The Persistence of Broadband User Behavior: Implications for Universal Service and Competition Policy

Andre Boik, Shane Greenstein, and Jeffrey Prince*

November 2018

Abstract

In several markets, firms compete not for consumer expenditure but consumer attention. We examine user priorities over the allocation of their time, and interpret that behavior in light of salient tensions in policy discussions over universal service, data caps, and related policy topics, such as merger analysis. Specifically, we use extensive microdata on user online choice to characterize the demand for the services offered online, which drives a household's supply of attention. Our data cover a period of time that saw the introduction of many new and notable sites and new devices on which to access them. In our analysis, we assess "how" households supply their attention along various dimensions, such as their concentration of attention across the universe of sites and the amount of attention expenditure per domain visit. Remarkably, we find no change in "how" households allocated their attention despite drastically changing where they allocated it. Moreover, conditional on total attention expenditure, demographics entirely fail to predict our key measures of attention allocation decisions. We highlight several important implications, for policy and beyond, stemming from the persistence and demographic orthogonality of our novel attention measures.

---

* University of California, Davis, Department of Economics (aboik@ucdavis.edu); Harvard Business School, Department of Technology and Operations Management (sgreenstein@hbs.edu); and Indiana University, Department of Business Economics and Public Policy, Kelley School of Business (jeffprin@indiana.edu). We thank the Kelley School of Business and the Harvard Business School for funding. We thank Michael Kummer, Scott Savage and Mo Xiao for excellent suggestions. We thank seminar audiences at Georgetown, Harvard, Northwestern, Oklahoma and the Federal Communications Commission, and conference participants at the American Economic Association Annual Meetings, the International Industrial Organization Conference, Silicon Flatirons, the Research Conference on Communications, Information, and Internet Policy, and the Searle Conference on Internet Commerce and Innovation. Philip Marx provided excellent research assistance, and Kate Adams provided excellent editorial assistance. We are responsible for all errors.

1 Introduction

Usage of services has long played a central role in communications policy in the United States. It remains central to the analysis of policies for universal service and merger review. For example, which users spend the most time online and what content do they consume? Are these choices of users visible in verifiable demographic features, such as low income or advanced age? When mergers among content providers such as Facebook and Instagram are proposed, how does this alter choices of users who previously substituted between content providers? To answer these questions requires detailed information about usage. Despite being a service that goes to more than one hundred million households in the US, and generating tens of billions of dollars in yearly access charges, it is surprising we know so little about online user behavior. There are few studies characterizing online usage and user priorities and the absence of user studies is a large gap.

At first glance, consumer attention allocation decisions across Internet sites seem to have much in common with other standard consumer choice settings: Because Internet user attention is a scarce resource that must be allocated across the Internet's vast supply of Web sites, users must make choices about where to allocate their finite budget of time. Nevertheless, first impressions are misleading. Consumer choice across Web sites lacks one of the standard hallmarks of a market, namely, relative prices reflecting scarcity and directing the allocation of a scarce resource. Most households pay for monthly service and then allocate online time among endless options without further expenditure. Unless a household faces a binding cap on usage, no price shapes any other marginal decision. Instead, choice depends on non-monetary considerations and the gains of the next-best choice. Evidence suggests that, until recently, only a small fraction of users face monetary constraints while deciding whether to use *additional* online resources (Nevo et al. (2016)). Relatedly, subscription services also play little role in user choice. In fact, only one of the top twenty sites (Netflix) is a subscription service, where the price of its Web site plays an explicit role in decision making.

In this paper, we step back from any specific policy debate, and offer a missing piece about users, a comprehensive picture of user priorities for online actions. We examine user priorities over the allocation of their time, and interpret that behavior in light of policy discussions over universal service, data caps, and related policy topics, such as merger analysis. We use extensive microdata on user online choice to help us characterize the demand for the services offered online, which drives a household's supply of attention. We ground the analysis in a specific time period, the allocation of U.S. household attention in the years 2008 and 2013, which was a time of enormous change in the supply of online options for the more than 70% of U.S. households with broadband connections to the Internet. During this five-year period, U.S. households experienced a massive expansion in online video offerings, social media, and points of contact (e.g., tablets, smartphones), among other changes.

Our dataset contains information for more than forty thousand primary home computers, or "home devices", at U.S. households in 2008 and more than thirty thousand in 2013. These data come from ComScore, a firm that tracks households over an entire year, recording all of the Web sites visited, as well as some key demographics. The unit of observation is a week's worth of choices made by households on the "home device"; unfortunately, we do not have data on other potential devices in the household. Using these data, we calculate household total expenditure of attention and attention shares across different site categories (e.g. social media, news, games). We focus on developing insights that will survive short term trends in surfing, and introduce two novel measures of attention, designed to capture "how" households allocate time online in terms of "breadth" and "depth": attention concentration across sites and attention expenditure per site, respectively. We analyze how these novel measures of online attention allocation decisions vary across demographics and across time during the time periods of 2008 and 2013. During that time period, we observe drastic turnover among the top 20 web sites, the introduction of ubiquitous video content, the iPhone and rival smartphones, and tablets. In contrast, from 2013 to 2016, there is no meaningful turnover in the top 10 web sites[1] and no drastic changes in points of contact come to mind.

