

DISCUSSION PAPER SERIES

No. 8558

THE IMPACT OF ECONOMICS BLOGS

David J. McKenzie and Berk Özler

DEVELOPMENT ECONOMICS



Centre for Economic Policy Research

www.cepr.org

Available online at:

www.cepr.org/pubs/dps/DP8558.asp

THE IMPACT OF ECONOMICS BLOGS

David J. McKenzie, World Bank and CEPR
Berk Özler, World Bank

Discussion Paper No. 8558
September 2011

Centre for Economic Policy Research
77 Bastwick Street, London EC1V 3PZ, UK
Tel: (44 20) 7183 8801, Fax: (44 20) 7183 8820
Email: cepr@cepr.org, Website: www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programme in **DEVELOPMENT ECONOMICS**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: David J. McKenzie and Berk Özler

CEPR Discussion Paper No. 8558

September 2011

ABSTRACT

The Impact of Economics Blogs*

There is a proliferation of economics blogs, with increasing numbers of economists attracting large numbers of readers, yet little is known about the impact of this new medium. Using a variety of experimental and non-experimental techniques, we try to quantify some of their effects. First, links from blogs cause a striking increase in the number of abstract views and downloads of economics papers. Second, blogging raises the profile of the blogger (and his institution) and boosts their reputation above economists with similar publication records. Finally, we find that a blog can transform attitudes about some of the topics it covers.

JEL Classification: A11 and A23

Keywords: blog, dissemination, impact evaluation and influence

David J. McKenzie
Development Research Division
World Bank
1818 H Street, NW
Washington, DC 20433
USA

Berk Özler
Development Research Division
World Bank
1818 H Street, NW
Washington, DC 20433
USA

Email: dmckenzie@worldbank.org

Email: bozler@worldbank.org

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=167024

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=174165

* 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; 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.

Submitted 26 August 2011

1. Introduction

Practically nonexistent a decade ago, blogs by economic scholars have become commonplace since. 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.¹ Cowen (2008) argues that blogs are the main way that the general public consumes economics in a given day and guesstimates 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 on the open web, 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 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, coupling the large readership of blogs with the argument of Cowen (2008) that the best ones are written at a level far higher than that of any major newspapers 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.² On the other hand, Sunstein (2008) argues that the blogosphere might be causing “group polarization” and creating “information cocoons” – making it unlikely that blogs would cause a significant change in the knowledge and attitudes of their readers.³ Bell (2006, p.75) summarizes another common perception of blogs, as “...a largely harmless outlet for extroverted cranks and cheap entertainment for procrastinating office workers.”⁴ Combined

¹ 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.

² 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 received 73% more citations than control articles during the first year after publication.

³ 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.

⁴ In response to a post by Henry Farrell on the blog “Crooked Timber” asking academics why they started blogging and/or reading blogs, one responded: “I started reading blogs for news and to fill the downtime while I’m pondering the next sentence to write (what used to be filled with Minesweeper and other logic-based games).”

with the possibility that blogging gives scholars the freedom to write about topics outside their area of expertise (what Jacob T. Levy called ‘public-intellectualitis’ in his blog) this would suggest that impacts of blogs are likely to be negligible.

Second, it also seems plausible that blogging might enhance the reputation of the blogger. Kahn (2011) states: “The shrewd academic uses his blog to market his ideas and to ‘amplify’ his new academic results. This is a type of branding.” Reading bloggers talk about the reasons they blog produces a list of benefits to themselves: blogging is an outlet for ideas and observations not suitable for other media; blogging takes the blogger, at least the successful ones, from relative obscurity to being well-known; instant feedback allows the scholar blogger to sharpen her arguments by having to defend them publicly; it’s fun to interact with other economists; blogging may lead to regular writing gigs or other professional opportunities, etc. Furthermore, blogging by individual scholars may have positive spillover benefits for their institutions. Conversely, almost every scholar blogger seems to have questioned at some point on his blog about why he is spending so much time blogging instead of revising that manuscript.

While revealed preference suggests there is value in blog posts to both the scholars who maintain them and to the large number of individuals who read them, there is, to date, no quantitative evidence of their impacts.⁵ 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? Third, do they cause changes in attitudes among their readers or lead to increased knowledge? We conduct event study analysis using download data from RePEc to answer the first question, and combine RePEc data with recent survey data on most admired economists to answer the second. To answer the third set of questions, we use evidence from a recent survey we conducted on the role of blogs in economics, and take advantage of a randomized experiment in which a random sample of the respondents of this survey were encouraged to read a new economics blog. The novel survey and the experiment are also used to shed light on the question of whether blogs improve the public awareness and the reputation of their creators, and of their affiliated institutions.

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. We also find evidence that blogs enhance the recognition of their creators, so that bloggers are more likely to be admired or respected among their peers than other scholars of similar publication records. Using the experiment that provided random encouragement to read a new blog produced by four researchers from the

⁵ 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).

World Bank, we find an improvement in the perceived quality of research produced at the World Bank and an increased interest by the survey respondents in working there as a researcher.⁶ However, the same experiment shows that the short-run impacts on attitudes and knowledge are relatively weak. While we detect some changes in attitude among predetermined sub-groups of interest, we find no evidence of impact on the full sample. On knowledge, reading the new blog seems to improve awareness of findings from recent studies for the average reader, but not for the marginal reader who only reads because of encouragement. To our knowledge, this paper provides the first quantitative evidence on the impacts of economics blogs and covers a number of the more important likely channels of influence.

The rest of the paper is as follows. Section 2 presents the impacts of blogs on dissemination of economics research using event study analysis, and Section 3 provides evidence on the effect of blogs on the reputation of their creators. Section 4 uses our new survey data to describe how blogs are used by graduate students in economics, junior faculty, and field workers in development. Section 5 presents impacts of reading a new blog, exploiting the randomized encouragement design. Section 6 provides a discussion of whether economics blogs influence policy, while Section 7 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 economics blogs have devoted considerable space to discussing public policy issues in the U.S. such as different plans for dealing with budget deficits, but it is difficult to assess how much any one particular blog post has influenced this debate. 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.

The typical economics working paper gets very few readers, especially after its first couple of months of release. For example, a random sample of papers released in the NBER working paper series in January 2010 shows that the median paper in this prestigious series received 21 abstract views and 12 downloads through RePEc services in the first two months, and then an average of 6-7 abstract views and 2-3 downloads per month through RePEc over the next year. Given these low readership levels, blog posts that draw attention to such research can potentially have large relative impacts on readership.

⁶ McKenzie and Özler are two of the four researchers who produce the “Development Impact” blog.

