The Death of a Technical Skill

John J. Horton
MIT Sloan and NBER

Prasanna Tambe*
Wharton, U. Penn

September 12, 2019

Abstract

In 2010, Steve Jobs announced that Apple would no longer support Adobe Flash—an at-the-time popular set of tools for creating Internet applications. In the years following Jobs' announcement, the use of Flash declined precipitously. However, using data from an online labor market, we show there was no detectable reduction in Flash hourly wages or even in the number of applicants per Flash job. We show that the reason wages stayed flat was that the negative demand shock for Flash was quickly met with a supply response: Flash specialists transitioned away from Flash, and new market entrants were less likely to specialize in Flash. A retrospective survey of affected Flash workers reveals them to be highly forward-looking, abandoning skills with no perceived future and picking up new skills, primarily through learning-by-doing. The implications for the spread of new technologies that require complementary labor are discussed.

*For the latest version, see `https://www.john-joseph-horton.com/papers/schumpeter.pdf`. We acknowledge valuable feedback from James Bessen and seminar participants at the NBER IT and Digitization Summer Institute, the INFORMS Conference on Information Systems and Technology, Cornell University, Arizona State University, the Rotterdam School of Management, the University of Texas-Austin, the University of Maryland, and the University of Utah. Thanks to Apostolos Filippas, Richard Zeckhauser, Chis Stanton, David Autor, Rob Seamans, Jacques Crémer, Adam Ozimek, Sebastian Steffen, Andrey Fradkin, Emma van Ingwen, Joe Golden, and Daniel Rock for helpful comments and suggestions. Prasanna Tambe is grateful for financial support from the Alfred P. Sloan foundation. John Horton is grateful for financial support from the Upjohn Institute. Thanks to Jeff Bigham and Elliot Geno for insight into the labor market for Flash skills. For a version of this paper with all code included as an appendix, see `https://www.john-joseph-horton.com/papers/schumpeter_with_code.pdf`.

1 Introduction

When the demand for a skill falls—or will foreseeably fall—workers with that skill must make a choice. If they choose to "exit," they can retire, learn a new skill, or work with another skill they already possess. If they choose to stay, they face whatever lower wages or worsened job-finding probability this choice (eventually) entails. Workers entering the labor market—but disposed to working with the declining skill—also have to make an analogous decision. The choices made by both incumbent and entrant workers presumably depend on factors such as their career time horizon, level of ability and current human capital. They also have to consider the anticipated course of the decline in demand, the costs of switching to a new skill, and so on. Workers might also consider the collective choices of all other similarly situated workers. In broad terms, this characterization of the human capital allocation problem is not novel (Ben-Porath, 1967; Chari and Hopenhayn, 1991; Jovanovic and Nyarko, 1996). However, there is relatively little fine-grained evidence on how individual workers adjust to a skill-specific demand shock. And yet it is the fine-grained details of adjustment that have implications not only for worker welfare, but also the diffusion of technology and rate of innovation (Acemoglu, 1998, 2010).

In this paper, we explore how the labor market for a technical skill responded to a negative shock in the demand for the associated technology. The skill that we study is Adobe Flash—a once-popular collection of software tools used for creating multimedia games, advertisements and applications delivered over the Internet. The decline in demand for Flash skills is widely attributed not to the emergence of a superior technology, but rather to a business decision made by Apple. On April 29th, 2010, Steve Jobs—the CEO of Apple at the time—published an open letter entitled "Thoughts on Flash" (TOF) which announced that Apple would no longer support Flash on iOS devices such as the iPhone, iPod, and iPad.[1] Despite Jobs' claims that the decision was made for technical reasons, this was viewed by many as a pretext—the "real" reason for this withdrawal of support was a desire for greater control over the experience on Apple devices, particularly the iPhone.[2] Even if Jobs' arguments had technical merit, it was clearly self-

[1]http://www.apple.com/hotnews/thoughts-on-flash/.

[2]The contemporaneous discussion of the announcement on Hacker News, a discussion site run by Y Combinator and known for insider takes on IT and the startup/tech industry—is illustrative, as it contains many arguments that Apple's choice was self-

serving for Jobs to make them at that particular moment in time.

Jobs' motivations and reasons notwithstanding, it is clear that this announcement was viewed by the developer community as an inflection point in the market for Flash development, and our data suggest it was an effective—albeit not immediately-acting—poison. It also appears that Apple's decision might have catalyzed other changes in the industry that hastened the demise of Flash. For example, in late 2011, it was announced that Flash would no longer be natively supported on the Android operating system, starting with the "Jellybean" release.[3] Although still used for some applications, the decline continued and in July 25, 2017, Adobe published a blog post announcing that they will remove all support for Flash by the end of 2020.[4]

The decline in Flash can readily be seen in a variety of data sources. Perhaps the best contemporary indicators of IT industry interest in a given technology are the questions being asked on StackOverflow, an enormously popular programming Q&A site. The left facet of Figure 1 shows the volume of questions per month for Flash and for comparison, a basket of very popular IT skills, all normalized to 1 in the TOF month.[5] The y-axis is on a log scale. We can see that Flash and our chosen basket of skills are growing more or less in lockstep pre-TOF, reflecting growth in the Q&A platform and the wider IT industry, but that after TOF-day Flash shows a strong absolute decline. There is some delay in the drop, likely reflecting the diffusion of the news of Apple's plans as well as the completion of already-planned projects.

To study how the decline in Flash affected workers specializing in Flash, we use data from a large online labor market (Horton, 2010). The decline in Flash is also readily apparent in the longitudinal data from this market: the right facet of Figure 1 plots the number of job opening posted per month for jobs requiring Flash skills and those requiring PHP (one of the "basket skills" from Figure 1). Both Flash and PHP are normalized to 1 for the TOF-month. We truncate the data to the start of StackOverflow (in 2008), even though the online labor market is considerably older. As we saw with the StackOverflow data, both Flash and PHP move closely together pre-TOF-

interested and made on flimsy technical arguments—see `https://news.ycombinator.com/item?id=1304310`.

[3]https://www.zdnet.com/article/adobe-ditches-mobile-flash-refocuses-on-html-5/

[4]https://www.theverge.com/2017/7/25/16026236/adobe-flash-end-of-support-2020

[5]HTML is a markup language for making websites; CSS is a language for styling websites and controlling their appearance; Java is a general purpose programming language, while PHP is another programming language frequently used for server-side scripting.

Figure 1: The decline of Flash on two platforms

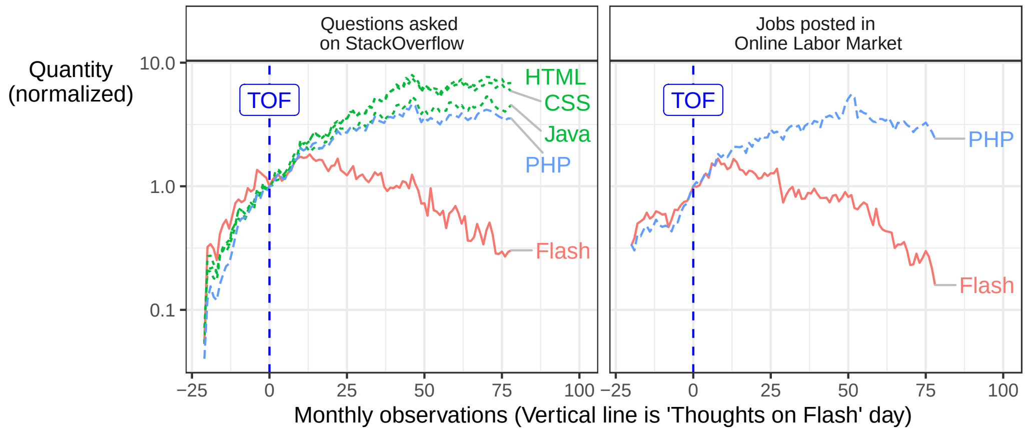

Notes: The left facet plots number of questions, by technical skill, posted on Stack Overflow between 2008 and 2016. The right facet shows the number of job posts requiring Flash or PHP in an online labor market. All series are normalized to have a value of 1 in the TOF-day period. The unit of observation is the month.

day and then diverge: following TOF-day, the number of Flash job openings began to decline relative to PHP, falling by more than 80% between 2010 to 2015.

As we will show, despite a large decline in the number of Flash openings posted, very little else about the market for Flash changed. There is no evidence employers were inundated with applications from out-of-work Flash programmers—the number of applicants per opening remained roughly constant.[6] There was no increase in the likelihood that Flash openings were filled, nor was there a reduction in the wages paid to hired Flash programmers. In short, despite a roughly 80% reduction in posted Flash jobs, we observe a reduction in the quantity of hours-worked, with no reduction in the price.

Big changes are visible when we look at individual Flash-specializing workers. We do see substantial movement away from Flash, as measured by counts of active Flash workers. As demand for Flash starts to fall, more workers begin exiting than entering the market for Flash, despite the platform itself growing rapidly. This early movement shows up in the time series of active Flash workers, with a net flow out of Flash even when the number

[6]We use the terms "worker" and "employer" for consistency with the economics literature, and not as a comment on the nature of the relationships created on the platform.

of jobs being posted was roughly flat. We also see this in the hours-worked per active Flash worker, which actually *increases* over time, contrary to the idea of a large pool of out-of-work Flash workers splitting a dwindling pool of jobs.

Market aggregates potentially mask compositional changes in the types of Flash jobs or the workers participating. For example, flat wages could actually be a real decline if the remaining Flash work is being done by the most experienced Flash workers. To understand the decision-making and outcomes of individual Flash workers, we constructed a panel of workers specializing in Flash prior to TOF-day, as well as a control group that were also active on the platform pre-TOF-day, but specializing in other non-Flash IT skills. Our empirical approach is essentially a difference-in-differences design, using the contemporaneous non-Flash developers as a control group. Rather than use all workers active on the platform as controls, we selected a sample with similar attributes at the time of the TOF-day. Because of the large number of potential "donor pool" workers on the platform, we obtain excellent balance on pre-TOF-day covariates, such as average wages, job search intensity, hours-worked, and experience.

Using the matched sample panel, we find no evidence that Flash workers were more likely to exit the online labor market post-TOF-day, making it unlikely that selection explains any of our results. As with the market-level time-series evidence, we find no reduction in the hourly wage. At the worker level, if anything, there is some evidence of an *increase* in the hourly wage for the most experienced Flash workers.

Our panel lets us explore when incumbent workers began moving away from Flash. We find evidence that workers specializing in Flash prior to TOF-day increased their job search intensity—sending more applications per month. Effects are concentrated among those workers who were more focused on Flash pre-TOF-day, as measured by the fraction of their hours-worked that were on Flash projects. In addition to increasing application intensity, Flash workers also shifted their job search *focus*, as measured by how similar an applied-to job is to the kinds of work the worker had completed in the past. We find strong evidence that they began moving away from their existing skill focus, applying to jobs that were more unlike jobs they had worked in the past. The movement away from Flash in job applications occurred when the fall-off in demand was still nascent and Flash openings were still fairly plentiful.

Despite Flash workers applying more broadly and more intensively—and seeming to obtain the same wage—we do find some evidence that total hours-worked (for jobs of any type) declined for Flash workers, which implies a fall-off in earnings, though these estimates are imprecise. If hours-worked did decline and it is not due to a change in preferences, it does suggest workers cannot work as many hours as they wish at their market wage, as we would expect in a fully competitive market. Despite very few traditional frictions in these markets, there is evidence that these markets have features inconsistent with a fully competitive characterization (Dube et al., 2019).

If the hourly wage for Flash work did not fall, why did incumbent Flash workers bother switching? One possible reconciliation of the data and economic theory is that Flash workers were infinitely elastic to the Flash market. At the level of the individual, this would seem improbable for a skill that requires a significant human capital investment. However, the supply curve could be *de facto* highly elastic if enough workers decide not enter and enough workers exit (to learn new skills) before a total collapse in demand. This kind of adjustment would be particularly effective if other skill markets are sufficiently large that Flash "refugees" are not likely to depress wages very much (Card, 1990).

To gain more insight into the adjustment process, we surveyed a sample of Flash workers that had been active on the platform at the time of TOF-day. Surveyed workers perceived Flash's dim future contemporaneously, and felt that they needed to switch to other skills. Respondents reported that their primary adjustment strategy was to learn new skills, and they did so by taking on projects in skills they wished to learn. This "earn while you learn" strategy (Tambe et al., 2018) was the most important method and was regarded as substantially more important than more traditional approaches, such as reading books and taking classes. Our results support the view that learning-by-doing is particularly important for new technologies (Bessen, 2003, 2016).

Surveyed workers reported a willingness to lower rates to obtain work in new skills—a margin of adjustment reported in other contexts but which might be particularly important in IT. In the IT sector, there is typically no additional physical capital required to acquire some skill, remote work/offshoring is comparatively easy and change is too fast for formal education to play much of a role (Barley and Kunda, 2011; O'Mahony and Bechky, 2006; Ang et al., 2002; Mithas and Krishnan, 2008). This feature of IT—rapid technological change the makes skills obsolete quickly—has recently been highlighted by Deming and Noray (2018).

A key finding of the paper—reinforced by both the survey evidence and the archival platform evidence—is that when demand fell, forward-looking workers move to other skills. This provides an explanation for why there was so little reduction in Flash wages at the market level—the negative demand shock was offset by a supply response. For a skill that requires substantial human capital—and with that human capital mostly acquired through on the job training—the future market for a skill matters to workers and strongly affects the present. We formalize this idea in a simple model that shows how wages for a dying skill can stay flat or even increase.

Our results have implications for understanding technology adoption. If both the demand and supply sides of a skill-specific labor market shift contemporaneously—as our evidence implies they do—there could be a coordination problem between firms adopting IT and workers trying to build the required human capital. Consistent with this view, in reporting which skills to move to post-Flash, workers report being sensitive to the risk of choosing the wrong skill in which to specialize—a theme that appears in several models of human capital formation (Krebs, 2003; Guvenen et al., 2014). To avoid this risk of picking the "wrong horse," survey respondents reported that they primarily relied on information from other programmers, technical discussion boards, and industry leaders to reduce this risk. Many looked to large technology companies for signs that a new technology is likely to become a standard, suggesting long-term demand for the skill—highlighting another benefit to standard-setting (Simcoe, 2012).

This is the first paper of which we are aware that uses longitudinal fine-grained wage, hours, and application data to study how a market adjusts when an information technology become "obsolete." Although there is a large literature on workers responses to adverse shocks (Jacobson et al., 1993; Kletzer, 1998; Couch and Placzek, 2010; Nedelkoska et al., 2015), our setting has several features that are advantageous when exploring human capital decision-making. Because of the online setting, the shock is "pure" without the other associated equilibrium effects that might occur from more geographically localized shocks (say from a plant closing), or that are confounded with other effects of the business cycle (von Wachter et al., 2009; Davis and von Wachter, 2011; Kahn, 2010). The project-based, relatively short-term nature of the work implies any losses from skill obsolescence are likely not due to firm-specific human capital (Neal, 1995) or even firm-specific rents (Goldschmidt and Schmieder, 2017). Our paper is most conceptually similar to (Edin et al., 2019), which looks at how individual outcomes are correlated with broad declines in occupations. They find surprisingly small reductions in earnings—far smaller than found from specific worker displacements—consistent with our results.

The role played by Steve Jobs is useful, as it is less likely that the fall-off in hours-worked is due to selected workers that made poor human capital choices *ex ante*. Despite Jobs' arguments having some technical merit, the Walter Isaacson biography of Jobs discusses the extent Jobs was motivated by personal grudges against Adobe that were completely unrelated to Flash (Isaacson, 2011).

