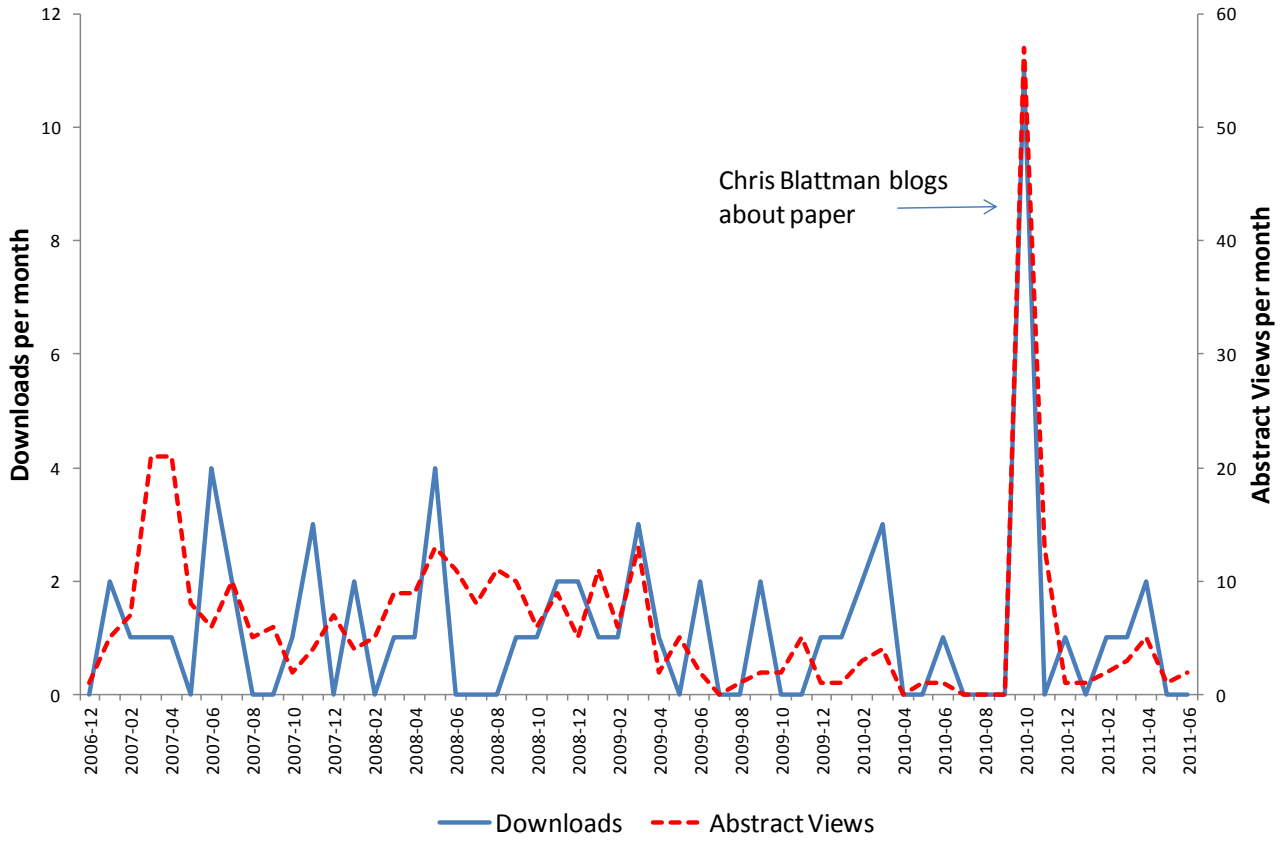


Figure 3: Abstract Views and Downloads per Month for Arai and Thoursie (2006)



Appendices (for ONLINE PUBLICATION ONLY)