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 restrict our analysis to blog posts which link directly to papers in the RePEc – the largest database of economics papers. Monthly abstract views and download statistics are publicly available from this site. The second issue is that, in some cases, papers are linked to when they are first released in a working paper series. It is harder to form a counterfactual in these cases, 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 make comparison to other papers in the same series or by the same author not necessarily a good counterfactual.

We therefore focus on blog postings to papers which have been out several months at the time of a blog posting, and which are listed in RePEc. Figure 1 then 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 2010.⁷ The paper then went back to averaging 0.8 downloads a month and 5.3 abstract views a month from April 2010 through March 2011.

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 1478 abstract views and 144 downloads in the month it was blogged about. 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 blogged about.

2.2 Formal Estimation

We systematically searched the blogs of the 50 most read finance and economics blogs (defined according to one list based on blogs with public traffic logs⁸) for references to research papers in RePEc. In addition we also looked at six other popular blogs that were not included on this list (Aid Watch, Chris Blattman, Economix, Freakonomics, Paul Krugman, and Dani Rodrik). We included only papers where the blogger directly linked to the RePEc version of

⁷ 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.

⁸ <http://www.gongol.com/lists/bizeconsites/> [March 2011 data, accessed August 2011].

the working paper, and where the paper had been out for at least three months before being blogged about. We also exclude blogs which only linked to one or two working papers at most.⁹ This resulted in a database of 107 research papers linked to on one of 8 blogs: Aid Watch, Baseline Scenario, Chris Blattman, Economix, Freakonomics, Marginal Revolution, Overcoming Bias, and Paul Krugman. Among the other blogs, the majority were either finance blogs or blogs based on discussion of macroeconomic issues typically without reference to working papers. 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 which 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. Then we estimate the impact of blog s blogging about 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) and cluster the standard errors at the paper level. 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 that a particular paper attracts a lot of attention for some unrelated reason in month t , resulting in a

⁹ 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.

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 blog post on July 8, 2010. 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 within up to 2 years on either side of the blogging month. 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.¹⁰ These statistical results are therefore consistent with the graphical illustrations seen in Figures 1-3, and show a big spike in abstract views and downloads in the month that the paper is blogged about.

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: 10.3 abstract views per month and 4.2 downloads per month from RePEc 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 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 which 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

¹⁰ 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.

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.¹¹ This seems consistent with the intuition that as an academic's blog expands readership to a larger and larger audience, the additional readers are less likely to be interested in the academic papers.

Finally, we should note that the estimates in Table 1 show the average impacts of being linked to by 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 blogged about compared to the mean abstract views over the months before the paper was blogged about 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 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. single line link; positive, neutral or negative link, etc.). Unfortunately there are insufficient data in our sample to explore this phenomenon systematically.

3. Probit Analysis of Influence

When economic bloggers discuss what the benefits of blogging are, much of the discussion tends to be self-centered – they view it as something beneficial to themselves. Thus blogs are viewed as providing a space to catalog and flesh out ideas, get feedback and a sense of popular applications for ideas, and to disseminate their own ideas to a wider audience (Nash, 2008). Yet it is difficult to quantify or measure many of these effects, although several bloggers attest that they occur.¹² In this section we attempt to at least provide modest evidence that the reputation of economic bloggers in the economics profession exceeds that of non-bloggers with similar publication records.

Davis et al. (2011) conducted a survey of academic economists in the U.S., with 299 (15%) responding. The survey asked these academics to list up to three living economists over the age of 60 and up to three under the age of 60 who they “regard with great respect, admiration, or reverence”. Gary Becker, Ken

¹¹ 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.

¹² See for example Daniel Drezner (<http://www.danieldrezner.com/archives/000727.html>) and Scott Sumner (<http://www.themoneyillusion.com/?p=4058>).

Arrow and Gary Solow were the top choices among the over 60s, and Paul Krugman, Gregory Mankiw and Daron Acemoglu the top choices among the under 60s. The under 60s list of 23 names contains a number of regular bloggers – in addition to Krugman and Mankiw are Steven Levitt, William Easterly, Nancy Folbre, Dani Rodrik, and Tyler Cowen.

We merge this list with a list of the top 500 economists according to the RePEc rankings (based on paper downloads, citations) and also code each of the RePEc top 500 according to whether they blog or not.¹³ This data is then used to estimate a probit model to see whether, conditional on RePEc ranking, individuals who blog are more likely to appear on the list of favorite or admired economists. Table 2 shows the results, for the pooled sample in column 1, and separately for under 60 and over 60 economists in columns 2 and 3 respectively. In all three columns we see that, conditional on their RePEc rank, regular blogging is strongly and significantly associated with being more likely to be viewed as a favorite economist. The point estimates are just over 0.4 in all three columns, indicating a 40 percentage point higher likelihood of being on the favorite economist list for economics bloggers. Blogging has the same size impact as being in the top 50 of RePEc rankings for the under 60 economists, and a larger impact for the over 60 economists.

This evidence is thus consistent with the view that blogging helps build prestige and recognition in the profession, with bloggers being more likely to be admired or respected than other academics of similar (or in many cases better) publication records. This is of course only a correlation, and there are several caveats to consider. First, to the extent that blogging serves to increase the RePEc ratings by increasing downloads (as seen in the previous section) and citations, the observed correlation will be a lower bound on the causal impact of blogging.¹⁴ However, if bloggers are also more likely to be engaged in other activities of a public intellectual, such as media appearances, writing books etc., and if these don't all arise directly as a result of blogging, the estimates will conflate the impact of blogging with the impacts of these other activities, thereby overstating the impact of blogs. Nevertheless, given the large magnitude of the coefficient observed, it does not seem likely that all of the observed impact of blogging just reflects omitted variables, and therefore we view this evidence as strongly

¹³ The economists from the favorite economist list who appear in RePEc outside the top 500 are coded based on their RePEc ranking. A small number of the favorite economists do not appear in RePEc. Their RePEc equivalent ranks were calculated based on total citation counts and publication h-statistics from Google Scholar.

¹⁴ Mixon and Upadhyaya (2010) show that among bloggers, there is a positive correlation between the readership levels of the blog and the academic citation levels of the bloggers, although it is not clear whether this reflects a causal relationship in either direction.

suggesting that blogging increases the influence, respect, or public image of the blogger.

4. 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 potential readership groups for blogs about development economics issues. Development economics is a useful case to examine for several reasons. First, it is a field in which a mixture of academic and less academic potential audiences can be identified, which may be more difficult with some other popular fields for blogging such as macroeconomics. Secondly, since the content is international in scope, it opens up the possibility of considering readers in a range of different countries. Third, 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.