Much of the existing literature on technological change and labor markets focuses on longer time scales than our study (Chin et al., 2006; Goldin and Katz, 2007). If technical change takes longer, we might expect different modes of adjustment, such as changes in the productive process with respect to broad types of labor inputs (Bresnahan et al., 2002; Autor et al., 2003; Michaels et al., 2014) or changes happening only at the generational level (Goldin and Katz, 1999). In our setting, the decline of Flash occurred within the career horizon of a large number of workers and so "within-worker" adjustment was necessary. This kind of rapid change is currently of significant interest due to the anticipated effects of emerging technologies—such as AI and sharing economy platforms—on the workforce (Filippas et al., forthcoming; Chen et al., Forthcoming; Hall and Krueger, 2018).

The labor market for technical skills is of particular interest, as this area is often characterized as having "shortages" despite the high wages. However, Deming and Noray (2018) show that the returns to STEM majors/occupations declines considerably over time, due to technological obsolescence. Our setting is an example of the key point Deming and Noray (2018) make—STEM occupations are characterized by rapid technological change, with workers having to adjust frequently and old human capital becoming less valued. This has implications for entry into STEM, as well as the micro-details of adjustment to these shocks, which is our focus.

Our setting does have limitations. Our samples are not the universe of all Flash workers or IT workers, and most of analysis is conducted in the context of a particular marketplace rather than *the* market (Roth, 2018). However, our evidence suggests that changes in the broader IT market are reflected in the marketplace. Existing data sources that would be representative of the whole labor market would not be appropriate for our research questions—none of the phenomena we explore would be detectable at the BLS level of occupation—our treatment and control would both be "software developers" or "computer programmers."[7] It is inevitable that statistical agency data collections elide over worker and job opening differences important to market participants, as the degree of specialization in a modern economy demands it. But this limitation of the data should not prevent us from realizing that matching and economic decision-making is often happening at finer scale that we can typically observe (Marinescu and Wolthoff, 2016) and that case studies at the right level of granularity can be more useful than "representative" samples.

The rest of the paper is organized as follows. Section 2 discusses the Adobe/Apple Flash controversy and the empirical context of study. Section 3 presents descriptive statistics on the Flash market. Section 4 discusses the economics of market adjustment to a fall-off in demand. Section 5 presents the outcomes and choices of individual workers. Section 7 presents our survey evidence. Section 8 concludes.

2 Empirical context

Flash applications can be technically complex, typically requiring a mix of programming, graphic design, and other complementary skills. A would-be Flash programmer must learn the underlying programming language, ActionScript, as well as Adobe's Flash authoring tools and best practices for building and debugging Flash applications. Even for programmers who are experienced in other programming languages, acquiring Flash skills would be a non-trivial investment, requiring months or even years of sustained effort.

To give a sense of the size of the human capital investment in learning Flash, the last published Flash user's guide is over 527 pages, and the full documentation (made up of HTML files) is nearly 17 *million* words. There are hundreds of books on Flash still listed on Amazon.[8] Given the potential size of

---

[7]See `https://www.onetonline.org/link/summary/15-1131.00`. See Bound et al. (2015) for an example of work that is possible on the IT labor market using the BLS level of occupation.

[8]As of June 22, 2019, there are still over 350 books returned on Amazon.com for the query 'Adobe Flash': `https://www.amazon.com/s?k=Adobe+Flash`https://www.amazon.com/s?k=Adobe+Flash. The user's guide is here: `https://www.adobe.com/support/documentation/archived_content/en/flash/cs3/flash_cs3_help.pdf`. The documentation word count was obtained by downloading `http://help.adobe.com/en_US/FlashPlatform/reference/actionscript/3/PlatformASR_Final_en-us.zip` and running: `find . -name *html -type f -print0`

the Flash human capital investment, it is not surprising that many developers focusing on Flash described themselves as "Flash developers" specifically and expressed anger about the decline of Flash in our survey.[9]

2.1 Apple's announcement

On April 29th, 2010, Apple's CEO Steve Jobs published an open letter, "Thoughts on Flash," announcing Apple's decision to no longer support Flash applications on the iPhone, iPad or iPod Touch. Immediately following the release of Jobs' letter, shares of Adobe fell more than 1%.[10] While Jobs claimed that technical considerations were driving Apple's decision-making, it was widely believed that these arguments were a pretext and that Apple wanted to kill Flash because it was unwilling to concede so much control over the user interface and device performance to third parties.[11] Tony Bradley, of PC World, wrote:[12]

> It boils down to Apple wanting to maintain tight, proprietary control over app development for the iPhone and iPad, and not wanting to share the pie. It also seems suspicious given Apple's foray into mobile advertising with the iAd platform competing directly with the fairly ubiquitous Flash-based ads.

In our survey of Flash workers, many respondents mentioned Apple's role in Flash's demise unprompted:

> "since Steve Jobs' infamous speech about Flash, there has been a decline in the use of Flash across the web"; "people are more interested in other laguages [sic] due to the reason that Flash is not supported by iphones and Ipads etc."; "Steve Jobs not letting

`| xargs -0 wc | tail -n 1` on the unzipped folder.

[9] One respondent wrote: "steve jobs wrote a nasty little hit piece before he died and that really was a turning point. before that me and my partner were working for big clients, making good money, even working with Adobe itself. hahaha he ruined everything. thanks for nothing steve!"

[10] Apple's Jobs slams Adobe's Flash technology.

[11] For example, "Steve Jobs is Lying About Adobe Flash", `htttp://www.businessinsider.com/steve-jobs-is-lying-about-flash-2010-4`, Business Insider, April 30th, 2010. and Adobe Hits back at Apple's "smokescreen", The Telegraph, April 30th, 2010.

[12] Apple v. Adobe: Something Just Doesn't Add Up

> Flash platform on Apple devices is the primary reason [for the decline].”; “steve jobs wrote a nasty little hit piece before he died and that really was a turning point.”

In the years following Jobs’ letter, the popularity of Flash waned considerably, though the effect was not immediate nor universally anticipated: on August 26th, 2010 the Wall Street Journal touted rising demand for Flash workers—a “trend” that seems to have had no empirical basis.[13] Today, Flash is a moribund technology confined to a small number of niche applications and will no longer be supported by Adobe at all by the end of 2020.[14] Even legacy uses of Flash are likely to become inoperable soon for most users, as modern browsers are beginning to block Flash applications by default—a development ghoulishly reported by the tech blog Gizmodo as “Google sticks another knife in Flash’s Corpse.”[15]

2.2 An online labor market

Our primary empirical setting for studying the demise of Flash is an online labor market. In this market, employers post job openings to which workers can apply. Employers can solicit applications by recruiting workers, or workers can just apply to openings they find. Employers then screen applicants and potentially make a hire. If a hire is made, the wage is observed, as well as the number of hours worked if the job was an hourly job. On the platform, hours-worked and earnings are measured essentially without error, as workers use a kind of digital punch clock to record hours.

Employers self-categorize job openings depending on the nature of the work. Employers also label each job opening with up to ten skills, which must match a “controlled vocabulary” of skills maintained by the platform. The controlled vocabulary of skills contains several thousand distinct skills, and new skills are frequently added (Anderson, 2017). As expected given the platform’s focus on technical work, many of the skills are programming languages, tools, and software frameworks, e.g., Flash, PHP, C, HTML, etc.

The skills added by an employer are used by workers to find jobs that match their skills. Would-be applicants use these listed skills—as well as the

[13] “Flash Back: Demand Up in Engineering Specialty,” The Wall Street Journal, August 26th, 2010.

[14] https://www.theverge.com/2017/7/25/16026236/adobe-flash-end-of-support-2020

[15] `https://gizmodo.com/google-sticks-another-knife-in-flashs-corpse-1836840400`

category of the job opening—to decide which jobs to apply to. Workers also list skills in their profiles. Employers can see the details of past jobs completed by the applicant, which are labeled with the skills selected by the original employer.

The extensive skill labeling on the platform allows us to characterize a worker's skill focus—particularly whether they used Flash. We also use the skill labeling to characterize a change in skill focus. To do this, we take all jobs that workers had prior to TOF-day and compute a vector of "shares" of skills. For example, if the worker had one project that used skill A and B and another that used B and C, we would compute their share vector as $(1/4, 1/2, 1/4)$ for (A, B, C). As we observe applications, we can then characterize how "close" an applied-to job is to a worker's skill history by taking the dot product of the skill vector for that job with the worker's share vector. We call this dot product "application similarity." For example, if the worker applied to a job that required skill B alone, the value would be 1/2 and if they applied to a job requiring D and E—skills they have not worked with at all, the value would be 0. As such, a lower dot product means the worker is applying to jobs farther "away" from his or her historical focus.

2.3 Representativeness

A natural question about our empirical context is whether it is representative of the broader labor market for technical skills. Figure 1 showed a close relationship between our platform and StackOverflow with respect to Flash, though of the skills being references on StackOverflow itself might be a selected sample compared to all IT skills being used in the wider industry. One way to assess representativeness—at least with respect to waxing and waning industry interest—is to compare job openings to search engine query volume. Given the ubiquity of Google, it is likely to be the best available source to measure interest economy-wide, providing something like "ground truth." Baker and Fradkin (2017) validate the idea that useful "real world" labor market information can be provided by Google search data. Unfortunately, this approach is challenging with "Flash" because [the] Flash is also the name of a popular DC Comics superhero. As one might imagine, superhero-related queries dominate software development queries on Google.

As a substitute skill to assess representativeness but still using Google search data, we use FBML (Facebook Markup Language), a now deprecated language for making Facebook applications. In relatively quick succession, the

Figure 2: Assessing representativeness by comparing FBML-related volume on three platforms

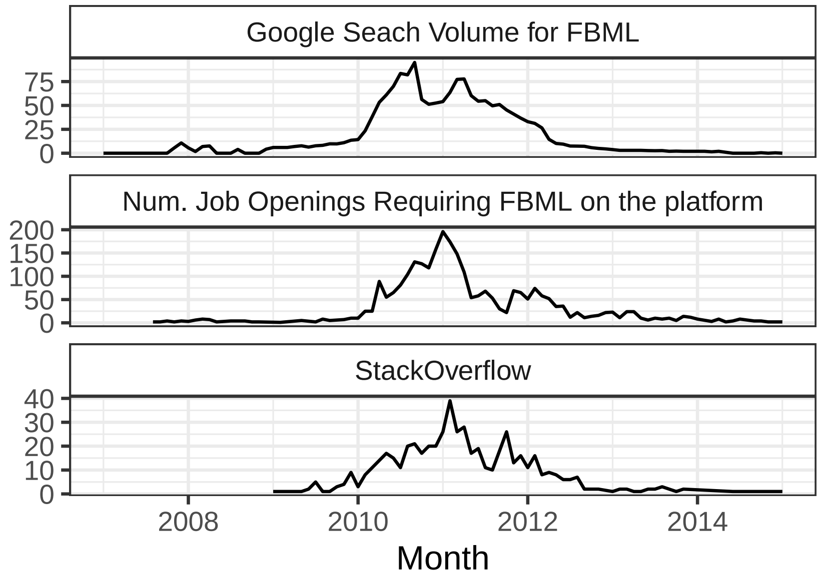

Notes: This figure shows the relative search volume for FBML on Google (top facet), job posts in an online labor market requiring FBML (middle facet), and number of questions on StackOverflow tagged with FBML (bottom facet). The unit of time is months.

technology was introduced, gained popularity and then was replaced with a different Facebook technology. Both the rise and fall of FBML occurred during the time frame we have online labor market data.

In Figure 2, we plot (1) the normalized search volume on Google for FBML, (2) the number of jobs requiring FBML on our online labor market, and (3) the number of questions asked on StackOverflow tagged with FBML. This comparison suggests that all the measures move together quite closely. The waxing and waning of the market for technologies appears to be well captured by the activity on the online labor market as well as on StackOverflow, at least for this one example where the two indicators of market activity can be cleanly measured.

3 The market for Flash skills

Simply plotting Flash market attributes over time illustrates some features of the decline in Flash. We first plot monthly time series about job openings that required either Flash or ActionScript. For comparison purposes, we do the same for PHP. PHP is a long-popular server-side scripting language for creating web applications (e.g., Facebook was originally written in PHP). It makes for an attractive comparison technology, as it was popular near the start of our time series and continued to be widely used up until the end of our data. It is also historically the mostly commonly-requested skill on the platform with, by far, the largest volume of hours-worked, making it less likely that any Flash "refugees" are affecting PHP measures. As not everything of interest can be measured at the job opening level, we also construct supply-side measures that allow us to characterize the flow in and out of Flash and the number of hours-worked per active worker.

3.1 Attributes of posted job openings

A number of monthly measures at the job opening level are plotted as indices in Figure 3. The indices are scaled such that they have mean value of 0 before TOF-day for both Flash and PHP.

In the top facet of the figure, the log number of openings is plotted, recapitulating what we saw in Figure 1. We can see the clear decline in Flash openings following TOF-day. The second facet from the top is the mean of the log of the number of applications per opening. The number of applications submitted for each PHP and Flash opening track very closely even post TOF-day. There is no evidence that Flash jobs were over-subscribed by out-of-work Flash specialists.

Despite no apparent post-TOF-day change in the number of applicants, perhaps Flash workers started bidding less, allowing more openings to be "filled"—defined as the employer hiring at least one worker. This does not appear to be the case—in the third facet from the top, the measure is the fraction of job openings that are filled, and if anything, the fill rate for Flash jobs seems to decline relative to PHP post-TOF-day. Again, this is opposite of what we would expect if the decline in Flash was regarded primarily as a negative demand shock only.

The conjecture that Flash workers started bidding less also does not seem to be borne out, at least as measured by the wage of the hired worker. In the

Figure 3: Attributes of Flash and PHP markets over time on the platform

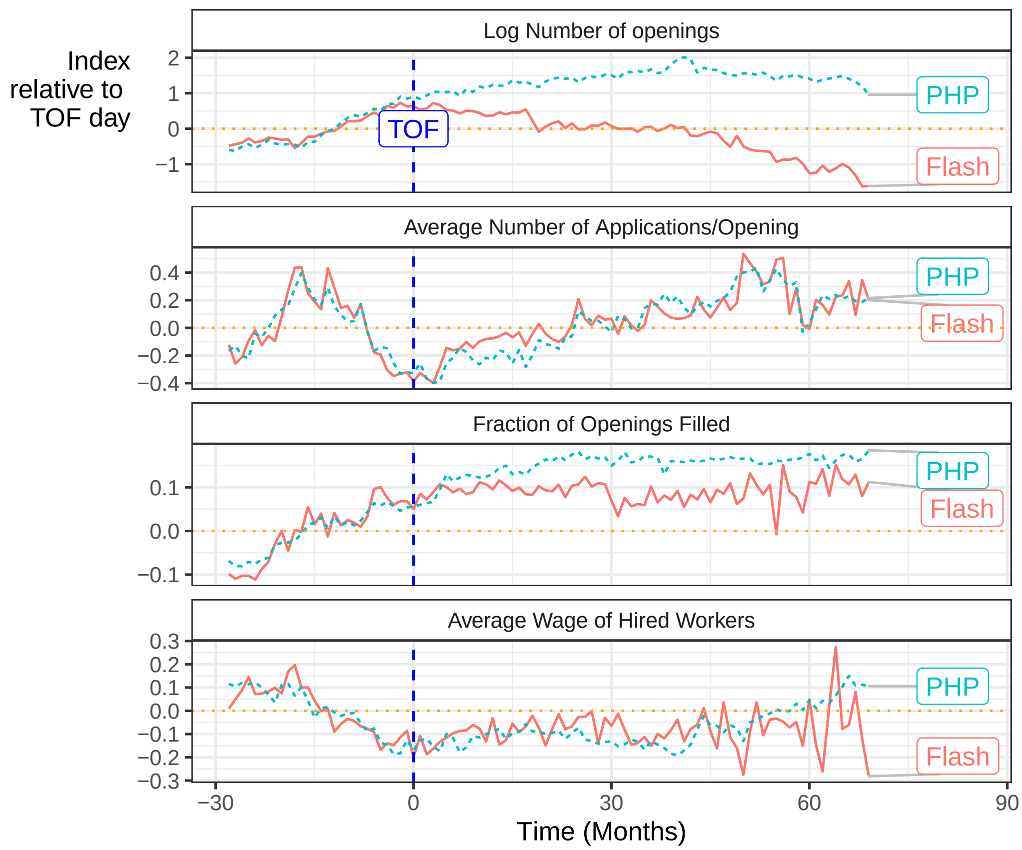

Notes: The figure shows the a variety of per-month attributes of job openings that require either PHP or Flash. The blue line indicates TOF-day.

bottom facet, we can see that average wages stay about the same relative to PHP. Near the end of the data, there are not many observations and so the estimates of the hourly earnings rate grows imprecise.