Appendix 1

Appendix Table 1: Impact of Blogs on NBER downloads

	(1)	(2)
<i>Month of blog post effects</i>		
Marginal Revolution	43.9** (16.7)	44.2** (15.2)
Paul Krugman	235.1*** (67.4)	233.0*** (68.3)
Baseline Scenario	45.7*** (7.0)	43.1*** (5.5)
<i>Month after blog post effects</i>		
Marginal Revolution	7.4 (10.6)	7.7 (8.0)
Paul Krugman	11.8* (6.1)	9.8 (5.7)
Baseline Scenario	3.0 (7.1)	0.5 (6.3)
<i>Month before blog post effects</i>		
Marginal Revolution	-10.2* (4.8)	-9.9*** (2.7)
Paul Krugman	-2.2 (4.7)	-4.2 (5.2)
Baseline Scenario	1.0 (3.9)	-1.5 (2.7)
Constant	32.4*** (1.1)	32.6*** (1.7)
Window on either side of blog date	12 months 6 months	
Observations	349	195

Notes: Robust standard errors in parentheses clustered at the paper level, *, **, and *** indicate significance at 10, 5 and 1% levels respectively.

Appendix Table 2: Balance on Observables for Full Sample and Follow-up Respondents

	Full Sample Randomized			Follow-up Respondents		
	Control	Treatment	p-value	Control	Treatment	p-value
<i>Stratification Variables</i>						
Ph.D. Student	0.54	0.54	0.961	0.58	0.59	0.757
Field staff	0.21	0.20	0.970	0.21	0.18	0.421
Masters student	0.26	0.26	0.983	0.21	0.23	0.682
Reads Chris Blattman's blog	0.39	0.38	0.944	0.41	0.44	0.571
Currently conducting experiment	0.32	0.31	0.790	0.31	0.33	0.704
<i>Other Variables</i>						
Age	27.66	27.75	0.781	27.71	27.83	0.747
Female	0.48	0.47	0.875	0.46	0.45	0.833
Goes to top 5 school	0.21	0.19	0.661	0.22	0.22	0.875
Lives in the United States	0.55	0.55	0.913	0.57	0.60	0.547
Wants to be an academic researcher	0.27	0.28	0.870	0.28	0.33	0.335
Reads Blogs Daily	0.15	0.14	0.769	0.13	0.17	0.279
Reads Economic Blogs	0.78	0.78	0.950	0.80	0.81	0.826
Number of Research Papers read (out of 12)	1.44	1.37	0.609	1.48	1.47	0.926
Currently doing a survey	0.33	0.31	0.543	0.33	0.33	0.929
Sample Size	311	308		233	212	

Appendix Table 3: Robustness of Matching Estimators in Table 4

	Control Group Mean	Table 4 Estimate	Excluding Blattman & Aid Watch readers	Falsification Test Impact of Reading Blattman
<i>Interest in Working as a researcher:</i>				
at World Bank	7.87	0.143 (0.316)	0.609 (0.533)	0.505 (0.343)
at IMF	5.18	-0.545 (0.471)	-1.364** (0.635)	-0.262 (0.384)
at top-10 research university	7.63	0.215 (0.402)	1.274** (0.613)	-0.782* (0.409)
at Liberal Arts University	5.55	-1.110*** (0.414)	-0.543 (0.768)	0.636* (0.380)
<i>Perception of Quality of Research Produced</i>				
at World Bank	7.73	0.483* (0.261)	0.338 (0.460)	-0.002 (0.243)
at IMF	6.39	0.288 (0.334)	0.048 (0.492)	-0.293 (0.312)
at Harvard, Yale and MIT	8.70	0.177 (0.154)	0.291 (0.353)	0.225 (0.184)
at selection of other schools	6.85	0.005 (0.207)	-0.051 (0.417)	0.005 (0.228)
<i>Perception of Extent to which World Bank staff face</i>				
Censorship over blog posts (1=high, 5 = low)	3.41	0.597*** (0.168)	0.811*** (0.250)	-0.318* (0.168)
<i>Awareness of Individuals</i>				
Proportion aware of 2 Development Impact bloggers	0.61	0.119* (0.069)	0.217* (0.111)	0.092 (0.073)
Proportion aware of 7 comparable researchers not mentioned on Development Impact blog	0.54	-0.038 (0.056)	0.120 (0.109)	0.132*** (0.050)
Sample Size		433	217	433

Notes: Robust standard errors in parentheses, *, ** and *** indicate significance at the 10%, 5% and 1% levels respectively.