4.1. Sample frames and response rates

The first group we identified consists of students in Ph.D. and Masters 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 participate online in the survey to their graduate students. This faculty list was comprised of members of the organization BREAD who teach development at a school with a Ph.D. or Masters program, as well as additional faculty who through personal contacts we knew to teach development at this level. Students were told that the purpose of the survey was to find out how the next generation of development economists find out about new studies in development economics, and about the role of social media such as blogs in these surveys. 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.

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 is the group 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 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).

The top panel of Table 3 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 are 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.

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

4.2 Self-reported uses of economics blogs in survey data

The second panel of Table 3 provides more detailed information on how individuals read and use blogs. Readership of economics blogs is high among all 5 groups surveyed, with between 76 and 84 percent of those surveyed having read

an economics blog in the past 6 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 view blog postings by entering the blog posting address in their browser and viewing the posting this way, as opposed to get a feed through an RSS viewer such as Google Read, or through 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 3 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 3 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. Given that preparing surveys is likely to be an infrequent event, it is still possible blogs are influencing questionnaire design over a longer period. 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.

5. 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 not to. This is difficult to do for existing well-known blogs, 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 Impact*¹⁵, which was launched at the start of April 2011, and conduct a randomized encouragement experiment with 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. In May and June 2011 it received a total of just under 50,000 page views, or just over 800 page views per day. Whilst 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 one ranking¹⁶. 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 a similar number of Google Reader subscribers as the Center for Global Development's Views from the Center blog.¹⁷ 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.

5.1. Randomized Encouragement

We took the 619 graduate student and field staff respondents to our baseline survey that had provided contact email addresses as the sample for this experiment.¹⁸ We stratified the data by three variables thought to potentially influence attitudes towards research methods and interest in the World Bank: whether they were a Ph.D. student, field staff, or Masters student; whether or not they said they read Chris Blattman's blog at baseline; and whether or not they said they were involved in a randomized experiment at baseline. Individuals were then

¹⁵ <http://blogs.worldbank.org/impactevaluations/>

¹⁶ <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.

¹⁷ Chris Blattman traffic statistics based on a blog posting in which he said he had 800,000 page views in 2010.

¹⁸ 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.

randomly assigned to treatment and control within strata, with Appendix Table 1 showing that this succeeded in balancing baseline observable characteristics.

The encouragement then consisted of two emails. The first was sent on April 6th, 5 days after the blog 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.”. They then received a second, shorter, email 3 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.

5.2. Follow-up Survey

A follow-up survey was then sent at the start of June 2011 to both treatment and control groups. This was therefore 2 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, 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 high 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 get them to reply were less likely to be frequent blog readers at baseline. Appendix Table 1 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.

5.3. 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 4 shows that the encouragement succeeded in increasing the proportion of respondents who read Development Impact by 9.9 percentage points, a more than 50 percent increase relative to the 18 percent of the control group who had read development impact in the last month. Columns 2 and 3 then re-estimate equation (1) by gender, and by whether or not the individual at baseline said they wish 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. 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).

5.4. Impacts of the Development Impact blog

We are then interested in using 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:

$$\text{Outcome}_i = \alpha_s + \gamma \text{Read Development Impact}_i + \omega_i \quad (3)$$

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 4). Moreover, this is potentially the parameter of interest

for answering questions like should blogs attempt outreach exercises to get 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 and Imbens (2006) to estimate the average treatment effect (ATE). We match on a wide range of baseline characteristics: age, gender, whether or not the individuals attend an elite (top 5 economics department¹⁹, whether they live in the U.S., whether they are a Ph.D. student, Masters student or field staff, whether they plan on a career as an academic researcher, whether they are currently involved in conducting a survey, whether they read economic blogs, the frequency of reading, and dummies for readership of the 4 most commonly read blogs in our survey, and the number of current research papers (out of 12) they had read at baseline. The identifying assumption is then that, conditional on this rich set 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, during the initial few months of a blog, whereby potential readers are coming across the blog by chance, 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.

Successful blogs are often argued to improve the reputation of the individuals and institutions producing these blogs, as suggested by the evidence on individuals in Section 4. We therefore examine in Table 5 whether reading the Development Impact blog changes readers opinions about the World Bank. 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 TOT using the encouragement experiment on the full sample, while column 3 shows the ATE estimated using matching. Columns 4 through 7 then look at the experimental effects for the two groups which responded 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 matching analysis and for the experimental estimates on the research-focused group. For this research-focused group, there is also a

¹⁹ 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.

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 5 looks at 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 quality perceptions for both the full sample, and for the male and research-focused subsamples. Blog readership seems to have spillover results on the reputation of the IMF's research (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. In contrast, there is no significant impact on the perceived quality of research at a range of other good, but not very top, economics programs. Taken together these results therefore show that, over a very short term, 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 more conversational tone than on most of the World Bank's blogs. 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 ATE 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 post.

Finally, we asked the survey respondents about their awareness of a list of 10 development researchers, all approximately 5-10 years post-Ph.D. Included amongst this list were the two other regular bloggers on the Development Impact blog (since the survey was being administered by ourselves, we didn't ask whether people had heard of us or our work). The final panel of Table 5 shows some evidence for greater awareness of the bloggers as a result of reading Development Impact. The experimental impact is positive and significant at the 10 percent level for males, but insignificant for the research-focused sub-sample. The ATE is strongly significant, which could reflect a strong average impact, or that individuals who already knew of our fellow bloggers were more likely to come across the blog and start reading it.

Table 5 therefore shows a number of significant changes in attitudes and general impressions towards the World Bank and its researchers as a result of readership. The mere existence of the blog and a casual reading of articles to get a sense of the issues being discussed may be sufficient to result in this level of

change. In Table 6 we look for changes in knowledge and attitudes which might only be expected to occur from more in-depth reading. To measure knowledge, we asked detailed questions related to 6 blog posts that had appeared on the Development Impact blog (appendix 2 provides the exact questions). These questions 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.

Two possible interpretations for this difference between the ATE and ITT/TOT suggest themselves. The first 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 baseline blog reading variables. Therefore we are controlling for a large number of characteristics that might well proxy for knowledge. A second explanation is therefore that reading the blog impacts knowledge for the average reader, but not for the marginal reader who only reads because of encouragement. This is plausible 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 and learn from it.

Finally, we examine 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.