3.2 Attributes of workers

We constructed a monthly panel of Flash and non-Flash workers. There are a total of 324,097 workers in the panel. Of these, the number that were active before TOF-day and had some Flash experience is 1,871. Flash workers are defined as those that had earned some amount of money on a Flash project before TOF-day. By TOF-day, they had collectively worked 248,491 hours on Flash projects. Some workers joined the platform post-TOF-day but still worked in Flash. A total of 5,640 workers worked at least some hours on Flash projects during the period covered by our data.

We already plotted the number of job openings requiring Flash in Figure 3, but not the actual quantity of Flash hours-worked delivered. With our worker panel, we can plot hours-worked. In Figure 4a, in the bottom facet, we plot the total number of hours-worked. We can see it peaks right near TOF-day and then declines. In the middle facet, the output is the total number of Flash workers active that month, with "active" defined as working at least some number of hours. It also peaks near TOF-day, but we can see that it seems to decline more steeply compared to the total hours-worked plot.

A greater relative decline in workers than hours is borne out in the top facet of Figure 4a, which plots hours-worked per active worker. We can see that it was about 40 hours per-month at TOF-day but then rises post-TOF-day, rising to nearly 60.[16] In contrast to a "lump of labor" conception, in which existing projects are split over more workers, we instead see greater concentration.

To explore the flow of workers in and out of Flash, we calculate the first and last months a worker worked on a Flash project. We then plot the number exiting (bottom facet), entering (middle facet) and the net flow (top facet) in Figure 4b. Prior to TOF-day, we can see that the flow is positive and that after TOF-day, it turns negative shortly after. The flow out reaches a maximum about 5 months after TOF-day.

The patterns in Figure 4 suggest that a change in Flash hours was better

[16] We do not plot early values of this ratio, as the market was relatively small and the estimate of hours-per-worker is imprecise.

Figure 4: By-month grand means for all Flash workers

(a) Hours-worked

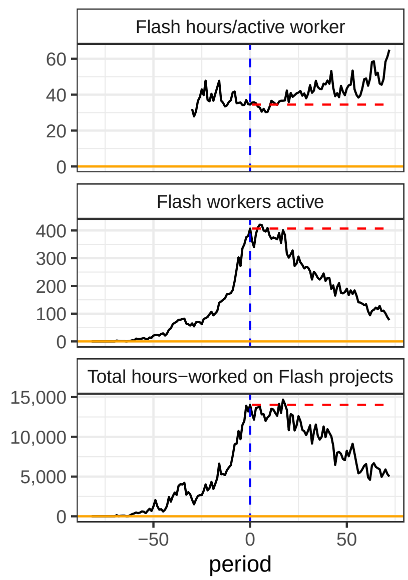

(b) Entry/exit

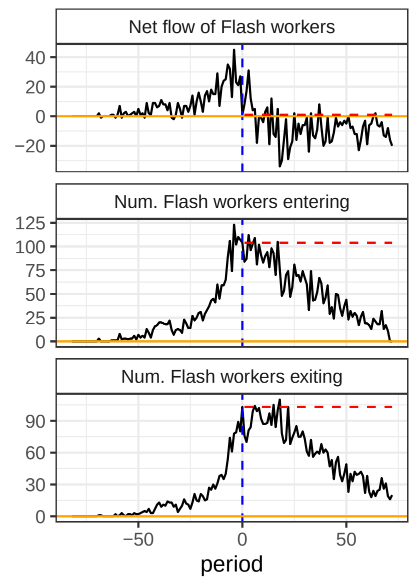

Notes: This figure shows the weekly time series of a number of attributes of workers active in Flash. The blue line indicates TOF-day. The dotted red line indicates the mean value of the attribute at TOF-day.

explained by some workers entirely abandoning Flash projects and other workers staying with it, rather than an even decline in Flash activity across all Flash workers. As we will show later, our survey evidence supports this notion of heterogeneous responses, with some respondents reporting no change in the amount they work in Flash, while others reporting entirely abandoning Flash.

4 Economics of a decline in demand for a skill

Having presented some of the key stylized facts of how the market for Flash skills evolved, we take a step back and discuss the economics that could explain the facts uncovered. If we assume a single labor market for Flash skills and set aside any elaborations about dynamics, a decrease in demand should not increase the quantity of hours-worked and not increase wages. At least one—hours-worked or wages—should fall, though the amount depends on the elasticity of the supply curve. Empirically, we observe a decline in hours-worked (bottom facet of Figure 4b), but no decline in wages (bottom facet of Figure 3). This outcome could occur if workers are infinitely elastic in their labor supply to the Flash market.

A near infinite labor supply elasticity seems improbable for a skill that requires a significant human capital investment—it would imply that Flash specialists have some other non-Flash skill that they are just as productive in and can readily switch to. This is hard to square with there being gains to specialization, though it .

As an alternative, consider a market supply curve consisting of the schedules of a collection of workers of different "vintages" planning to work for multiple periods. These workers value experience in the present for the human capital it imparts, which will be valuable in the future. With this dynamic perspective, a negative demand shock could be met with a supply response other than movement along the supply curve: would-be entrants switch to other skills with more of a "future" and some incumbents with a long-enough horizon exit as well. This could create a highly elastic market supply curve—or even shift the curve in so wages actually rise—despite costly switching and workers having "standard" labor supply elasticities to the labor market as a whole. We make this argument more precise with a simple model, which bears some similarities in spirit to Artuç et al. (2010), which estimates of model of switching costs inferred from labor flows in response to trade-shocks.

4.1 A simple model of skill obsolescence and transition

Workers use some skill to produce an output. Workers live two periods and have a discount factor of $\beta < 1$. In the first period, they are still learning the skill and have productivity $y < 1$, and so get paid yw, where w is the market wage in efficiency units. If incumbents stick with that skill for the second period, they have productivity 1 and get the "full" wage w. We will refer to workers that have already worked a period as "incumbents" and those first entering the market as "entrants." If an incumbent switches skills for their second period, he or she only has productivity y in that new skill, and never gets to be experienced. We assume there is an intensive margin elasticity of supply: $S(w) = wn$ where n is the measure of units of productivity in that market (i.e., each entrant worker provides y and each incumbent worker provides 1). A worker providing any amount of output is sufficient to raise his or her personal productivity from y to 1.

There are two skills, A and B. The A skill market is small, and the B skill market is vast, in the sense that if all A workers joined the B skill market, the wage in the B skill market would not change. Each market has completely inelastic demand.

The two markets have been in a steady state, with a measure of 1 workers joining the A market each period. If both incumbents and entrants are in a skill market, they have aggregate productivity $n = 1 + y$. Demand in A was 1 prior to the shock. Market clearing requires $S(w_A^0) = w_A^0(1 + y) = 1$, where w_A^0 is the pre-shock wage. For entrants to be indifferent between A and B, pre-shock, as required in a steady-state, $w_A^0 = w_B$, where w_B is the pre-shock wage in B and so $w_B = 1/(1 + y)$.

From a steady state, demand in A then has a foreseeable shock, in the sense that workers can see two periods of demand coming up: the demand will be d in the next period, with $d < 1$ and 0 the period after that. The wage in A when demand in the period is d will be w_A^1, which is endogenous. In the period after that, the market will no longer exist.

Would-be A entrants have to decide whether to enter A and earn yw_A initially, and then $\beta y w_B$, (working a period in A and then moving to B) or just start their careers in B, earning yw_B and then βw_B. Incumbents in A have to decide whether to stick with A and earn w_A or switch to B and earn yw_B. It is clear that if entrants choose to enter A in spite of the impending fall-off, it must be because w_A is higher than yw_B.

First consider the scenario where some entrants still choose to enter A and all incumbents stick with A. Let the fraction of new entrants choosing A to be x, with $1-x$ going to B, but because B is vast, w_B is unchanged by the influx. Market clearing in A is now $w_A^1(1+yx)=d$. As all new entrants have the option of joining B, incentive compatibility requires that entrants are indifferent, or $yw_A^1 + y\beta w_B = yw_B + \beta w_B$, which we can re-write as

$$\frac{w_A^1}{w_B} = 1 + \left(\frac{1-y}{y}\right)\beta.$$

Let $\overline{w_A}$ be the A skill wage that satisfies this constraint. Note that $\overline{w_A} > w_B$, implying that wages *rise* in the dying skill.[17]

Now consider a scenario with a larger shock in demand. All new entrants switch to the new technology, but no incumbents switch. However, because employers are completely inelastic, the wage in A falls until incumbents are indifferent between switching or staying, i.e., $w_A^1 = yw_B$. Note that the larger the y i.e., the smaller the cost of learning the skill, the less the wage declines in the dying skill. This creates a discontinuity in the A wage—before it was pinned down by the new entrants; now it is pinned down by the incumbents.

Now consider an even larger shock, such that at least some incumbents have to exit for the market to clear. Let z of the incumbents switch and all the new entrants choose B. Market clearing in A is $zw_A^1 = d$. As all incumbents have the option of joining B, incentive compatibility implies that $w_A^1 = yw_B$. Let $\underline{w_A}$ be this wage in A that satisfies incentive compatibility.

To summarize, there are three possible equilibria that depend on the size of shock:

1. For $d > \overline{w_A}$, the wage will be $\overline{w_A}$ and some entrants still choose A, getting a premium in their first period.

2. For $\underline{w_A} < d < \overline{w_A}$, all entrants choose B but all incumbents stick with A. The market wage is $\underline{w_A} = yw_B$, or the wage of the novice B worker.

3. For $d < \underline{w_A}$, even some incumbents choose to switch, but with the wage also staying at the novice B worker level, $\underline{w_A} = yw_B$.

The possible equilibria as a function of the size of the demand shock are shown in Figure 5. The x-axis is the quantity of labor in the A market (in

[17] Though we should expect this effect to be tempered by demand curves being sloped rather than vertical, as we assume.

efficiency units) during the d period; the y-axis is the wage in the A market during the d period. The wage in the B market is w_B is indicated and is $1/(1+y)$. The two relevant wages, $\overline{w_A}$ and $\underline{w_A}$, are plotted as horizontal lines. If the demand level is greater than $\overline{d}$, the market clears by only x of the would-be entrants choosing A. Below $\overline{d}$ but above $\underline{d}$, the market clears by all incumbents staying in the market but reducing their output on the intensive margin and the wage being at $\underline{w_A}$. For levels of demand below $\underline{d}$, the market clears by some incumbents choosing to exit.

As β increases (i.e., workers value the future more), we have a higher $\overline{w_A}$, meaning entrants need a larger premium in A to not switch. This also means that the region where entrants are indifferent shrinks. As the $\underline{w_A}$ is pinned down with incumbents that have no future, β is irrelevant for very large demand shocks. Although the quantity is increasing in d, the wage only changes once, at $d = \overline{w_A}$.

For a larger y, $\overline{w_A}$ decreases, but $\underline{w_A}$ increases. The two different effects reflect the different perspectives of incumbents and entrants. For entrants, a higher y means the forgone benefit of being experienced in B in the next period is lower, and so a relatively lower w_A wage is needed to keep the worker indifferent. For incumbents, a higher y means jumping to B is less costly, and so a higher wage in A is needed to keep them from moving.

Returning to our Flash context—the fact that we observed large reductions in hours-worked but flat wages—and some evidence of continued entrance by workers post TOF-day—the data seem consistent with the "small shock" equilibria i.e., $d > \bar{d}$. Of course, the actual situation is not the stylized two period world of the model, and the wages in the "B" market are relevant to the adjustment decisions of Flash incumbents.

In the model presented, gaining skills in A does not help in B—this is surely a over-simplification, particularly in our empirical context, where knowledge in Flash is likely to teach programming and graphic design skills more generally. There are also language similarities between ActionScript and JavaScript, a technology that is at the core of the modern web.

Instead of our stark "pick a single industry" assumption, we could allow workers to split a unit of time between A and B, spending t_A in A and $1-t_A$ in B, leading to human capital in the second period of z_A and z_B respectively. Work in one skill would translate into human capital at a rate $1-y$ in that skill, and a rate $\alpha < 1-y$ in the "other" skill, under the assumption that the best experience in a skill is actually working in that particular skill. This

Figure 5: Demand shocks and possible equilibria

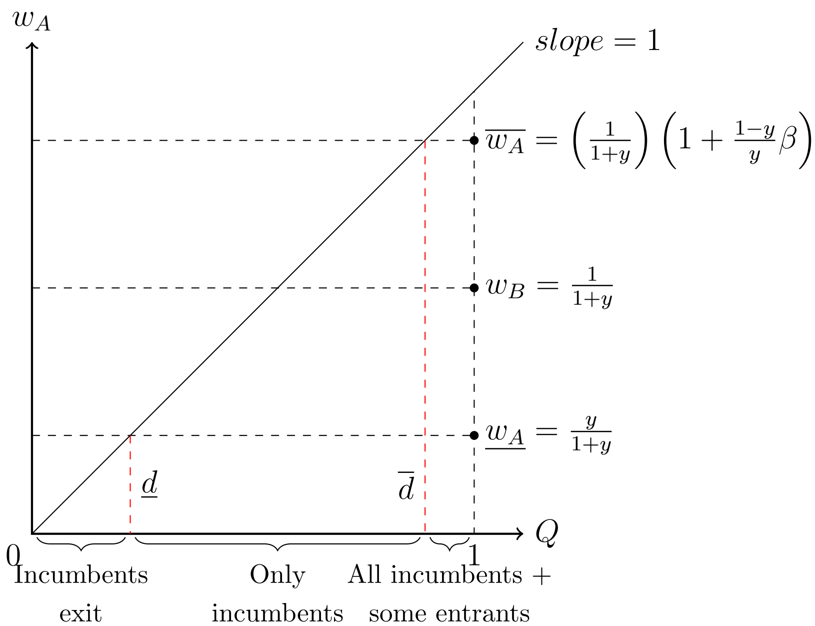

Notes: The x-axis is the quantity of labor in the A market during the d period; the y-axis is the wage in the A skill market during the d period. The wage in the B market is w_B is indicated and is $1/(1+y)$. The two relevant wages, $\overline{w_A}$ and $\underline{w_A}$, are plotted as horizontal lines. If the demand level is greater than $\overline{d}$, the market clears by only x of the would-be entrants choosing A. Below $\overline{d}$ but above $\underline{d}$, the market clears by all incumbents staying in the market but reducing their output on the intensive margin and the wage being at $\underline{w_A}$. For levels of demand below $\underline{d}$, the market clears by some incumbents choosing to exit.

would create an investment frontier, as in Murphy (1986), defined by

$$z_A = (1-y)t_A + \alpha(1-t_A) + y$$
$$z_B = (1-y)(1-t_A) + \alpha t_A + y.$$

As in Murphy (1986), risk neutral workers would still pick a corner solution of specializing in A or B, but the higher the value of α, the lower the cost of transition—at $\alpha = (1-y)$, learning would be the same in both, making it irrelevant which skill was chosen.

5 Outcomes and choices of individual workers

We now switch our focus away from the Flash market as a whole and explore the outcomes of individual workers active prior to TOF-day, which we can think of as the "incumbents" in the model. We select those workers from the panel we constructed for Figure 4 who had at least 40 hours of on-platform experience prior to TOF-day and at least some number of those hours were spent working with Flash. For these workers, we construct a binary indicator, AnyFlash. We also record the fraction of their total hours-worked that were with Flash, which is FracFlash.

