
Today's News

Thursday, July 17, 2008

Dispute Over the Economics of File Sharing Intensifies

By [DAVID GLENN](#)

In 2004, two economists released a working paper with a contrarian message: Sharing music files over computer networks isn't to blame for the recent collapse in CD sales.

It was a provocative attack on an assertion that has loudly been made by the music industry. A revised version of the paper ran last year in the *Journal of Political Economy*, which by many measures is one of the top five economics journals.

But now a rival scholar says that some of the paper's arguments are transparently shoddy and that the journal has failed to alert its readers to those flaws. This month he has taken his campaign to the Web, posting (among other things) e-mail messages and internal referee reports from the paper's authors and one of the journal's former editors.

The paper's authors—Felix Oberholzer-Gee, now an associate professor at Harvard Business School, and Koleman S. Strumpf, now a professor of economics at the University of Kansas—had persuaded an open-source file-sharing network known as OpenNap to give them a huge cache of raw user data from the last several months of 2002. They then counted songs that were illegally downloaded and the sales of CD's that contained those songs. They looked especially at how CD sales varied during weeks when songs were more easily available on the file-sharing network (because European students were on vacation and therefore spending more time online). The paper's bottom line: File sharing had no net effect on sales.

The paper seemed like a model piece of empirical social science for the *Freakonomics* era. Unusual data source analyzed with "instrumental variables"? Check. Counterintuitive conclusion? Check. Implications for hot-button policy debate? Check. The scholars filed an amicus brief in defense of file-sharing companies in the U.S. Supreme Court's *Metro-Goldwyn-Mayer Studios Inc. v. Grokster* case in 2005. When a revised version of their working paper appeared in the February 2007 issue of the *Journal of Political Economy*, it was the lead article.

But ever since the working paper first appeared, Mr. Oberholzer-Gee and Mr. Strumpf have been hounded by a critic: Stan J. Liebowitz, a professor of economics at the University of Texas at Dallas. Mr. Liebowitz, whose research has often been financially supported by the record industry, faults the scholars for not sharing their OpenNap data, which would allow colleagues to replicate their analyses. Insofar as he can check the paper, he says, he finds its arguments to be weak and misleading.

Until recently, the three scholars' quarrel had been on a low simmer—a testy e-mail exchange here, an

awkward conference panel there. But in late May, the *Journal of Political Economy* rejected a critique that Mr. Liebowitz had submitted of Mr. Oberholzer-Gee and Mr. Strumpf's paper. (In economics parlance, critiques of previously published papers are known as comments.) The editor who sent the rejection notice was Steven D. Levitt, a professor of economics at the University of Chicago and the co-author of the best-selling *Freakonomics: A Rogue Economist Explores the Hidden Side of Everything*. Mr. Levitt's term as editor ended on June 1.

The rejection enraged Mr. Liebowitz for two reasons: First, he believes that he has demonstrated serious flaws in the file-sharing paper and that the journal owes it to its readers to discuss those flaws. Second, he discovered that the journal had used Mr. Strumpf himself as one of the two anonymous referees for his comment. (The second referee gave a positive report and advised that the comment be published.) Mr. Strumpf's role as referee was revealed when a reporter for the German newspaper *Handelsblatt* asked Mr. Strumpf to reply to Mr. Liebowitz's criticisms. He e-mailed to the reporter a response that was substantially identical to the negative referee report.

In recent weeks, Mr. Liebowitz has stepped up his campaign by posting new critiques of the file-sharing paper on the Social Science Research Network, along with the referee reports and other correspondence with Mr. Levitt and the two authors.

In an e-mail message to *The Chronicle*, Mr. Strumpf defends his work and suggests that Mr. Liebowitz's zeal stems from the fact that an academic center he directs, the Center for the Analysis of Property Rights and Innovation, receives grants from the Recording Industry Association of America and other commercial interests. "One might ask why Professor Liebowitz has remained so engrossed with our study," he writes.

Anonymous Self-Defense

Mr. Liebowitz says in an interview that he is no longer quite so angry about the journal's choice to use Mr. Strumpf as an anonymous referee. That decision was criticized by *Handelsblatt*, and it has been the subject of chatter on economics blogs. Mr. Liebowitz says that of course he expected Mr. Oberholzer-Gee and Mr. Strumpf to see his comment and to be given an opportunity to reply. But he is not pleased that the journal used one of them as an anonymous referee, particularly given that it used only two referees.

"When I first found out, I thought this was outrageous," he says. "But I've spoken to a few people who say that they've been in similar situations, and they've responded to comments anonymously in the third person. I think a lot of people would frown on this, but I guess it's not actually so rare."

The Chronicle contacted six editors at other highly ranked economics journals, and all said that the journal's behavior here was reasonably standard procedure.