The bottom of Table 6 shows that blog readership has not changed many of these attitudes towards methodology, with no significant experimental changes in the full sample. Amongst the subsamples, the most significant change occurs in the male sample, where there is an increase in the proportion who believe that it is difficult to succeed as a development economist on the job market without having a randomized experiment. The first two months of postings focused heavily on experimental studies, which may have led to this impression, although the ATE estimated through matching is negative and marginally significant. There is also some evidence among the research-focused subsample that more agree with the statement that external validity is no more of a concern in experiments than in

most non-experimental studies (something discussed in a well-read blog post)²⁰. Nevertheless, given the number of outcomes tested here, only the change for males would continue to be significant once p-values are multiplied by 8 to account for testing 8 different attitude questions in the sample.

Given the short period of time over which results are observed, it seems likely that the estimates obtained here are lower bounds for what the impact of blog readership may be over the medium to long-term. The fact that we find a number of large and significant changes even over this shorter period therefore suggests that blogs can play a role in changing opinions, attitudes, and knowledge.

6. Do blogs influence policy?

Assessing the impacts of blogs on policy faces many of the same problems as assessing the extent to which economics research in general, and the media in particular, influence policy. Discussions in blogs may provide talking points or inform the discussion, but it is difficult to directly trace particular blog posts to particular policy actions. For example, Austin Frakt argues that blogs help connect the world of research with that of journalism and policy by being able to connect the results from good research studies to policy questions at the right point in time, thereby influencing the debate more effectively.²¹ Even small blogs may have an impact to the extent that their stories or points of view are picked up by elite bloggers, as influential people read the latter (Drezner 2005). Drezner and Farrell (2008) provide anecdotal evidence that blogs might be playing an important role in politics by vigorously prodding media attention to certain events, by helping shape political campaign tactics, by affecting legal outcomes, and by influencing policy outcomes.

Since we are unaware of any attempt to provide even the same level of systematic anecdotal evidence for economics blogs, we reviewed major blogs and solicited feedback from both bloggers and blog audiences as to what policy impacts blogs might actually be having. We conclude from this that any direct influence on policy is rare. For example, Gregory Mankiw replied “nothing specific comes to mind,” while Steven Levitt said that “I can think of lots of cases we thought we should affect policy, but none where we actually did!”. Dani Rodrik also could not point to any specific cases, but said he is amazed “how widely the stuff in the blog is read, and I have had many policymakers tell me about one or other point I had made in the blog.”²²

²⁰ <http://blogs.worldbank.org/impac-tevaluations/a-rant-on-the-external-validity-double-double-standard>

²¹ See <http://theincidentaleconomist.com/wordpress/on-blogging-academyhealth-blog/>

²² Personal communications with authors.

Among the cases where it does appear blogs influenced policy, the role of the blog has been more as investigative journalism than as disseminator or originator of economic analysis. For example, David Roodman blogged about *Kiva*, a popular peer-to-peer microlending website, and explained through his post that it operated in a different way than implied, leading to coverage in the New York Times and a modification in Kiva's website.²³ Justin Wolfers blogged on Freakonomics about a paper that showed that the process that the U.S. Census Bureau uses to make public-use micro-data anonymous was inducing systematic errors in these data, which led to coverage in the Wall Street Journal and most likely caused the Census bureau fixing the mistake and re-issuing the data.²⁴

However, even if the vast majority of blog posts do not directly influence policy, it only takes the right reader reading the right post(s) to potentially shape policy in important ways. Tim Worstall has blogged on multiple occasions about the gender pay gap in the United Kingdom, and how the main cause is due to child care. He reports, although he has no documented proof, having been told that as a result of one reader of these posts being involved in the Liberal Democrat Party, the party was urged to, and did adopt, a policy that there should be transferable maternity/paternity leave rather than just maternity leave – a policy that was ultimately adopted by the Government.²⁵ The survey evidence in Section 5 shows that a sizeable share of the World Bank operational economists and NGO staff say they have changed their views of the effectiveness of particular interventions as a result of blog posts. Given that some of these individuals are in a position to indeed influence policy change, it seems likely that blogs are sometimes reaching the right readers to at the right time. These cases illustrate that it may be possible for blog posts to influence policies. However, given that we were able to gather only a few examples from discussions widely held on several economics blogs with thousands of readers, it appears that such cases may in fact be very rare.

²³ See http://blogs.cgdev.org/open_book/2009/10/kiva-is-not-quite-what-it-seems.php. Kiva made it appear as if lenders were funding specific projects, who would receive loans if enough people funded them, when in fact the projects shown were almost all funded in advance.

²⁴ <http://www.freakonomics.com/2010/02/02/can-you-trust-census-data/>

²⁵ See the comments in <http://blogs.worldbank.org/impactevaluations/the-impact-of-blogs-part-ii-blogging-enhances-the-blogger-s-reputation-but-does-it-influence-policy>

7. Conclusions

Economic blogs are doing more than providing a new source of procrastination for writers and readers. Using a variety of empirical approaches, we have provided the first quantitative evidence that they are having impacts. There are large impacts on dissemination of research – a link on a popular blog results in a substantial increase in abstract views and downloads, while a majority of economics blog readers say they have read a new paper in the past month as a result of a blog. There also appear to be benefits in terms of the bloggers becoming better known and more respected within the profession – bloggers are over-represented relative to their academic publication records in a poll of favorite economists, and readers of a new blog have become more aware of the researchers writing it. Finally, we find some evidence from our experiment that blogs influence attitudes and knowledge: readers of the new Development Impact blog think more highly of World Bank research and are more interested in working in this institution, and the average reader may have gained knowledge about the contents of recent research papers as a result of reading this blog.

The fact that blog posts are available worldwide immediately after posting poses challenges for evaluating their impact. A further contribution of this paper is therefore in illustrating a variety of methods that can be used to assess the impacts of a blog. These methods can therefore be readily adapted in future work to consider the impact of other economics blogs not considered here, as well as blogs in other academic disciplines.

One natural response to our results is to ask why, given these benefits, more economists don't blog? Tyler Cowen has argued that the answer to this question is "because they can't, at least not without embarrassing themselves rather quickly, even if they are smart and very good economists. It's simply a different set of skills".²⁶ However, our results show that there are a number of positive externalities from economics blogs that are unlikely to be captured by the blogger him or herself: bloggers increase the dissemination of other people's research (in addition to their own work), and can have positive effects on the reputation of their institutions. The presence of these externalities, coupled with costs to blogging (such as the time cost), suggest that there may be an undersupply of good economics blogs.