Our goal will be to compare the trajectories and choices of these workers with Flash experience to a counterfactual group of workers not specializing in Flash. As a comparison group, we could simply use all active workers, but given that the market is fairly strongly divided into technical and non-technical work with large associated differences in hourly earnings (Horton, 2017), this approach would tend to mix workers that are not truly comparable. Instead, we use a matching approach first to construct a sample of comparable workers, then proceed with a standard panel analysis, using both a worker/time fixed-effects specification in both a "static" and distributed lag specification (Borusyak and Jaravel, 2016). We also estimate dynamic panel data models, allowing worker outcomes to depend on based values of the outcome, which give long-run effects very similar to the static specifications. This analysis is an Appendix B.3.

5.1 Identification and the nature of the treatment

To identify the effects of TOF-day on Flash workers, we use non-Flash workers as counterfactuals for all outcomes. Our baseline identifying assumption is that, conditional upon worker- and time-fixed effects, Flash workers are observationally equivalent to workers in the control group. This assumption is analogous to the one made in the literature on the effects of job loss (von Wachter et al., 2009; Couch and Placzek, 2010; Jacobson et al., 1993). In our setting, because we have a measure of Flash concentration, we can allow effects to differ based on how focused workers were on Flash pre-TOF-day. This is in contrast the displacement literature where job loss is binary.

If we can obtain a sample of counterfactual workers, all that should differ is exposure to the "treatment" which is having a technical skill put on a path towards zero demand. It is important to note that the realized treatment—the actual effects on the demand for Flash at the actual pace observed—is inherently a fact particular to our setting. It is easy to imagine other skills being made obsolete more or less quickly, which in turn could prompt quantitatively or even qualitatively different responses. Our view is that elucidating in detail one particular example and focusing on the economic mechanisms is the only viable path towards generalizability.

In studies of displaced workers, the consensus seems to be that the fixed effect approach works best with long observation windows and relatively mature workers whose hourly earnings likely reflects their market earnings potential (von Wachter et al., 2009). We have some insight into a worker's age and off-platform experience (which we will discuss), but in this market, the short duration of projects and the lack of other factors that could affect earnings (amenities, benefits, relational contracts, etc.) implies that worker productivity is frequently "marked to market." We have a relatively long panel, and so we can observe workers for some time prior to TOF-day, though we saw earlier, entry was peaking just prior to TOF-day.

5.2 Construction of a matched sample

We take all workers that were active on the platform at TOF-day, with "active" defined as working at least 40 hours on platform before TOF-day. We then take those workers with some Flash experience (more than 1%) and then compute the fraction of hours-worked that were on Flash projects, FRACFLASH, which is about 25%.[18] This is likely a lower bound, as not all employers label their openings with skills, instead discussing them in the job text.

We match the Flash workers to a control group of non-Flash workers (those with no pre-TOF-day Flash experience), (Sekhon, 2011). For matching, we construct a cross-sectional dataset by computing worker-level summary statistics at TOF-day, including (1) platform tenure, (2) average wage, (3) cumulative hours-worked, (4) cumulative earnings and (5) cumulative number of applications sent. We break up the matching into quartiles by Flash pre-TOF-day Flash focus.

In Table 1 the two samples are compared. Panel A reports details about the full distribution, while Panel B reports t-tests for the means. In Appendix A, we plot the full distributions. As the table makes clear, we obtain excellent balance on the covariates we match on. With these cross-sectional matches, we then restrict our monthly panel to only those treated workers and their matched counterparts.

As our samples are drawn from a marketplace, a natural concern is how these two groups interact, potentially creating a SUTVA violation (Blake and Coey, 2014). As we can observe all applications, we can characterize the degree to which the two groups were competing. In over 70% of openings, there was no overlap in applicants—in the remainder where the was some overlap, there was typically just one applicant from the other group. As workers send many applications and at least some Flash developers have at least some non-Flash skills, this speaks to how "clustered" this market is by skill. Furthermore, this "overlap" fraction falls over time, contrary to the notion that out-of-work Flash developers were crowding out workers in the control. One of the reasons why SUTVA is probably not a first order concern is that the Flash sample was relatively small compared to the market as a whole—at the time of TOF-day, the total number of workers who had applied to at least job was over 250,000.

5.3 Time series for the matched sample

Although only matched on a cross-sectional dataset of what workers "looked like" on TOF-day, we can compare the matched samples over-time, pre-

[18] For FRACFLASH: the min is 0.01 and the max is 1; the 25th, 50th and 75th percentiles are 0.045/0.131/0.34; the mean is 0.25 and standard deviation is 0.28.

Table 1: Comparison of the matched samples

Panel A: Summary statistics

	N	Min	25th	Mean	Median	75th	Max	StDev
Total billed hours-worked								
Non-Flash	1,342	0	144	821	360	431	9,351	1,173
Flash	1,397	0	153	829	397	497	9,280	1,185
Number of organic apps								
Non-Flash	1,342	0	59	387	222	261	2,985	456
Flash	1,397	0	64	406	232	275	5,240	484
Total tracked hours-worked								
Non-Flash	1,342	0	128	738	331	390	9,172	1,045
Flash	1,397	0	138	721	355	434	8,660	991
Hourly earnings								
Non-Flash	1,342	0	1,474	11,148	4,169	5,018	168,534	19,148
Flash	1,397	0	1,534	11,671	4,607	5,533	255,585	20,940

Panel B: Means comparisons

	Means		Difference		
	Flash	Non-Flash	Diff.	SE	t-stat
Total billed hours-worked	829	821	7.9	45.1	0.2
Number of organic apps	406	387	19.4	18.0	1.1
Total tracked hours-worked	721	738	−17.4	39.0	−0.4
Hourly earnings	11,671	11,148	522.9	766.2	0.7

Notes: This table reports group summary statistics for several worker attributes at TOF-day, as well as a t-test comparing those means.

TOF-day. If the two groups moved similarly pre-TOF-day, it is a promising sign that the groups are truly comparable. In Figure 6, we plot the monthly averages for treated and control workers, demeaning each worker by their pre-period mean, and then demeaning each aggregated series with the value at TOF-day (so each series mechanically has a value of 0 for that period). Focusing on the pre-period, we can see that Flash and Non-Flash samples move together for all outcomes.

Post-TOF-day, the groups begin to diverge on some measures, previewing some of our main regression results—this concordance is unsurprising but reassuring, given that our demeaning and averaging essentially re-create the regression mechanics. In the top facet we see there is no evidence of a decline in average wages for Flash workers. It is important to note that this is the average wage for all kinds of work, including both Flash and non-Flash projects. Although there is no evidence of a wage decline, in the facet below, there is some evidence of a decline in hours-worked. This decline comes despite evidence of an increase in applications sent, which we can see in the third facet from the top. In addition to perhaps increasing application intensity, there is also evidence of a move away from skills relied on in the past: in the bottom facet, we see that Flash workers begin applying to jobs that are less like their previous jobs, skill-wise. Interestingly, there is also a decline in similarity for non-Flash workers, consistent with the Deming and Noray (2018) claim of frequent STEM obsolescence.

5.4 Estimates of the effects of TOF-day on Flash workers

To obtain estimates of the effects of TOF-day, using our matched panel, we estimate the model

$$y_{it} = \left(\textsc{AnyFlash}_i \times \textsc{Post}_t\right)\beta + \gamma_i + \delta_t + \epsilon_{it}. \tag{1}$$

where y_{it} is an outcome of interest for worker i at period t, $\textsc{Post}_t$ is an indicator for whether the period was after TOF-day, and $\textsc{AnyFlash}_i$ is an indicator for whether the worker has worked in Flash prior to TOF-day. The γ_i represents a full set of worker fixed effects, which will absorb the impact of time-invariant worker characteristics. The δ_t represents a full set of period fixed effects. The error term, ϵ_{it}, includes all other time-varying unobservable shocks to the worker outcome. The β coefficient is interpretable as the average monthly effect of the TOF-day on Flash over the full post-period.

Figure 6: Matched sample monthly averages for treatment and control groups, demeaned

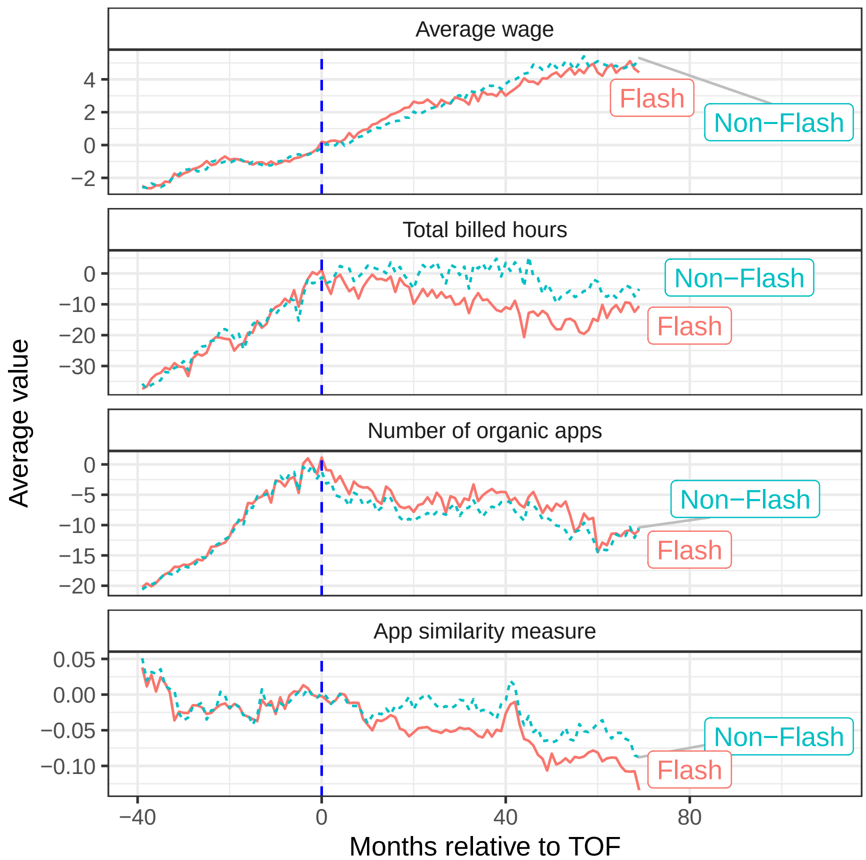

Notes: This figure plots the monthly averages for workers from our matched sample. Before calculating each monthly average, we demean each worker's value by their pre-TOF-day value. From these monthly means, we subtract the value on TOF-day month so that each series has a value of 0 in that period.

We can also allow the effects of TOF-day to depend on how focused the worker was on Flash pre-TOF-day. For each worker, we compute the fraction of hours-worked that were spent on Flash projects, FracFlash. This is zero for the control group. We can interact this measure with the post indicator, creating an independent variable $\text{FracFlash}_i \times \text{Post}_t$. This changes the interpretation of β to a measure of the effect for workers spending 100% of their time on Flash projects, with effects linearly scaled down for workers less-focused on Flash.

Our preferred specification is to restrict the panel to Flash workers with (a) at least 1% of their pre-TOF-day hours in Flash (b) have at least 40 hours of platform work in total and (c) work in at least two distinct months prior to TOF-day. We also truncate the end of the panel by one quarter, as some panel measures are "incomplete" near the end of our data. We also truncate the start of the panel to 40 months before TOF-day, as this corresponds to the true launch of the platform (nearly the entire sample joined the platform after this date). We restrict hourly wages to be positive and less than \$100/hour—this removes a very small number of observations that likely reflect users who were not on *bona fide* hourly contracts but were instead using the time-tracking features of the platform or were using lump sum payments but billed them as hours-worked. We report point estimates using other sample definitions—as well as alternative panel specifications—in Appendix B, which shows our estimates are not sensitive to these choices.

A first question is whether the workers were differentially likely to be active on the platform. We find no evidence that they differed from the control on several definitions of "exit"—this analysis is reported in Appendix B.2.

In Figure 7, we report estimates of β from Equation 1, for both independent variable specifications(AnyFlash and FracFlash). The left facet outcomes are in logs (with 0s removed), in Figure 7a, whereas the right facet outcomes are in levels, in Figure 7b. Mirroring Figure 6, the outcomes from top to bottom are average wages, hours-worked, the number of applications sent, and application similarity. In all regressions, standard errors are clustered at the level of the individual worker. The regressions are run with unbalanced panels when the outcome cannot be computed (such as an average wage or application similarity measure if the worker did not work or send any applications that period, or when the outcome is the log of a count, such as hours-worked).

In the top facet, the outcome is the average wage (in logs in Figure 7a and levels in Figure7b). For average wages in logs, we can see the effect of

Figure 7: Effects of TOF-day on Flash worker outcomes

(a) Log outcomes (b) Level outcomes

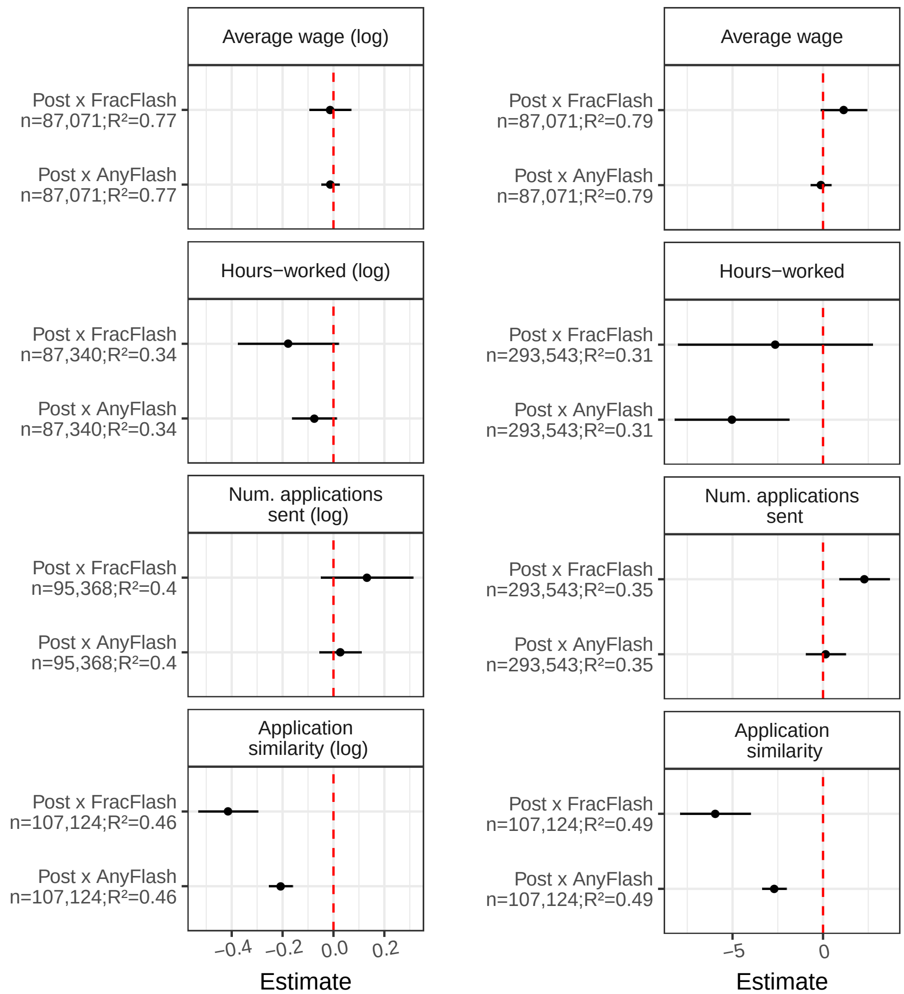

Notes: This figure reports estimates of β from Equation 1. We estimate models in which there is a single treatment indicator, AnyFlash, and models where the independent variable is FracFlash, which is the fraction of pre-TOF-day hours-worked that were on projects requiring Flash. Sample sizes for each regression are reported under the point estimates. Differences in sample sizes reflect the fact that the panel is unbalanced.

