

The Impact of Economics Blogs*

David McKenzie, *World Bank, BREAD, CEPR and IZA*

Berk Özler, *World Bank*

Abstract

There is a proliferation of economics blogs, with increasing numbers of famous and not-so-famous economists devoting a significant amount of time to writing blog entries, and in some cases, attracting large numbers of readers. Yet little is known about the impact of this new medium. We utilize event study analysis, regression, original survey evidence, and a randomized experiment to measure several key impacts of blogs. Striking results are found in terms of dissemination of research, with blog postings causing a large increase in the number of abstract views and downloads of linked papers. There is also evidence to suggest blogging raises the profile of the individuals producing this content, resulting in bloggers being viewed more favorably in the profession than other economists of similar publication records. Survey evidence indicates that many researchers say they have read new papers or changed their research design as a result of a blog post. Finally, data from an experiment which randomly encouraged readership of a new economics blog shows that even in a few months a blog can change impressions about the quality of an institution, raise awareness of the people producing it, and change attitudes about some of the topics covered.

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; and all the individuals who took time to participate in this survey. All views expressed are those of the authors only. This document contains only some sections of the manuscript.

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 survey 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*¹, 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 ranking². 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.³ As such, while the findings of the

¹ <http://blogs.worldbank.org/impactevaluations/>

² <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.

³ Chris Blattman traffic statistics based on a blog posting in which he said he had 800,000 page views in 2010.

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.⁴ 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 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 took more attempts to respond 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,

⁴ 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.

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:

$$Read\ Development\ Impact_i = \alpha_s + \beta Treat_i + \varepsilon_i \quad (1)$$

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:

$$Outcome_i = \alpha_s + \gamma Read\ Development\ Impact_i + \omega_i \quad (2)$$

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 department⁵, 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 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

⁵ 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".

– 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)⁶. 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 <http://blogs.worldbank.org/impactevaluations/a-rant-on-the-external-validity-double-double-standard>

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

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?
- Students in both the treatment and the control groups over-report schooling**
 - Students in the control group, but not the treatment group, over-report schooling
 - Students in the treatment group, but not the control group, over-report schooling
 - Students in both the treatment and control groups accurately report schooling
 - 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?
- When the autocorrelation of the outcome of interest is close to zero.**
 - When the autocorrelation of the outcome of interest is close to 0.5
 - When the autocorrelation of the outcome of interest is close to 1
 - When the variance of the outcome of interest is very small
 - 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?
- Take-up of bednets was as high with consumer loans as it was with free distribution
 - Despite increasing bednet purchases, microcredit did not increase usage of bednets
 - Neither microloans or free nets led to any measureable health improvements**
 - Microloans led to a 25% reduction in malaria episodes among households offered the loans
 - 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?
- Women were much more likely to use contraceptives and reduce unwanted births if they were seen separately from their husband**
 - There was no impact of increased access to contraception, suggesting high family sizes are optimal
 - Women needed to have their husbands present at the counseling sessions in order for the contraceptive intervention to have an effect
 - Women given contraceptives engaged in riskier sexual behavior
 - Don't know/never heard of this study.