²⁶ <http://marginalrevolution.com/marginalrevolution/2011/08/does-blogging-help-ones-professional-reputation-as-an-economist.html> [accessed August 22, 2011]

References

- Abadie, Alberto and Guido Imbens (2006) "Large sample properties of matching estimators for average treatment effects", *Econometrica* 74(1): 431-97.
- Arai, Mahmood and Peter Thoursie (2006) "Giving up foreign names: An empirical examination of surname change and earnings", *Stockholm University Department of Economics Working Paper*.
- Angrist, Joshua, Guido Imbens and Donald Rubin (1996) "Identification of causal effects using instrumental variables", *Journal of the American Statistical Association* 91(434): 444-455.
- Ball, Samuel and Gerry Bogatz (1970) *The First Year of Sesame Street: An Evaluation*. Educational Testing Service, Princeton, NJ.
- Bell, Tom W. (2006) "The Impact of Blogging on the Practice of Law: Hit the Snooze Button," *Nexus* 11(1): 75-78.
<http://tomwbell.com/writings/Snooze.pdf>
- Bruhn, Miriam and David McKenzie (2009) "In pursuit of balance: Randomization in practice in development field experiments", *American Economic Journal: Applied Economics* 1(4): 200-32.
- Campbell, John, Andrew Lo and A. Craig MacKinlay (1997) *The Econometrics of Financial Markets*. Princeton University Press. Princeton, NJ.
- Cowen, Tyler (2008) "Why everything has changed: the recent revolution in cultural economics", *Journal of Cultural Economics* 32: 261-73.
- Davis, William, Bob Figgins, David Hedengren, and Daniel Klein (2011) "Economics professors' favorite economic thinkers, journals, and blogs (along with party and policy views)", *Econ Journal Watch* 8(2): 126-46.
- Deaton, Angus (2010) "Instruments, randomization, and learning about development", *Journal of Economic Literature*, 48(2): 424-55
- Dellavigna, Stefano and Ethan Kaplan (2007) "The Fox news effect: Media bias and voting", *Quarterly Journal of Economics* 122(3): 1187-1234.
- Drezner, Daniel (2005). "Confessions of a Scholar-Blogger." *University of Chicago Magazine*, 97(3).
- Drezner, Daniel and Henry Farrell (2008) "Introduction: Blogs, politics and power: a special issue of Public Choice", *Public Choice* 134: 1-13.
- Glenn, Owen (2003) "Scholars who blog", *Chronicle of Higher Education*, June 6.
- Irwin, Douglas (1997) "Higher tariffs, lower revenues? Analyzing the fiscal aspects of the "Great tariff debate of 1888", *NBER Working Paper no. 6239*.
- Kahn, Matthew (2011) "Where are the female economics bloggers?", *Christian Science Monitor*, March 26 <http://www.csmonitor.com/Business/Green-Economics/2011/0326/Where-are-the-female-economics-bloggers> [accessed August 2, 2011].

- Landry, Craig, Andreas Lange, John List, Michael Price, and Nicholas Rupp (2006) "Toward an understanding of the economics of charity: Evidence from a field experiment", *Quarterly Journal of Economics* 121(2): 747-82.
- McKenna, Laura and Antoinette Pole (2008) "What do bloggers do: an average day on a political blog", *Public Choice* 134: 97-108.
- Mixon, Franklin and Kamal Upadhyaya (2010) "Blogometrics", *Eastern Economic Journal* 36: 1-10.
- Mortensen, Dale and Christopher Pissarides (1994) "Job creation and job destruction in the theory of unemployment", *Review of Economic Studies* 61(3): 397-415.
- Nash, Betty Joyce (2008) "Economic blogs", *Richmond Federal Reserve Region Focus*, Spring/Summer p. 21-24.
- Paluck, Elizabeth (2009) "Reducing Intergroup Prejudice and Conflict Using the Media: A Field Experiment in Rwanda", *Journal of Personality and Social Psychology* 96(3): 574-87.
- Ravallion, Martin (2009) "Should the Randomistas rule?", *The Economists Voice* 6(2).
- Phillips, David P., Elliot J. Kanter, Bridget Bednarczyk, and Patricia L. Tastad (1991) "Importance of the Lay Press in the Transmission of Medical Knowledge to the Scientific Community," *The New England Journal of Medicine* 325(16): 1180-1183.
- Ribstein, Larry (2006) "From Bricks to Pajamas: The Law and Economics of Amateur Journalism," *William & Mary Law Review* 48(1): 185-249. <http://scholarship.law.wm.edu/wmlr/vol48/iss1/3>
- Rockey, James (2009) "Who is left-wing, and who just thinks they are?", University of Leicester Working Paper.
- Sunstein, Cass (2008) "Neither Hayek nor Habermas", *Public Choice* 134: 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)	448.9*** (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: Does blogging increase reputation?

Marginal effects from Probit estimation of the likelihood of being a favorite economist

	(1) All	(2) Under 60	(3) Over 60
Blog regularly	0.426*** (0.127)	0.413*** (0.158)	0.406* (0.245)
In Repec top 50	0.341*** (0.0748)	0.418*** (0.0958)	0.243** (0.118)
In Repec rank 51-100	0.107* (0.0600)	0.157* (0.0894)	0.0269 (0.0761)
In Repec rank 101-200	0.00750 (0.0335)	0.0750 (0.0479)	-0.129*** (0.0395)
Proportion of sample on favorite list:	0.093	0.069	0.146
Observations	514	350	164

Notes: Robust standard errors in parentheses,

*, **, and *** indicate significance at 10, 5 and 1% levels respectively.

Table 3: 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 4: Did the Encouragement Work, and for Whom?

Dependent Variable: Read Development Impact Blog in last month

	(1) Full sample	(2) Males	(3) Females	(4) Research Focused	(5) 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 5: Impact of Reading Development Impact Blog on Perceptions of Institutions

	Control Group Mean	Full sample			Males		Research-focused	
		(1) ITT	(2) TOT	(3) Matching	(4) ITT	(5) TOT	(6) ITT	(7) TOT
Interest in Working as a researcher:								
at World Bank	7.87	0.122 (0.213)	1.243 (2.151)	0.514* (0.301)	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)	-0.534 (0.430)	-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)	0.282 (0.370)	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)	-0.364 (0.377)	-0.278 (0.343)	-1.810 (2.264)	-0.893** (0.420)	-4.464* (2.423)
Perception of Quality of Research Produced								
at World Bank	7.73	0.309** (0.156)	2.968* (1.681)	0.442* (0.232)	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)	0.052 (0.312)	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)	0.346** (0.169)	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)	0.010 (0.186)	0.258 (0.184)	1.825 (1.360)	0.113 (0.244)	0.524 (1.071)
Perception of Extent to which World Bank staff face								
Censorship over blog posts (1=high, 5 = low)	3.41	0.130 (0.107)	1.147 (0.921)	0.711*** (0.149)	0.208 (0.144)	1.296 (0.860)	0.537*** (0.188)	2.465*** (0.931)
Awareness of Individuals								
Proportion aware of 2 Development Impact bloggers	0.61	0.0120 (0.0410)	0.114 (0.379)	0.168*** (0.062)	0.107* (0.0575)	0.837* (0.491)	-0.0206 (0.0740)	-0.102 (0.358)
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 6: Impact of Reading Development Impact Blog on Reader Knowledge and Attitudes