TOF-day was a precise zero for both specifications—there is no evidence of a decline in hourly wages for Flash workers. This matches what we observed in Figure 6. In levels, there is some limited evidence that the most experienced Flash developers had an increase in earnings—the FRACFLASH specification is positive and nearly significant.

In the second facet from the top, the outcomes are log hours-worked and hours-worked. Every point estimate is negative. In logs, the ANYFLASH specification point estimate is about 10% fewer hours, whereas the FRACFLASH estimate is 20% fewer hours-worked. Given that more Flash-focused workers had to make a larger transition, skill-wise, this could reflect a greater difficulty in finding work, assuming preferences have not changed and this is not an intensive margin response to changed wages. In levels, we also observe a decline but the magnitudes are "switched" with a larger reduction in the ANYFLASH specification, though both estimates are fairly imprecise.

In the second facet from the bottom, the outcome is the log number of applications sent and the number of applications sent. For the number of applications sent—a measure of job search effort—we can see that point estimates are always positive, though for the ANYFLASH group the estimates are close to zero. In contrast, for the FRACFLASH group log regressions, the estimates imply about a 10% increase, conditional upon sending any. For the ANYFLASH group the point estimate is close to zero.

In the bottom facet, the outcome is application similarity (scaled from 0 to 100). There is a marked decline for Flash workers in all specifications. The effects are stronger in the FRACFLASH group, implying that those workers most specialized in Flash had to make the largest adjustments in application focus.

5.5 Distributed lag model

Rather than assuming a single effect as in Equation 1, we can estimate a distributed lag model. The benefit of this specification is that we can trace out the evolution of adjustments. The downside is that we have more parameters to estimate. We estimate

$$y_{it} = \sum_{k=-a}^{b} \left(\textsc{Post}_{ik} \times \textsc{AnyFlash}_i \right) \beta_k + \gamma_i + \delta_t + \epsilon_i, \quad (2)$$

where $\textsc{Post}_{ik}$ is the value of the post indicator k months away, using $a = 12$ and $b = 24$ to give us a one year pre-period and a 2 year post period.

Cumulative effects for our outcomes (in levels) are plotted in Figure 8, imposing a restriction that cumulative effects on $t = -1$ are zero and using the covariance matrix to construct confidence intervals. We also plot effects using FracFlash instead of AnyFlash.

In Figure 8, we can see that across outcomes, the FracFlash estimates are less precise than the AnyFlash estimates. Reassuringly, there is no graphical evidence of pre-trends for any outcome, for either specification.

In the top facet, the outcome is the average wage. Matching our static estimates, there are no effects on average wages with the AnyFlash specification and perhaps some limited evidence of an *increase* in wages in the FracFlash specification.

In the second facet from the top, the outcome is hours-worked. For both FracFlash and AnyFlash, total hours-worked are lower, matching what we observed in the static estimates. The estimates are imprecise and the confidence intervals frequently include zero. Furthermore, there is no clear temporal pattern—particularly no evidence that the gap shrinks by the end of the period.

For application intensity, matching the static results, we see large increases in application intensity, but only in the FracFlash specification. The increase in intensity occurs shortly after TOF-day and stays consistently higher. In contrast, for the AnyFlash group, the effects appear flat and are quite close to zero.

For application similarity, we see a post-TOF-day decline with both specifications, though effects are larger in the FracFlash specifications. The movement away from Flash is not right after TOF-day, only turning clearly negative about 8 months later, which is right around the "Jellybean" announcement. Note the larger effects for FracFlash. This is not mechanical, as we could have seen these incumbents more focused on Flash being disproportionately likely to "stay" in Flash, giving a high similarity measure.

6 Application-level analysis of bidding

Although our matched sample analysis is done with a panel, we observe the individual applications made by workers. For these applications, we know the wage bid offered by the worker. In Table 4, we explore how TOF-day affected wage bidding. In each regression, the outcome is the log wage bid. We start with only the Flash workers (AnyFlash = 1). Standard errors are clustered

Figure 8: Distributed lag model estimates of TOF-day effects

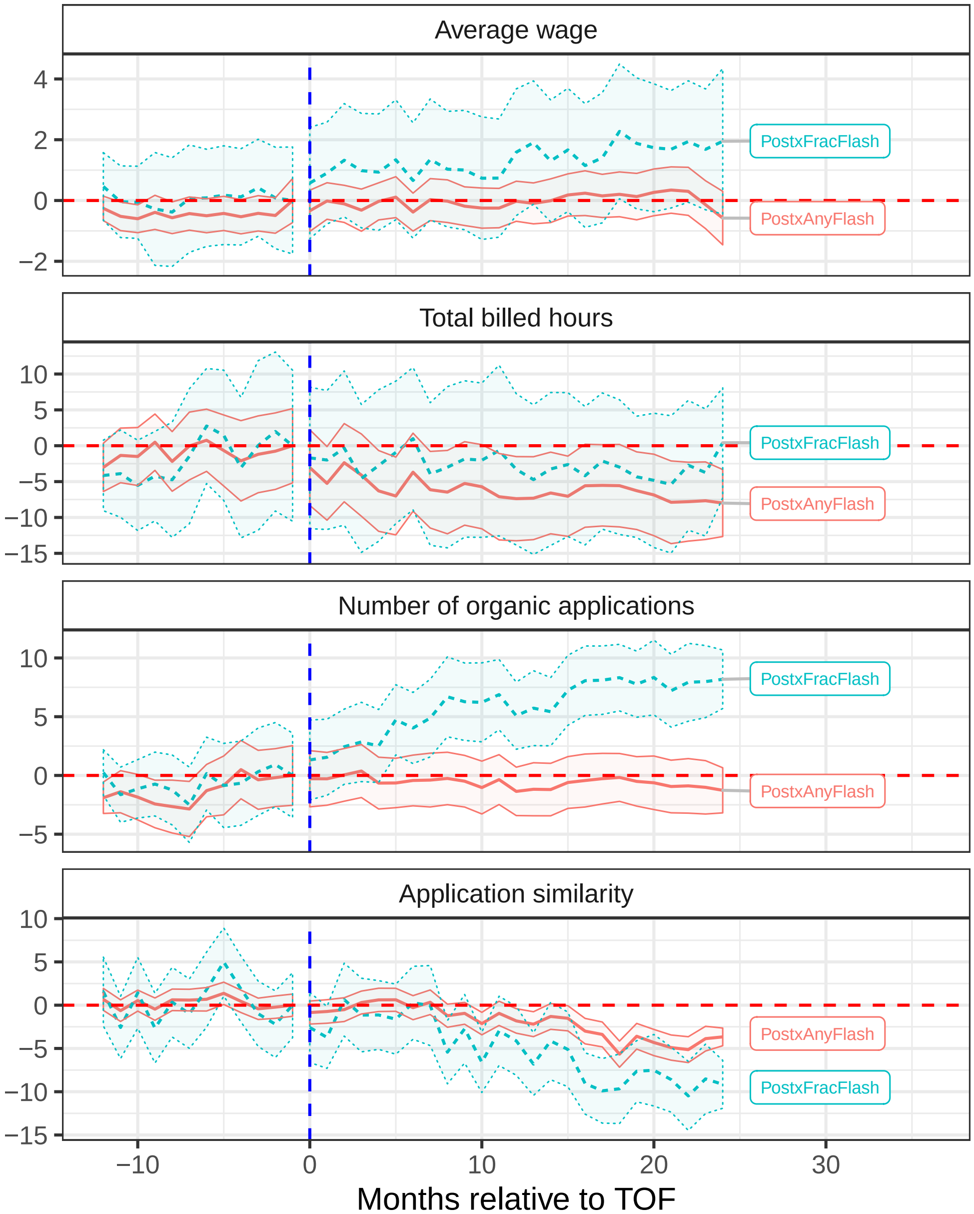

Notes: This figure reports estimates of cumulative effects from Equation 2.

at the level of the individual worker.

In Column (1), we include worker-specific fixed effects and an fixed effect for the period in which the application was made. With this specification, we include an indicator for whether the application was made to a Flash opening. We can see some evidence that among Flash workers, they applied a slight premium to their wage bids (about 2%) when the opening was a Flash opening, pre-TOF-day. This would be consistent with them having a higher productivity for Flash. In the post period, the effect is somewhat reduced (about -1%), but the implied difference is far from significant. In short, as the application level, we have no evidence that Flash specialists substantially reduced their wage bids for Flash openings. We also performed the same analysis, but allowing for effects to differ by FracFlash—we found no evidence that pre-TOF-day concentration mattered. This is consistent with the other panel evidence and the market-level evidence for no decline in wages.

There is a possibility that a change in composition of Flash jobs that might be affecting for the results. For example, if more difficult Flash openings are being posted after TOF-day, perhaps the lack of changes in wage bids could be explained by the inclusion of a compensating differential. However, as we have the entire application graph, we can exploit the matched nature of our data and include both a worker and an opening-specific fixed effect (Abowd et al., 1999). In Column (2), we report this regression, with the period fixed effect dropped. We also drop the Flash app indicator (as it would be absorbed by the opening fixed effect) but we can still include the Flash opening indicator interacted with a Post indicator. The coefficient is negative, consistent with workers bidding less, but the effect is small and imprecise.

By only including Flash workers, we cannot detect a general decline in wage bids, only a difference in wage bids for Flash and non-Flash openings. In Column (3), we expand the sample to applications from all workers, but still include an opening-specific FE and a worker-specific FE. We can see there is no evidence of a decline in wage bids from Flash workers postTOF-day.

6.1 Which Flash workers exited quickly?

Using just the treated Flash workers, we explore when they "exited" Flash. We define exiting as sending their last Flash application. For each worker, we compute this last period and then regress it on select worker attributes

Table 4: Effects of Flash decline on log wage bid to Flash and Non-Flash job openings

	Dependent variable:		
	Log wage bid		
	(1)	(2)	(3)
Flash app	0.020**		
	(0.008)		
Post×Flash app	−0.014	−0.007	
	(0.012)	(0.042)	
Post×AnyFlash			−0.014
			(0.014)
Worker FE	Y	Y	Y
Month FE	Y	N	N
Opening FE	N	Y	Y
Observations	1,239,155	1,239,155	3,325,526
R^2	0.692	0.879	0.881

Notes: This table reports regressions showing how worker bidding changed post-TOF-day Significance indicators: $p \leq 0.10$: †, $p \leq 0.05$: $*$, $p \leq 0.01$: $**$ and $p \leq .001$: $***$.

at TOF-day.

Table 5, the outcome is the last period a Flash application was sent. In Column (1) the main independent variable FracFlash. We can see that workers with higher Flash concentration prior to TOF-day stayed active longer—the implied effects are that a worker focusing 100% on Flash stayed about 5.5 months longer than a worker only doing a small amount of Flash.

Table 5: Period of last Flash application from treated group

	Dependent variable:			
	Months Post-TOF stil active			
	(1)	(2)	(3)	(4)
FracFlash	5.570**		−1.492	
	(1.906)		(2.452)	
Flash bid premium pre-TOF×FracFlash			34.350*	
			(17.376)	
Flash bid premium pre-TOF		−7.505	−17.522*	
		(5.417)	(7.451)	
Age				−0.022
				(0.145)
Constant	33.173***	36.862***	37.196***	35.672***
	(0.566)	(0.659)	(0.887)	(3.521)
Observations	1,842	1,027	1,027	1,385
R^2	0.005	0.002	0.006	0.00002

Notes: This table reports regressions where the outcome is the number of months post-TOF that the Flash worker sent his or her last application to a Flash job opening. Significance indicators: $p \leq 0.10 : \dagger$, $p \leq 0.05 : *$, $p \leq 0.01 : **$ and $p \leq .001 : ***$.

We can also compute the pre-TOF-day difference in wage bids to Flash and non-Flash jobs, and then see if workers who were bidding a larger premium stayed longer. We compute "Flash bid premium pre-TOF" by taking only applications made pre-TOF-day and then averaging wage bids for Flash and non-Flash applications. From Columns (2) and (3), we see some evidence that workers with (a) larger Flash "premium" and (b) high concentration in Flash stayed in the Flash market longer. In Column (4), we use a worker's imputed age at TOF-day (using their self-reported education records, assuming the earliest degree was completed when they were 18). There is no evidence that a worker's age matters in speed of adjustment, though the platform skews very young—the average age is just 24 and the 75th percentile is just 26—that almost no workers are close to likely retirement. In our setting, Flash focus is likely a better proxy for how much Flash-specific human capital a worker has developed rather than biological age.

6.2 Discussion

Our analysis answers one puzzle—the large change in quantities but little change in prices—but it tells us little about how workers made this adjustment. Did they perceive the change in the market contemporaneously? Were they sensitive to the same signal that Jobs' memo sent to the demand side of the market? What prompted them to pursue other skills? How did they decided what skills to pursue? How did they learn those skills? For these questions, we switch methods and report the results of a survey we conducted, with respondents drawn from the AnyFlash = 1 group.

7 Surveys of Flash workers

We surveyed Flash workers that were active on the platform prior to TOF-day. Survey questions were based on responses from two prior rounds of pilot surveys, which elicited free-text responses on the nature of Flash development, technical skill obsolescence, and the impact that changes in the Flash market had on them. Survey respondents were paid $10 to complete the survey. The data were collected from respondents in late June and July of 2017. Of those who were invited to take the survey, 43% participated, giving us a total survey sample size of 186.

7.1 Exiting or staying?

We asked respondents how perceived changes in the market for Flash skills affected their use of Flash. What we wanted to know if the experiences of our surveyed workers matched our empirical results of large declines in active Flash workers. The question we asked was: "If you noticed changes [in the labor market for Flash], how, if at all, did this impact your use of Flash for development?" We restricted possible answers to: 1) continued to use Flash to the same extent, 2) mostly continued to use Flash but switched over to

Figure 9: Responses to the question: "If you noticed changes [in the labor market for Flash], how, if at all, did this impact your use of Flash for development?"

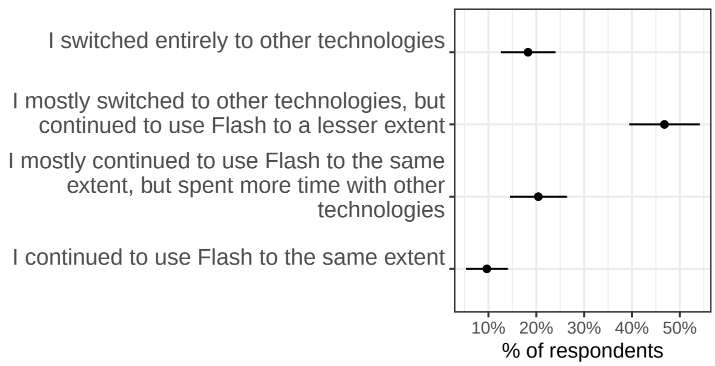

Notes: Responses to question on skills used post-TOF-day.

some other technologies, 3) mostly switched but continued to use some Flash, or 4) switched entirely to other technologies.

The percentage of respondents choosing each possible answer is plotted in Figure 9. Of the respondents, about 65% reported switching to other technologies either completely or to a large extent. 90% reported at least some reduction in the extent to which they used Flash. Survey responses clearly support the notion of a decline in Flash on the supply side of the labor market, but it also is consistent with at least some workers being "stayers."