Appendix Table 4: Robustness of Matching Estimators in Table 5

	Control Group Mean	Table 5 Estimate	Excluding Blattman & Aid Watch readers	Falsification Test Impact of Reading Blattman
<i>Knowledge</i>				
Number of questions correct about 6 papers	0.91	0.815*** (0.167)	0.598** (0.237)	0.220 (0.160)
Changed how they feel about effectiveness of an intervention in the past month	0.25	0.162** (0.082)	0.225 (0.149)	0.093 (0.081)
Sample size		439	220	439

Notes: Robust standard errors in parentheses, *, ** and *** indicate significance at the 10%, 5% and 1% levels respectively.

Appendix 2: Knowledge Questions (correct answers in bold)

1. In Alfredo Burlando's study of the impact of a black-out on infant health in Zanzibar, what did he find?
 - a. Infants born during the black-out were more likely to die in their first two weeks
 - b. Infants born 7-9 months after the black-out weighed less**
 - c. Infants conceived during the blackout weighed less
 - d. Mothers who knew they were pregnant weren't able to protect their fetuses from income shocks
 - e. All of the above
 - f. Don't know/have never heard of this study

2. Consider a randomized experiment in which only 25% of those in the treatment group take up the intervention, and that 0% of the control group get the treatment. Assume the treatment and control groups are the same size, and a constant treatment effect. How much does the sample size need to be to get the same power as you would get with a sample size of 1000 and 100% take-up?
 - a. 1250
 - b. 1500
 - c. 2000
 - d. 4000
 - e. 8000
 - f. 16000**
 - g. Don't know

3. In Barrera-Osorio, Bertrand, Linden and Perez-Calle's study of the impact of a conditional cash transfer program in Colombia, they look at impacts on both self-reported schooling and administrative schooling data. Which of the following do they find?
 - a. Students in both the treatment and the control groups over-report schooling**
 - b. Students in the control group, but not the treatment group, over-report schooling
 - c. Students in the treatment group, but not the control group, over-report schooling
 - d. Students in both the treatment and control groups accurately report schooling
 - e. Don't know/never heard of this study.

4. Consider an impact evaluation you are designed which uses a baseline and is deciding between doing one or two follow-up surveys at close intervals to one another. When will adding a second follow-up survey at a close interval be most useful?
 - a. **When the autocorrelation of the outcome of interest is close to zero.**
 - b. When the autocorrelation of the outcome of interest is close to 0.5
 - c. When the autocorrelation of the outcome of interest is close to 1
 - d. When the variance of the outcome of interest is very small
 - e. Don't know
5. In the study of Tarozzi, Mahajan and others on the impacts of introducing microfinance loans to buy bednets in Orissa, India, which of the following is a finding of the study?
 - a. Take-up of bednets was as high with consumer loans as it was with free distribution
 - b. Despite increasing bednet purchases, microcredit did not increase usage of bednets
 - c. **Neither microloans or free nets led to any measureable health improvements**
 - d. Microloans led to a 25% reduction in malaria episodes among households offered the loans
 - e. Don't know/I've never heard of this study.
6. In Ashraf, Lee and Field's work on increasing access to contraception in Zambia, which of the following is a finding of the study?
 - a. **Women were much more likely to use contraceptives and reduce unwanted births if they were seen separately from their husband**
 - b. There was no impact of increased access to contraception, suggesting high family sizes are optimal
 - c. Women needed to have their husbands present at the counseling sessions in order for the contraceptive intervention to have an effect
 - d. Women given contraceptives engaged in riskier sexual behavior
 - e. Don't know/never heard of this study.