	Control Group Mean	Full sample			Males		Research-focused	
		(1) ITT	(2) TOT	(3) Matching	(4) ITT	(5) TOT	(6) ITT	(7) TOT
Knowledge								
Number of questions correct about 6 papers	0.91	-0.103 (0.0982)	-1.038 (1.159)	0.655*** (0.151)	-0.0367 (0.133)	-0.267 (0.975)	0.0273 (0.183)	0.140 (0.889)
Attitudes or Beliefs: Agree or Strongly agree that...								
structural models rely too much on functional form assumptions and are unlikely to yield reliable estimates	0.41	-0.0713 (0.0482)	-0.619 (0.458)	-0.098 (0.077)	-0.00947 (0.0683)	-0.0614 (0.443)	-0.107 (0.0890)	-0.510 (0.456)
the economics profession focuses too much on identification	0.31	-0.0164 (0.0462)	-0.183 (0.518)	-0.005 (0.074)	0.0456 (0.0652)	0.362 (0.526)	0.0324 (0.0802)	0.152 (0.375)
they are likely to reject any paper that relies on propensity-score matching for identification	0.11	0.00452 (0.0346)	0.0543 (0.417)	-0.005 (0.056)	0.0234 (0.0512)	0.154 (0.342)	-0.0462 (0.0721)	-0.193 (0.303)
development economics has moved too far away from its core purpose of understanding the drivers of economic growth	0.21	-0.0270 (0.0392)	-0.235 (0.349)	-0.023 (0.060)	0.0528 (0.0593)	0.361 (0.430)	-0.0547 (0.0676)	-0.250 (0.325)
experiments have no special ability to produce more credible knowledge than other methods	0.12	0.0490 (0.0340)	0.454 (0.351)	-0.013 (0.050)	0.0127 (0.0492)	0.0908 (0.349)	0.0518 (0.0652)	0.254 (0.331)
development researchers are letting methodology drive the questions they answer	0.59	0.0218 (0.0483)	0.209 (0.473)	-0.152* (0.078)	0.100 (0.0655)	0.721 (0.552)	0.0942 (0.0879)	0.447 (0.443)
that externality validity is no greater in most non-experimental studies than it is in most experiments	0.43	-0.00739 (0.0487)	-0.0642 (0.425)	-0.006 (0.077)	-0.0151 (0.0685)	-0.103 (0.471)	0.203** (0.0882)	0.827* (0.446)
it is difficult to succeed as a development economist on the job market without having a randomized experiment	0.34	0.0737 (0.0504)	0.656 (0.519)	-0.134* (0.074)	0.193*** (0.0684)	1.168** (0.562)	0.102 (0.0861)	0.533 (0.518)
Sample size		445	445	439	239	239	135	135