7.2 Reported effects of the Flash decline on wages and hours

A key finding in our empirical work was large declines in hours-worked in Flash but minimal changes in wages. We asked Flash workers how they perceived changes in the Flash market with respect to wages and project availability. We asked: "How have any changes in the market for Flash projects over the last several years impacted the overall wages you have earned using Flash?" and "How have any changes in the market for Flash projects over the last several years impacted the total number of hours you have spent on projects that use Flash?" Responses were limited to "Decreased significantly," "Decreased somewhat," "No significant change," "Increased

Figure 10: Responses to the questions: "How have any changes in the market for Flash projects over the last several years impacted X" where X = "the overall wages you have earned using Flash?" and X = "the total number of hours you have spent on projects that use Flash?"

(a) Fraction of respondents choosing each characterization, by hours and wages

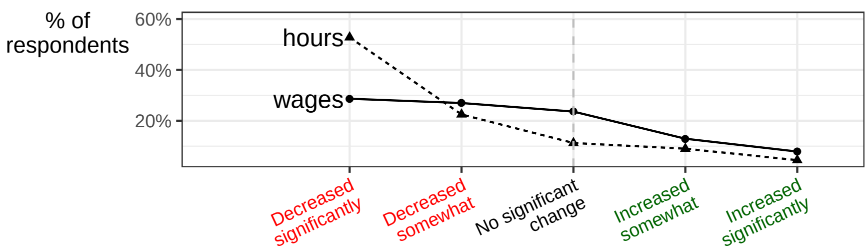

(b) Ordered logit model with cut-points and latent index effects

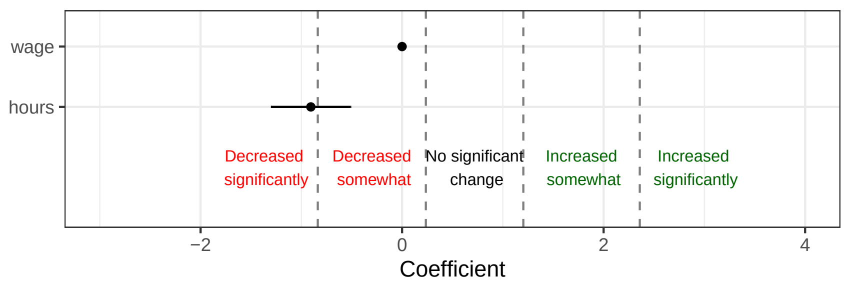

Notes: Top facet reports the percentage of respondents selecting the various possible responses. The bottom facet shows coefficients and cut-points for an ordered logit model.

somewhat," and "Increased significantly." The percentages reporting each response are shown in Figure 10a, for both hours and wages.

Their responses suggest there was a substantial decline in both wages and hours, but more workers report declines in hours than wages. About 50% of respondents experienced a significant decline in hours, but only about 30% of workers experienced a significant decline in wages. A smaller set of workers, about 10%, also reported experiencing a *rise* in hours and wages. That some workers report no decrease in their hours-worked is consistent with some number of incumbents sticking with Flash and our finding that hours-per-worker active in Flash actually rose post-TOF-day.

Given that our responses are ordered, we also estimate an ordered logit. The outcome is the respondent's characterization ("Decreased significantly", "Decreased somewhat" etc.) and the independent variable is whether it was the "hours" question or the "wages" question. An advantage of also estimating an ordered logit is that we can be more confident that the difference in responses we see in the counts is likely not due to sampling variation.

The ordered logit coefficient on question type is shown in Figure 10b, along with the point estimates for the cut points between levels. Mechanically, the confidence interval on the coefficient for the omitted category is 0, as there is also uncertainty in the cut-point estimates (which we do not plot). The ordered logit coefficients give clear evidence that the survey respondents perceived larger declines in hours that wages.

In their free text answers, respondents noted the drop in availability of Flash jobs, but some also pointed out that this was accompanied by a rise in offered wages. When asked to describe changes in the market for Adobe Flash projects between 2010 and 2017,[19] some selected responses were:

> "Definitely a drop in available jobs/job postings. I have found the offer rate has gone up (maybe since the talent pool is smaller now)."

> "The number of available jobs is far fewer in 2017 than in 2010. The types of Flash projects I get are more interesting and difficult than I used to, which may be in part because I have more experience, but may also be because the available jobs in ActionScript are highly specialist, as the more general usage of Flash has shifted to HTML5."

> "Since Steve Jobs' infamous speech about Flash, there has been a decline in the use of Flash across the web. The number of jobs posted has declined, the types of projects were focused mainly on games & video players, interestingly, the budgets for Flash projects have risen due to the decline in the number of available Flash developer."

[19] We asked "Describe any changes you have noticed in the market for Adobe Flash projects between 2010 and 2017—either the numbers of jobs posted, the offer rate, or types of projects. If you began working with Adobe Flash after 2010, describe any changes you have seen in the market since you joined."

Like this last respondent, others also cited the post by Steve Jobs as a triggering event in the decline of demand for Flash projects.[20]

> "Steve Jobs not letting Flash platform on Apple devices is the primary reason. It was not big issue, but the way flash was portait [sic] as evil, people could not understand the whole scenario. Secondary is Android ditching Flash."

> "there was a campaign against flash for a long time and adobe really fumbled the ball, never defending it or improving it appropriately. steve jobs wrote a nasty little hit piece before he died and that really was a turning point. before that me and my partner were working for big clients, making good money, even working with Adobe itself. hahaha he ruined everything. thanks for nothing steve! (but it's really adobes fault for failing to handle the situation at all)."

7.3 Learning new skills or falling back on old skills?

We wanted to understand whether workers were moving towards new skills or relying more on old skills that they were perhaps not using when focused on Flash. We asked: "If you perceived a change in the market for Flash that compelled you to make adjustments, how important was each of the following adjustment strategies on a scale of 1 to 5?" We asked about four specific skill adjustment strategies–switching to existing skills, enhancing existing skills, researching new skills to learn, or learning new skills–on a Likert scale ranging from 1 (*not very important*) to 5 (*extremely important*).

Figure 11 reports the distribution of responses. All four strategies were relatively common in our sample of workers, but researching and switching to new skills was more common than falling back on existing skills. About 35% of respondents reported that they switched to existing skills in response to changes in the Flash market, whereas this number was over 60% for learning new skills.

[20] We asked "If you have noticed changes, what business, economic, or technological factors, do you feel, were the primary CAUSES of these changes in the market for Flash projects?"

We can see the relatively greater importance of new skills in Figure 11b, where the ordered logit coefficient on the "Learn new skills" indicator is significantly higher than the two point estimates for the responses about relying on old skills.

7.4 Deciding which new skills to pursue

Given the large fraction of respondents reporting the need to learn new skills we wanted to explore how they made the choice which specific skills to pursue. We asked the question "On a scale of 1 to 5, how important are each of the following factors when deciding which technical skills to learn?" and then presented 7 possible factors: buzz/word-of-mouth, difficulty to learn, life span, maturity/stability, market wage, future demand and current demand.

Figure 12a plots the percentages choosing each rating for the 7 factors. Current and future demand are both ranked as being important by nearly every respondent, with the majority selecting "Extremely important." Other factors, such as the wage, maturity and lifespan are all ranked as being important, but not as much as the two "demand" factors. Difficulty of learning the skill and buzz were only somewhat important. In short, respondents report being very forward-looking with respect to demand when selecting what skills to learn. This is also apparent in the ordered logit coefficients in Figure 12b, where the two demand-related coefficients—both current and future—stand out as particularly important.

7.5 Methods for learning new skills

After deciding which new skills to learn, workers have to choose how to learn those skills. Using the same 1-5 importance scale, we asked the question "When you switch to a new technology, how important are each of the following when learning the new skill?" about a number of possible methods. The possible sources were "Classroom," "Friends," "Online courses," "Books," "Unpaid projects," "Online forums," and "On-the-job."

The distribution of respondents' answers to this question is shown in Figure 13a. We can see that learning on-the-job and visiting online forums (e.g. websites such as Stack Overflow) were the two most common ways for respondents to learn new technologies. In Figure 13b, we can see that the ordered logit point estimates are both in the "Extremely important" region. The rest of the methods are of middling importance (near level 3) except

Figure 11: Responses to question: "If you perceived a change in the market for Flash that compelled you to make adjustments, how important was each of the following adjustment strategies on a scale of 1 to 5?"

(a) Fraction of respondents choosing each possible response, by strategy

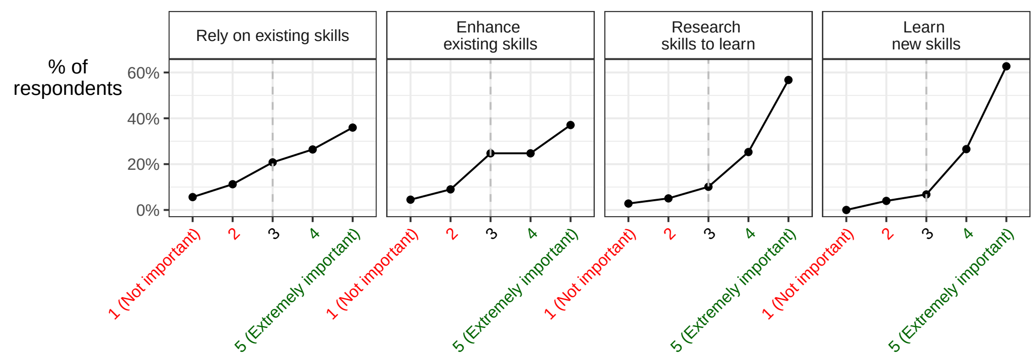

(b) Ordered logit model with cut-points and latent index effects

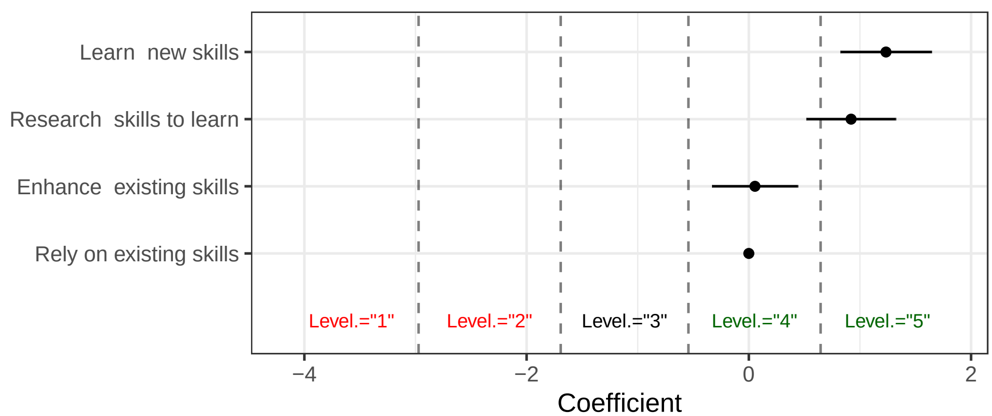

Notes: The top panel are the fractions of respondents choosing each category. The bottom panel are the point estimates and cut points for an ordered logit.

Figure 12: Responses to the question: "On a scale of 1 to 5, how important are each of the following factors when deciding which technical skills to learn?"

(a) Fraction of respondents choosing each possible response, by factor

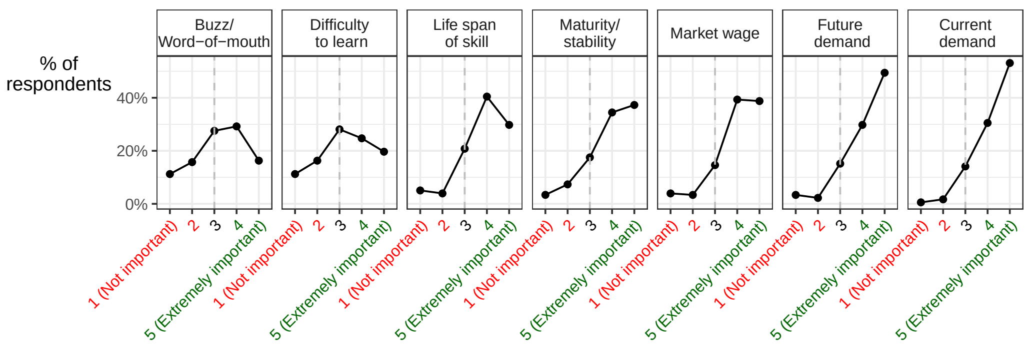

(b) Ordered logit model with cut-points and latent index effects

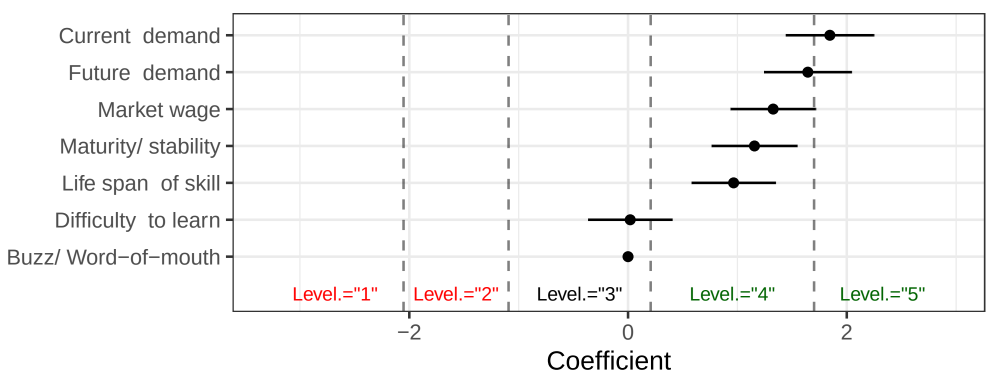

Notes: The top panel shows the fraction of respondents; the bottom panel the result of an ordered logit.

for the classroom, which is ranked as much less important for learning new skills.

7.6 Bidding and learning on-the-job

Given that the most common way workers reporting learning new skills was through on-the-job training, there is a bit of a chicken-and-egg problem—how do workers get work in a skill they do not yet possess?[21] One possibility is that workers take jobs while they are still relatively unskilled and/or offer a wage discount. We asked two questions that try to understand this approach: "When you work with new technical skills, how comfortable do you prefer to be–on a scale from 0 (never used it before) to 100 fully proficient with the technology) – before starting to bid on projects that require that skill?" and "When you are beginning to work with new technical skills, how much of a discount, relative to the market wage, do you apply to your bid? Your answer should be on a scale of 0 (no discount, I bid at the FULL market wage) to 100 (I discount the whole price and do the project for FREE to sharpen my skills)?"

The distribution of the responses to these questions are shown in Figure 14 as histograms. On average, respondents noted that they only need to be about 66% comfortable with a new technology before beginning to work on projects requiring use of that technology. Moreover, the spread around this average is fairly wide, with a fair number of respondents saying that they only needed to be between 20%-60% comfortable with a skill to apply to projects requiring that skills.

With respect to the wage, workers, on average, also noted that they lower their bids when bidding on projects for which they are using new skills. The mean value in the distribution of how much they lower their bid is about 40%. In other words, they bid at about 60% of the level at which they would normally bid when the skill is new. Although we did not "pick up" this lowering wage bidding empirically, it could be because these "learning" jobs are relatively small and as a stepping stone to higher wages which could mask this bidding down dynamic.

These responses are consistent with the notion that an exchange of wages for capital deepening affected workers' choices as they switched to new projects.

[21] The Theodore Roosevelt quote is apt here—"Whenever you are asked if you can do a job, tell 'em, 'Certainly I can!' Then get busy and find out how to do it."

Figure 13: Responses to the question: "When you switch to a new technology, how important are each of the following when learning the new skill?"