Stephen Morris, a professor of economics at Princeton University and editor of *Econometrica*, writes in an e-mail message that when his journal receives comments on previously published articles, "we will often ask an author of the original article to write a referee report. If the original author/referee wishes to remain anonymous, we allow him/her to do so. I think it is fair to characterize this as standard practice in economics. We take into account the original authors' views, whether presented directly to the comment authors or as an anonymous referee report, but of course take into account their potential prejudice."

Dealing With Rejection

But even if it was kosher for the journal to use Mr. Strumpf as a referee, there is still the question of whether the journal acted wisely when it rejected Mr. Liebowitz's comment. The second referee report, after all, urged publication, saying that Mr. Liebowitz had correctly pointed out several flaws in the paper.

In an essay he posted two weeks ago, Mr. Liebowitz writes that Mr. Strumpf's anonymous referee report "fails to present a cogent defense, or virtually any defense for that matter," against the accusations made in the comment.

Mr. Strumpf replies that most of Mr. Liebowitz's criticisms are trivial, even if correct. "Almost every point raised in Liebowitz's various pieces," he says, "involves incidental points which are not central to our conclusions."

Mr. Liebowitz, however, says his objections cast doubts over the entire study. Because he has not been able to scrutinize the OpenNap data at the heart of the study, his criticisms are largely aimed at a few non-OpenNap-based auxiliary tests that are presented at the end of the paper. But those tests are so weakly constructed, he says, that they call into question the validity of all the work.

One of the tests has to do with seasonal variations in record sales. "The number of file-sharing users in the United States drops 12 percent over the summer ... because college students are away from their high-speed Internet connections," Mr. Oberholzer-Gee and Mr. Strumpf write. Because of that seasonal dip, we might expect to see higher CD sales during the summer. But summer CD sales, as a proportion of the full year's sales, have not risen during the file-sharing era. Ergo, there doesn't seem to be much relationship between file sharing and CD sales.

That all sounds logical enough. But Mr. Liebowitz points out that one of the argument's premises—"the number of file-sharing users in the United States drops 12 percent over the summer"—is profoundly misleading.

The authors' citation is to a report on monthly file-sharing usage prepared by BigChampagne, a company that measures traffic on peer-to-peer networks. Mr. Liebowitz has a copy of the same report, which covers the period from August 2002 through May 2006. It's true that summer file sharing drops by an average of 12 percent (11.7 percent, to be precise) during the three summers—2003, 2004, 2005—covered in that report. But all of that effect comes from a severe drop in the summer of 2003, during a much-publicized wave of industry lawsuits against file sharers. During the summer of 2004, file sharing was flat, and during the summer of 2005, it actually rose slightly. So the test Mr. Oberholzer-Gee and Mr. Strumpf have set up—which is based on the ratio of summer-to-full-year CD sales—tells us nothing, Mr. Liebowitz says.

"If one of my undergraduates did that, I would fail him," says Bruce D. McCullough, a professor of decision sciences at Drexel University who became interested in the dispute because he is a proponent of data transparency in economics publishing. "To take one decline, one flat, and one advance, and to suggest that it always goes down in the summer is just wrong."

Mr. Strumpf concedes that the only true summer decline in the BigChampagne data came in 2003, but he says that the summer-sales test has been "misconstrued." Because the *rate of increase* in file sharing drops during the summer, he says, we would still expect to see a boost in the summer share of CD sales, if file sharing injured CD sales. The summer share of sales has not risen, and Mr. Strumpf takes that as evidence that file sharing does not affect sales.

But it is still not clear that that is a sensible test for this particular data set. In the 2004 and 2005 BigChampagne data, the relatively flat summer usage levels are just blips in a much larger upward trend in file-sharing activity. (The number of file-sharing users almost doubled between January 2004 and December 2005.) In each of those two years, it is the January-through-March period, not summer, that has by far the lowest average file-sharing levels. So a more logical test here might be whether the *winter* share of CD sales is rising, relative to the pre-file-sharing era.

(Mr. Strumpf also points to usage records from Internet2, which indicate a stronger pattern of summer declines in file sharing. But he and Mr. Oberholzer-Gee did not cite Internet2 data in their paper.)

The summer-sales question is one of only a dozen similar criticisms that Mr. Liebowitz puts forward. Among other quarrels, he disputes the authors' arguments that worldwide CD sales have not been falling, that file sharing has not injured sales for any particular musical genre, and that changes in retail-inventory practices explain most of the decline in CD sales in the United States. In several cases, the authors have not replied to Mr. Liebowitz's criticisms, either in public or in Mr. Strumpf's referee report.

The Chronicle asked Mr. Strumpf to clarify how the paper analyzed sales of CD's in different musical genres. Here, Mr. Liebowitz's concern is that the paper appears to lump many albums into vague categories like "current" and "catalog," which correspond to certain *Billboard* charts but which, in Mr. Liebowitz's words, "are not really categories of music appealing to particular audiences." (The positive referee agreed with Mr. Liebowitz on this point.) In reply, Mr. Strumpf offered some plausible criticisms of an alternative method of analyzing genres presented by Mr. Liebowitz. But he did not directly explain or defend his own paper's genre analysis.