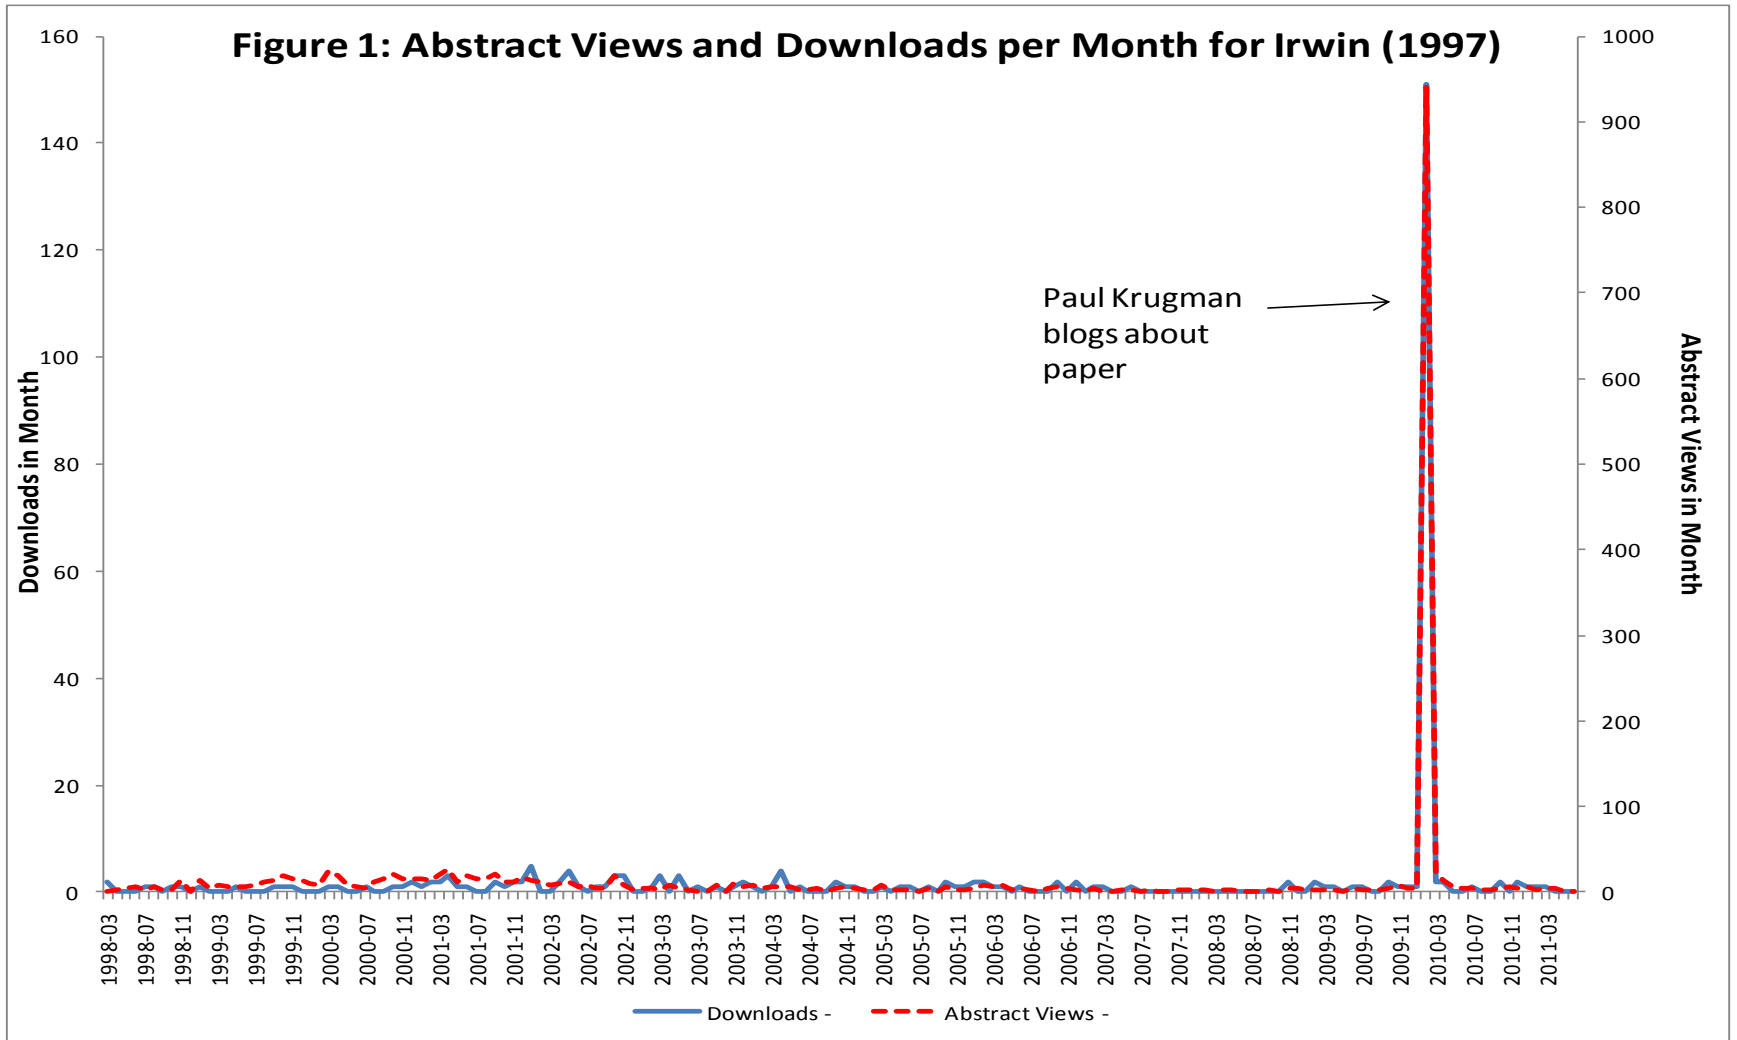


Figure 2: Abstract Views and Downloads per Month for Landry et al. (2006)