(a) Fraction of respondents choosing each possible response, by method

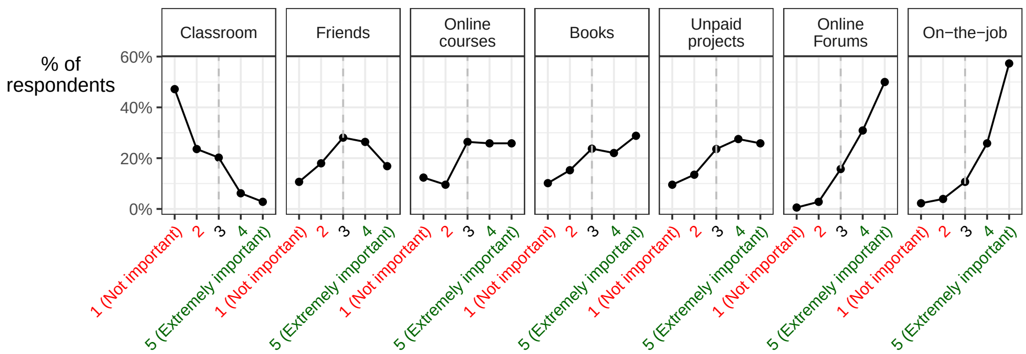

(b) Ordered logit model with cut-points and latent index effects

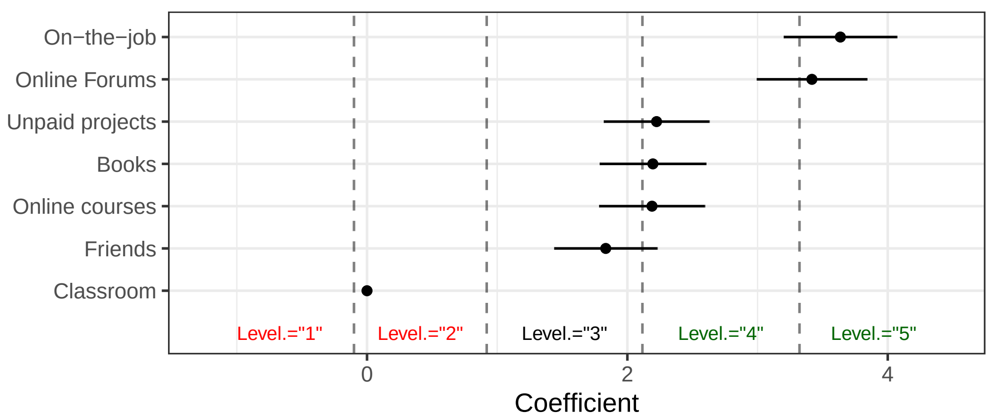

Notes: This question asked what methods workers used to learn new skills. The top panel shows the distribution of responses. The bottom panel plots coefficients and cut-points for an ordered logit model.

Figure 14: Subjective measures of proficiency before bidding on work with a new skill and the amount of discount applied in bids

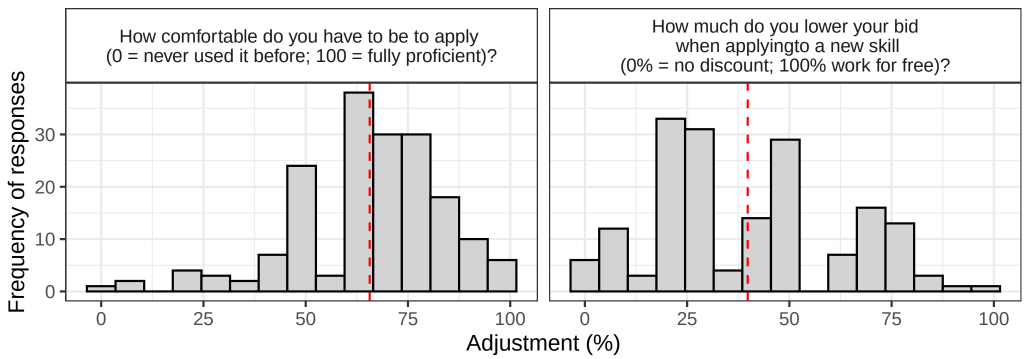

Notes: The left panel shows the distribution of responses to the question: "When you work with new technical skills, how comfortable do you prefer to be—on a scale from 0 (never used it before) to 100 fully proficient with the technology)—before starting to bid on projects that require that skill?" The right panel shows the distribution of responses to the question: "When you are beginning to work with new technical skills, how much of a discount, relative to the market wage, do you apply to your bid. Your answer should be on a scale of 0 (no discount, I bid at the FULL market wage) to 100 (I discount the whole price and do the project for FREE to sharpen my skills)?"

The discounts are large, implying that for new skills, the learning-by-doing component of the work is worth almost as much as the hourly wage itself, at least for these initial projects. If a large fraction of the work being done in any skill is being done by workers with an eye towards the future, the notion of a future negative demand shock propagating back to the "present"—say by causing exit or deterring entry—is understandable.

The free text responses also shed light on this learning-by-doing approach. For those workers who chose to switch to a new skills, we asked about project strategies.[22] A number of workers noted that when switching to a new technology, it was important to start with simpler projects, and build skill by working on projects.

> "Switching to new technology implies that you cannot directly go to expert level projects. It requires starting with easy level and enhance your learning as you work more on live projects."

> "When working with a new technology, I start with simpler projects compared to what I would normally work in using a technology I have already been using for a long time. As I get more projects in with the new tech, I start applying for more complex jobs."

To put this in terms of our model, the respondents are claiming that $y < 1$ but that working in the technology can raise productivity. Finally, we asked workers about their bidding strategies when switching to new projects.[23] The responses to this question were supportive of the notion that workers tend to bid lower for projects where they are using new skills. In fact, a key theme of the responses to this question was the idea of being paid while learning.

> "When switching to a new technology, I start with lower bids first, and then move on to high bids as I gain experience."

> "IF I was switching to say some javascript I would bid low to build up experience and a portfolio in that technology but still be paid something while learning. pretty good deal to get paid anything at all if you're learning a bunch while you go."

[22] We asked "How, if it all, does switching to a new technology affect the types of projects you select?"

[23] We asked "How, if it all, does switching to a new technology affect the amount you bid on these projects?"

8 Conclusion

Our main empirical finding is that despite a large reduction in demand for Flash skills, wages changed very little, due in part to how rapidly workers adapted to this change. The supply of Flash workers proved to be remarkably elastic with no discernible evidence of a decline in wages, in part because of rapid adjustment by the supply side of the market. The adjustment was rapid because workers were forward-looking about human capital choices and their primary strategy of developing human capital being learning-by-doing—making quick adjustment possible.

In terms of generalizability, our results are derived from a particular context where technological shocks are commonplace, the pace of technological change is quick, and there are relatively few frictions associated with moving to new technologies or employers. Furthermore, there are many "nearby" skills that workers who are affected can switch to. Other types of skills might offer fewer adjustment options to affected workers, though our results do suggest what features seem to help adjustment—namely a relatively easy way to gain on-the-job training. The relatively low stakes of hires and the short-duration of the relationships on the platform we study might be important factors for helping workers adjust. That being said, it is disconcerting to note that even years after the decline of Flash, the most Flash-focused workers seem to have lower hours-worked and higher application intensity. The cause is unlikely to be a change in preferences.

The waxing and waning of the demand for specific human capital is commonplace in a dynamic economy, and the waxing an waning of IT or STEM skills is particularly important. Many new technologies have a complementary labor component, and the diffusion of that technology could, in principle, depend on how quickly complementary skills appear. If hiring workers with a waning skill became dramatically cheaper, we might see more use of that technology than we otherwise would. Instead, if firms face a considerable pay premium to use a legacy technology, the pace of diffusion of new technologies will not be constrained by cost disadvantages. Our results are consistent with the labor market not being a significant impediment to technology adoption: the "old" skill did not become relatively cheaper, and the price of "new" (at least to Flash workers) skills would if anything fall, as supply increased and/or workers were willing to work at a discount to increase their human capital. If workers quickly move to where the demand will be, the notion of a general skills mismatch seems unlikely (Marinescu and Rathelot, 2018).

References

Abowd, John M, Francis Kramarz, and David N Margolis, "High wage workers and high wage firms," *Econometrica*, 1999, *67* (2), 251–333.

Acemoglu, Daron, "Why do new technologies complement skills? Directed technical change and wage inequality," *Quarterly Journal of Economics*, 1998, pp. 1055–1089.

_ , "When does labor scarcity encourage innovation?," *Journal of Political Economy*, 2010, *118* (6), 1037–1078.

Anderson, Katharine A, "Skill networks and measures of complex human capital," *Proceedings of the National Academy of Sciences*, 2017, *114* (48), 12720–12724.

Ang, Soon, Sandra Slaughter, and Kok Yee Ng, "Human capital and institutional determinants of information technology compensation: Modeling multilevel and cross-level interactions," *Management Science*, 2002, *48* (11), 1427–1445.

Artuç, Erhan, Shubham Chaudhuri, and John McLaren, "Trade shocks and labor adjustment: A structural empirical approach," *American economic review*, 2010, *100* (3), 1008–45.

Autor, David H, Frank Levy, and Richard J Murnane, "The skill content of recent technological change: An empirical exploration," *Quarterly Journal of Economics*, 2003, *118* (4), 1279–1333.

Baker, Scott R and Andrey Fradkin, "The impact of unemployment insurance on job search: Evidence from Google search data," *Review of Economics and Statistics*, 2017, *99* (5).

Barley, Stephen R and Gideon Kunda, *Gurus, hired guns, and warm bodies: Itinerant experts in a knowledge economy*, Princeton University Press, 2011.

Ben-Porath, Yoram, "The production of human capital and the life cycle of earnings," *Journal of Political Economy*, 1967, *75* (4, Part 1), 352–365.

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan, "How much should we trust differences-in-differences estimates?," *Quarterly Journal of Economics*, 2004, *119* (1), 249–275.

Bessen, James E, "Technology and learning by factory workers: The stretch-out at Lowell, 1842," *Journal of Economic History*, 2003, *63* (1), 33–64.

_ , "Information Technology and Learning On-the-Job," 2016.

Blake, Thomas and Dominic Coey, "Why marketplace experimentation is harder than it seems: The role of test-control interference," in "Proceedings of the fifteenth ACM conference on Economics and computation" ACM 2014, pp. 567–582.

Borusyak, Kirill and Xavier Jaravel, "Revisiting event study designs," *Working Paper*, 2016.

Bound, John, Breno Braga, Joseph M Golden, and Gaurav Khanna, "Recruitment of foreigners in the market for computer scientists in the United States," *Journal of labor economics*, 2015, *33* (S1), S187–S223.

Bresnahan, Timothy F., Erik Brynjolfsson, and Lorin M. Hitt, "Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence," *Quarterly Journal of Economics*, 2002, *117* (1), 339–376.

Card, David, "The impact of the Mariel boatlift on the Miami labor market," *ILR Review*, 1990, *43* (2), 245–257.

Chari, Varadarajan V and Hugo Hopenhayn, "Vintage human capital, growth, and the diffusion of new technology," *Journal of Political Economy*, 1991, pp. 1142–1165.

Chen, M Keith, Judith A Chevalier, Peter E Rossi, and Emily Oehlsen, "The value of flexible work: Evidence from Uber drivers," *Journal of Political Economy*, Forthcoming.

Chin, Aimee, Chinhui Juhn, and Peter Thompson, "Technical change and the demand for skills during the second industrial revolution: Evidence from the merchant marine, 1891-1912," *Review of Economics and Statistics*, 08 2006, *88* (3), 572–578.

Couch, Kenneth A and Dana W Placzek, "Earnings losses of displaced workers revisited," *American Economic Review*, 2010, *100* (1), 572–89.

Davis, Steven J and Till von Wachter, "Recessions and the costs of job loss," *Brookings Papers on Economic Activity*, 2011, pp. 1–73.

Deming, David J and Kadeem L Noray, "STEM careers and technological change," 2018.

Dube, Arindrajit, Jeff Jacobs, Suresh Naidu, and Siddharth Suri, "Monopsony in online labor markets," *American Economic Review: Insights*, 2019.

Edin, Per-Anders, Tiernan Evans, Georg Graetz, Sofia Hernnäs, and Guy Michaels, "Individual Consequences of Occupational Decline," Technical Report, Centre for Economic Performance, LSE 2019.

Filippas, Apostolos, John Horton, and Richard Zeckhauser, "Owning, using and renting: Some simple economics of the sharing economy," *Management Science*, forthcoming.

Goldin, Claudia and Lawrence F Katz, "Human capital and social capital: The rise of secondary schooling in America, 1910–1940," *Journal of Interdisciplinary History*, 1999, *29* (4), 683–723.

_ **and** _ , "The race between education and technology: The evolution of US educational wage differentials, 1890 to 2005," 2007.

Goldschmidt, Deborah and Johannes F Schmieder, "The rise of domestic outsourcing and the evolution of the German wage structure," *Quarterly Journal of Economics*, 2017, *132* (3), 1165–1217.

Guvenen, Fatih, Burhanettin Kuruscu, and Serdar Ozkan, "Taxation of human capital and wage inequality: A cross-country analysis," *Review of Economic Studies*, 04 2014, *81* (2), 818.

Hall, Jonathan V and Alan B Krueger, "An analysis of the labor market for Uber's driver-partners in the United States," *ILR Review*, 2018, *71* (3), 705–732.

Horton, John J, "Online labor markets," in "International workshop on internet and network economics" Springer 2010, pp. 515–522.

_ , "The effects of algorithmic labor market recommendations: Evidence from a field experiment," *Journal of Labor Economics*, 2017, *35* (2), 345–385.

Isaacson, Walter, *Steve Jobs*, Simon and Schuster, 2011.

Jacobson, Louis S, Robert J LaLonde, and Daniel G Sullivan, "Earnings losses of displaced workers," *American Economic Review*, 1993, pp. 685–709.

Jovanovic, Boyan and Yaw Nyarko, "Learning by doing and the choice of technology," *Econometrica*, 1996, *64* (6), 1299.

Kahn, Lisa B, "The long-term labor market consequences of graduating from college in a bad economy," *Labour Economics*, 2010, *17* (2), 303–316.

Kletzer, Lori G, "Job displacement," *Journal of Economic Perspectives*, 1998, *12* (1), 115–136.

Krebs, Tom, "Human capital risk and economic growth," *Quarterly Journal of Economics*, 2003, *118* (2), 709–744.

Marinescu, Ioana and Roland Rathelot, "Mismatch unemployment and the geography of job search," *American Economic Journal: Macroeconomics*, 2018, *10* (3), 42–70.

_ **and Ronald Wolthoff**, "Opening the black box of the matching function: The power of words," Technical Report, National Bureau of Economic Research 2016.

Michaels, Guy, Ashwini Natraj, and John Van Reenen, "Has ICT polarized skill demand? Evidence from eleven countries over twenty-five years," *Review of Economics and Statistics*, 2014, *96* (1), 60–77.

Mithas, Sunil and Mayuram S Krishnan, "Human capital and institutional effects in the compensation of information technology professionals in the United States," *Management Science*, 2008, *54* (3), 415–428.

Murphy, Kevin, "Specialization and Human Capital." PhD dissertation, University of Chicago 1986.

Neal, Derek, "Industry-specific human capital: Evidence from displaced workers," *Journal of Labor Economics*, 1995, pp. 653–677.

Nedelkoska, Ljubica, Frank Neffke, and Simon Wiederhold, "Skill mismatch and the costs of job displacement," *Workin paper*, 2015.

O'Mahony, Siobhan and Beth A Bechky, "Stretchwork: Managing the career progression paradox in external labor markets," *Academy of Management Journal*, 2006, *49* (5), 918–941.

Roth, Alvin E, "Marketplaces, markets, and market design," *American Economic Review*, 2018, *108* (7), 1609–58.

Sekhon, Jasjeet S, "Multivariate and propensity score matching software with automated balance optimization: the matching package for R," *Journal of Statistical Software*, 2011.

Simcoe, Timothy, "Standard setting committees: Consensus governance for shared technology platforms," *American Economic Review*, 2012, *102* (1), 305–36.