Hidden Data

At this point, the dispute has reached such a high pitch that Mr. Liebowitz has given up hope of ever seeing Mr. Oberholzer-Gee and Mr. Strumpf's OpenNap data.

He says that their reasons for not sharing the data have shifted over the years. In a 2004 e-mail message, Mr. Strumpf said that the raw user data was so sensitive that it might be subpoenaed in industry lawsuits and that his university's legal office had urged him not to share the data. (Mr. Liebowitz sympathizes with that concern but says that it would have been easy for the authors to strip away the users' IP addresses.) This year, Mr. Strumpf told *Handelsblatt* that he and Mr. Oberholzer-Gee had signed a confidentiality agreement with their OpenNap source that prevented the sharing of the data. Mr. Strumpf declined to show a copy to *Handelsblatt* or to *The Chronicle*.

Mr. Strumpf says that he has always been candid with Mr. Liebowitz about the impossibility of sharing the data. He showed *The Chronicle* an April 2004 e-mail message in which he told Mr. Liebowitz about both the legal concerns and about his promise to OpenNap not to distribute the data.

If Mr. Oberholzer-Gee and Mr. Strumpf's paper were submitted today, it is not clear whether it would be accepted by the *Journal of Political Economy*. In 2006, partly at Mr. McCullough's urging, the journal adopted transparency rules that require authors to post their data and statistical models on the journal's Web site. (The file-sharing paper was grandfathered in because it was submitted before the rule was announced.)

In cases in which authors use confidential or proprietary data, the journal sometimes makes allowances. "We deal with proprietary data on a case-by-case basis," writes Robert Shimer, a professor of economics at the University of Chicago and one of the journal's editors, in an e-mail message to *The Chronicle*. Mr.

Shimer says that he, personally, insists that authors with proprietary data must explain to their colleagues how they can purchase an identical data set. Since no identical data set from OpenNap is available, Mr. Oberholzer-Gee and Mr. Strumpf's paper would not clear that hurdle. (Mr. Shimer left that point implicit; he did not directly comment on this dispute.) But Mr. Shimer added that his practice is not an official journal policy and that other editors there might have different standards.

As far as Mr. McCullough is concerned, the authors' inability to share the data will forever cast a cloud over the paper. "If they are not making their data and code available, then I have to think that they have something to hide," he says. "There is a lot of nonreplicable research published in economics. We need to change the profession so that readers can expect that there actually does exist data and code that will reproduce the published results. Right now, at many journals, we cannot expect that."

But Mr. Strumpf insists that even if the OpenNap data set cannot be duplicated, other scholars can replicate his paper's general analysis by acquiring similar data sets from other file-sharing networks or from commercial vendors like BigChampagne.

He points to a 2004 paper by Tatsuo Tanaka, an associate professor of economics at Keio University, in Japan. Mr. Tanaka used data from a Japanese file-sharing network known as Winny, and he found results similar to those of Mr. Oberholzer-Gee and Mr. Strumpf: a limited effect, at most, on CD sales.

A Different Consensus

Which brings us back to the core question. Does file sharing decrease CD sales, or doesn't it? A widely publicized study released last year by the Canadian government found that file sharing actually increases CD sales. But that study is an outlier. A majority of economic studies have concluded that file sharing hurts sales, though often to a more modest degree than the record industry would like the public to believe.

David Blackburn, a senior consultant at NERA Economic Consulting, a New York-based firm, analyzed file sharing when he was a graduate student at Harvard five years ago. His analysis found that file sharing helps artists who are relatively little-known because downloaders have a chance to sample their music. "That only makes sense, right?" he says in an interview. "File sharing can't hurt your sales if you don't have any sales to begin with. All it can do is introduce more people to your music, which gives you a chance to make sales."

But Mr. Blackburn also found that top-selling, well-known artists take a substantial sales hit from file sharing—and that in terms of total CD sales, that negative effect outweighs the positive "sampling" effect for fledgling artists. "Most album sales are from ex-ante well-known artists," he says, "so the net impact on the industry is actually quite negative."

Mr. Blackburn says that he has long been fascinated by the argument presented in Mr. Oberholzer-Gee and Mr. Strumpf's paper. "When you think about it, that's a kind of shocking conclusion," he says. "And I'm not sure that they put forward a good explanation of why there would be no impact. In large measure, I think that's why there's been so much pushback against the paper."

[Copyright](#) © 2008 by [The Chronicle of Higher Education](#)

[Subscribe](#) | [About The Chronicle](#) | [Contact us](#) | [Terms of use](#) | [Privacy policy](#) | [Help](#)