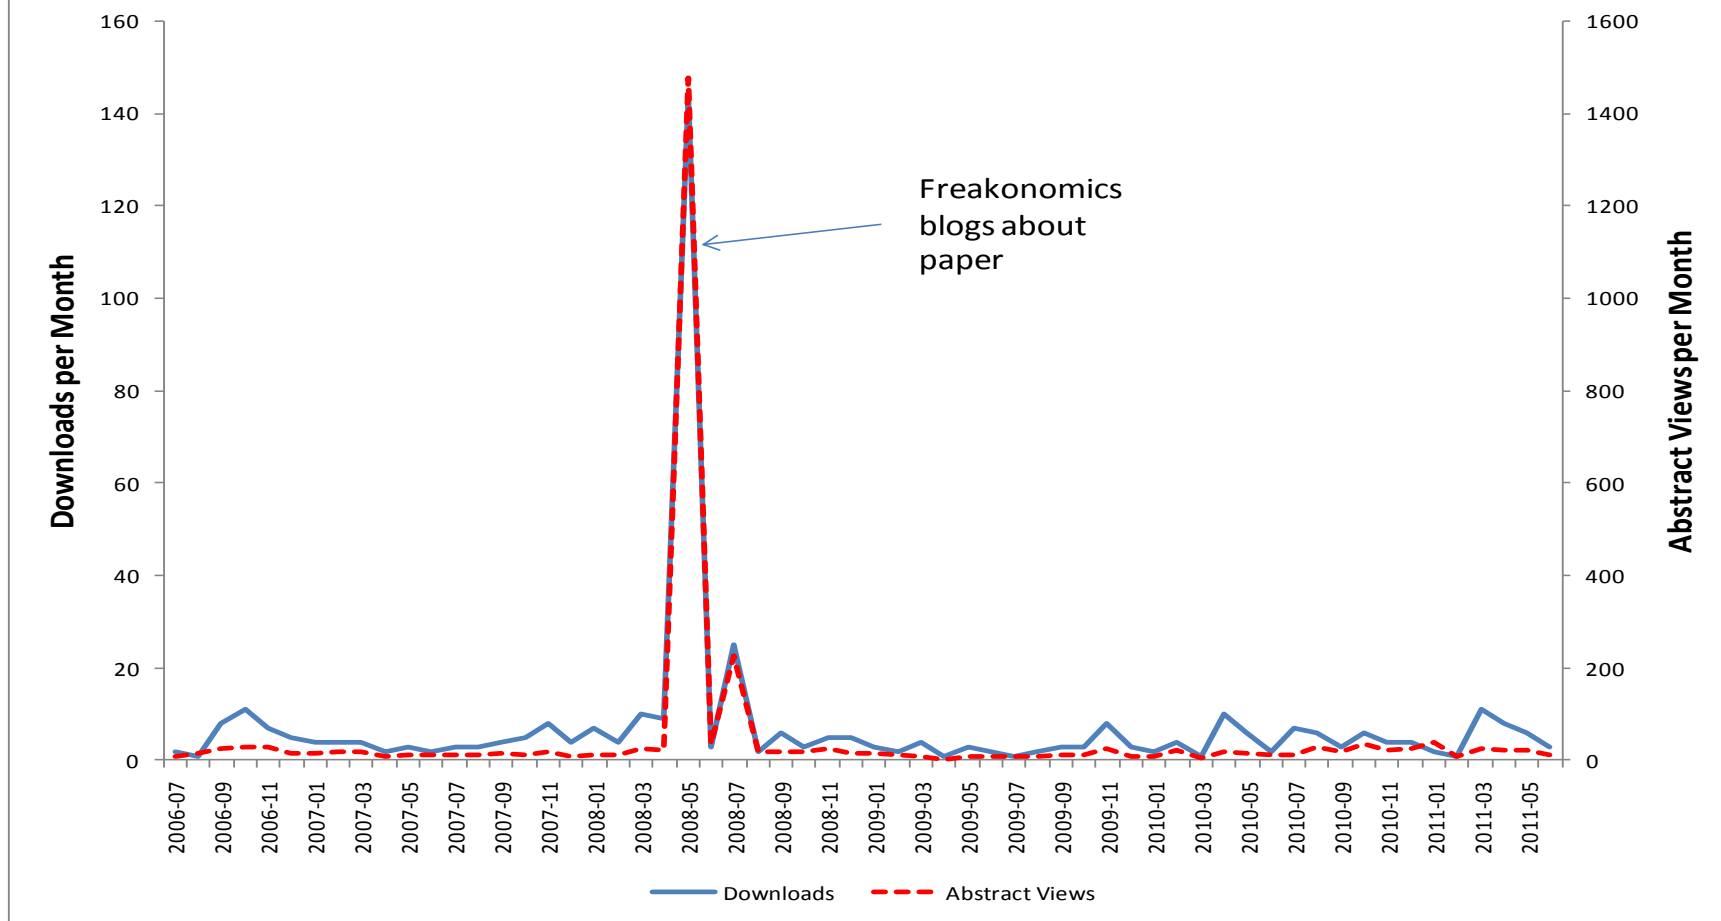
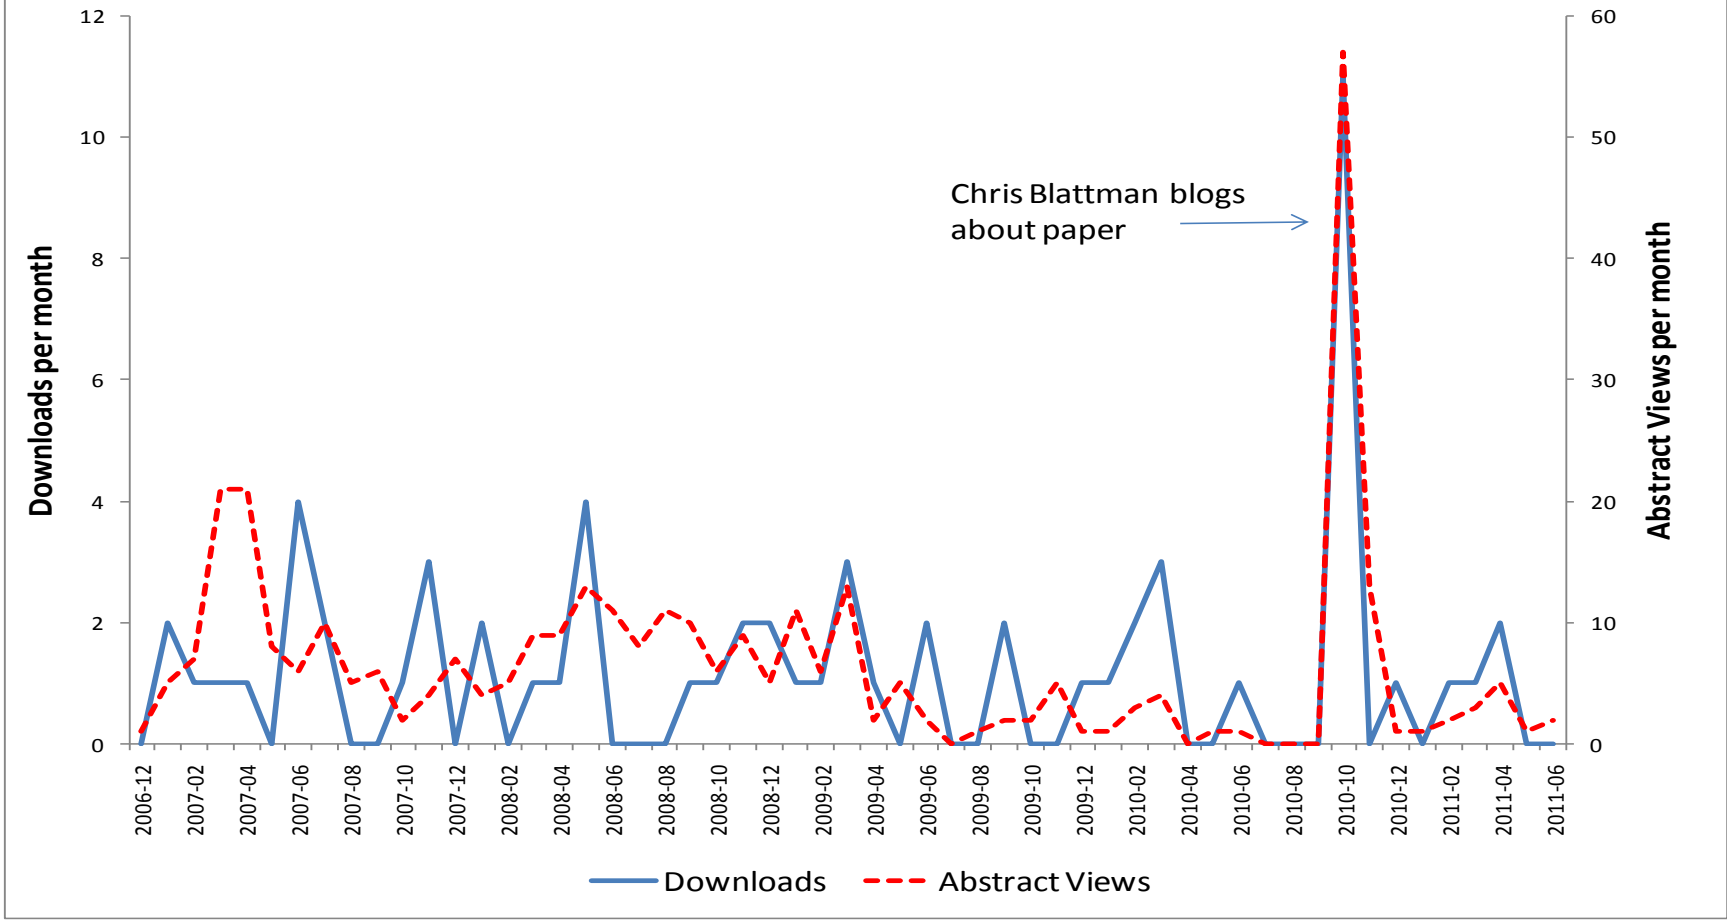


Figure 3: Abstract Views and Downloads per Month for Arai and Thoursie (2006)



Appendix 1

Appendix Table 1: 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 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.