Tambe, Prasanna, Xuan Ye, and Peter Cappelli, "Paying to program? Engineering brand and high-tech wages," *Working paper*, 2018.

von Wachter, Till, Jae Song, and Joyce Manchester, "Long-term earnings losses due to mass layoffs during the 1982 recession: An analysis using US administrative data from 1974 to 2004," *Working paper*, 2009.

A Balance of matched sample

Figure 15 shows the distribution of worker attributes at TOF-day for the full sample, the Flash worker sample, and their matches. We can see that the full sample differs dramatically from the Flash sample, but that after matching, the distributions are quite similar.

Figure 15: Comparison of worker attribute distributions at TOF-day for Flash sample, all workers and the matched control group

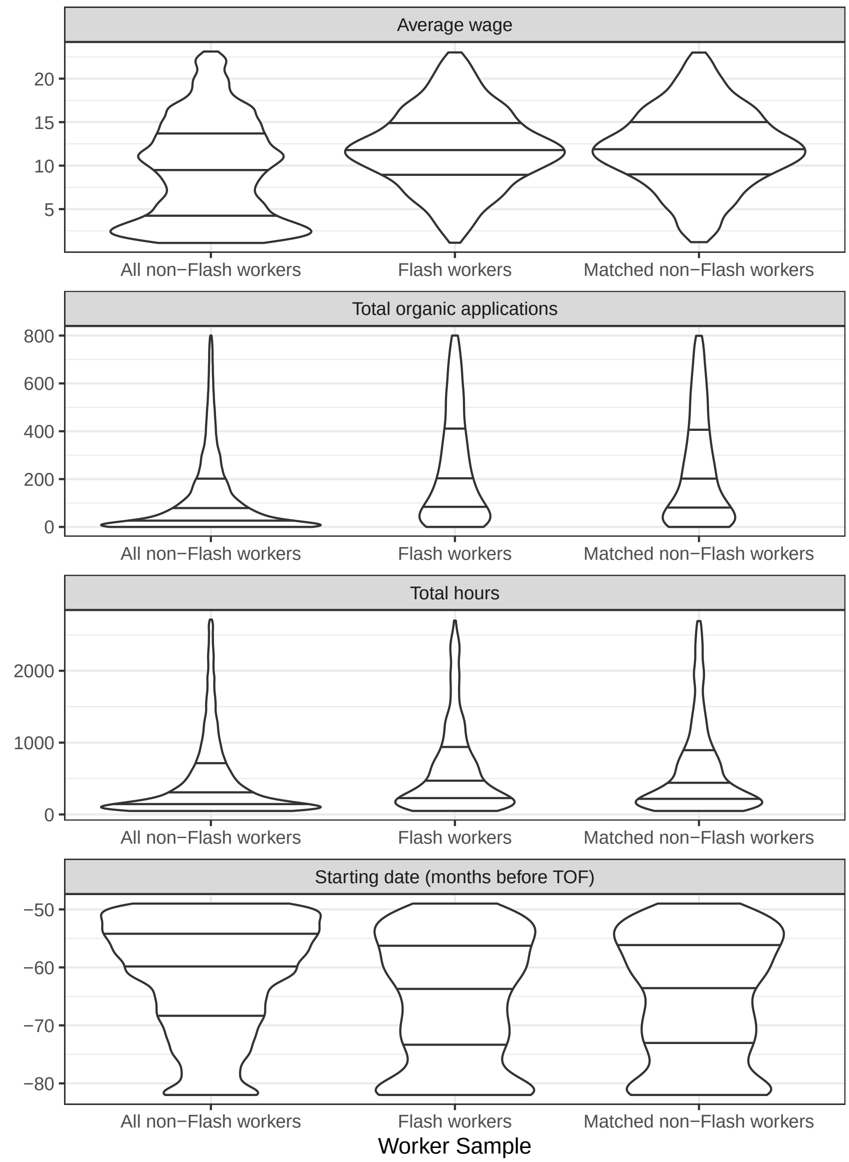

Notes: This figure shows shows the distribution of worker attributes at TOF-day for the full sample, the Flash worker sample, and their matches.

B Robustness and sensitivity

B.1 Panel definition and specification

In addition to choosing the independent variable, another alternative to this "long" panel form is to collapse everything into a pre/post period, creating a short panel. With this structure, rather than have hours-worked per month as a dependent variable, the outcome is the total hours-worked before or after TOF-day. This structure has the nice benefit of making the panel more "balanced" for outcomes that are conditionally defined (e.g., the hourly wage, which we can only calculate for a given month if the worker works some number of hours). It also has the advantage of obviating Bertrand et al. (2004) problems. With this short panel approach, for outcomes that are counts (such as hours-worked), we divide the sum by the number of active periods make effects comparable to the "long" panel estimates. In the short panel specification, we replace the period indicators with a POST indicator but maintaining the worker-specific fixed effects. Figure 16 plots all the point estimates using the long and short panels, the two independent variable specifications and different sample selection criteria. With the exception of the "any" outcomes (which get collapsed to nearly 100% in the short panel), all the point estimates are broadly similar.

B.2 Extensive margin

Figure 18 plots effects of TOF-day on whether a worker was active on the platform. We use two measures of activity (1) did they earn any money and (2) did they send any applications. All the point estimates are close to zero and are not conventionally significant. There is some evidence of a reduced probability of earning some amount of money, but also some evidence of them being more likely to send an application.

B.3 Dynamic panel data model

We estimate the model

$$y_{it} = (\textsc{AnyFlash}_i \times \textsc{Post}_t)\,\beta + \sum_{p=1}^{k} y_{i(t-p)}\lambda_p + \gamma_i + \delta_t + \epsilon_{it}. \tag{3}$$

Figure 16: Point estimates for all models

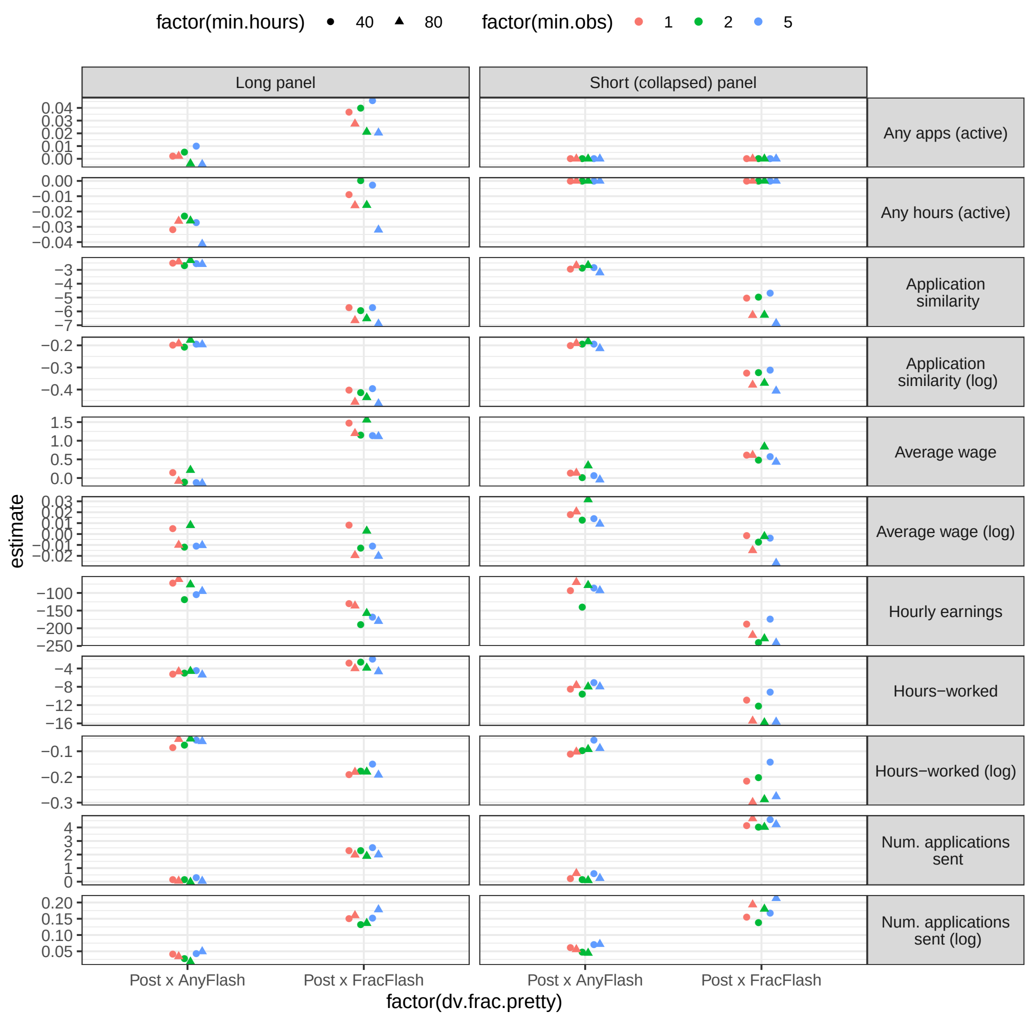

Notes:

Figure 17: Effects of TOF-day letter on Flash worker outcomes

Figure 18: outcomes

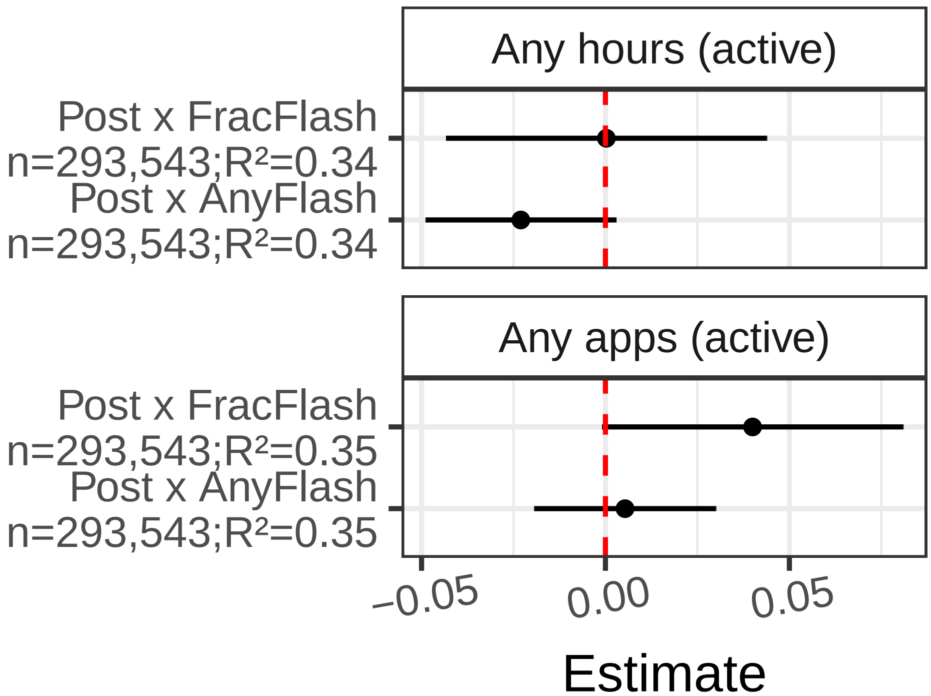

Notes: This figure reports estimates of β from Equation 1. We estimate models in which there is a single treatment indicator, AnyFlash and models where the independent variable is FracFlash, which is the fraction of pre-TOF-day hours-worked that were on projects requiring Flash. Sample sizes for each regression are reported under the point estimates. Differences in sample sizes reflect the fact that the panel is unbalanced.

We assume sequential exogeneity

$$\mathbf{E}[\epsilon_{it}|y_{it-1}, \ldots y_{i0}, \textsc{Post}_t, \ldots \textsc{Post}_0, \delta_t, \gamma_i] = 0. \tag{4}$$

In Figure 19, we report the long-run estimates of the TOF-day, $\beta/(1 - \sum_k \lambda_p)$. We use values of k from 1 to 4. We also report our baseline no lag version for comparison.

Figure 19: Effects of Flash decline using lagged dependent variable models

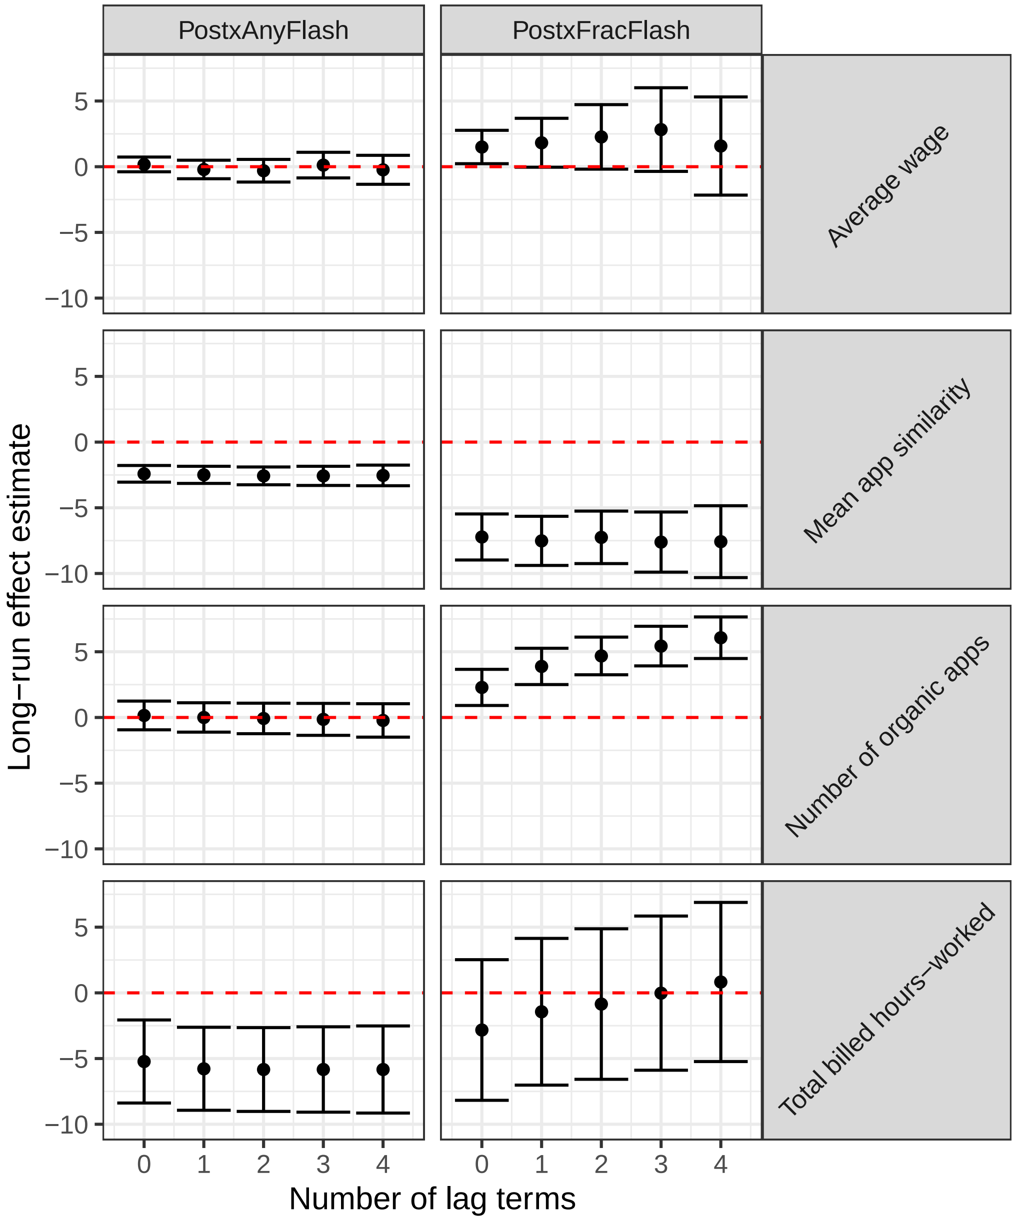

Notes: Dynamic panel data estimates.