In our analysis, we first show how total attention varies with standard demographics. Here we show, consistent with prior work (Goldfarb & Prince, 2008), that total attention declines with income, and this result is strikingly consistent over time. Further, we find major shifts in attention across categories, with movement away from chat and news and toward social media and video. Beyond these basic measures, we illustrate graphically the joint density

[1] See Appendix Table 1 for the Top 20 Domains in 2016.

of the breadth and depth of attention in 2008 and 2013, and assess how these measures depend on basic household demographics.[2] In doing so, we uncover a feature of households' attention allocation behavior that we consider to be remarkable and highly relevant for both policymakers and practitioners. From 2008 to 2013, while we document changes in household total online attention and drastic changes in the shares of attention across site categories, we observe absolutely no meaningful changes in either the breadth or depth of attention. Additionally, the breadth and depth of attention are largely independent of demographics. In other words, we find that these key measures of "how" households allocate attention online are roughly constant across time and can be minimally predicted using demographic differences across households. Hence, our breadth and depth measures appear to be not just novel but quite important, as they capture a basic, persistent household characteristic, beyond what demographics can tell us.

Our findings have implications for broadband and competition policy. Regarding broadband policy, our findings are informative for efforts toward universal service through the Connect America Fund. For example, if non-adopters or underserved users are less costly to serve (in terms of any costs related to internet *usage*, as opposed to connection per se), universal service and subsidies could be tailored to these households. Our findings suggest lower total usage with income and that this relationship is highly stable, casting doubt on the idea that households that are often among the underserved would generate lower costs with respect to usage. Further, the very limited connection between age and income (among other demographics) with expenditure per site visit suggests little (usage) cost difference for service to the underserved conditional on total time online, at least to the extent that costs are linked to the data intensity of usage (e.g. video). , .

Regarding competition policy, our results suggest there is competition among websites to fill users' exogenous slots of time, in addition to competition via website characteristics/offerings. In other words, the competitive environment may be shaped as

[2] Our measure of attention breadth or concentration is a measure frequently used in the economics literature called the Herfindahl-Hirschman Index (HHI). If a household spends most of its time online at few websites, the household's usage is said to be concentration whereas if the household spreads its time substantially across many websites it is said to be unconcentrated. Depth or attention expenditure per site simply measures how long households typically stay at websites that they visit. The details of how we calculate these measures appear in section 3.

much by the relative time demands of websites as the types of content they provide. For policy, this adds a valuable perspective when considering the possible impact on new competitors from non-neutral ISP pricing and zero rating.

Our findings also have implications for advertising, and point towards an additional new approach to privacy policy. For advertising, our results indicate that demographics are weak determinants of duplications, whereas our breadth and depth measures (which are largely orthogonal to demographics) may be much more informative. As privacy laws threaten to limit the ability to track online behavior directly, general information about breadth and depth for a website's visitors may be particularly valuable.

2 Literature review

2.1 Contribution to prior literature

The commercial Internet supports enormous amounts of economic activity, and it has experienced increases in online offerings throughout its short existence. Starting from modest beginnings in the mid-1990s, today this sector of the U.S. economy supports tens of billions of dollars of advertising revenue and trillions in revenue from online sales. Not surprisingly, that phenomenon has spawned an extensive literature, and it has grown so much that it merits handbooks to cover the research (Peitz and Waldfogel 2012). These handbooks organize the literature around many sub-topics, such as the supply and demand for infrastructure, online and offline competition (Lieber and Syverson 2012), and the supply and demand for online advertising (Anderson 2012).

One theme cuts across many of these topics: All households redirect time from other non-Internet leisure activities and allocate it toward online usage, and different online activities compete with each other in the household's budget for time. While researchers recognize that users pay an opportunity cost during online time by withdrawing from other leisure activities or household production activities (Wallsten 2013; Webster 2014), the household's time for, and attention to, its online activities remains incompletely characterized. Hence, there is no widely accepted baseline model of aggregate demand for online activity (and supply of attention) built from a common understanding of online behavior. We contribute to that characterization by discovering a remarkable stability in two measures of households' attention allocation decisions:

attention concentration and attention expenditure per site visit.

Our findings have implications for prior work in this area that are based upon what we might call a "standard" economic model of time allocation in a frictionless labor/leisure framework and which, when estimated, delivers parameters characterizing demand for time across households (Brynjolfsson and Oh 2012; Goolsbee and Klenow 2006). We uncover behavior strongly inconsistent with the frictionless framework, suggesting such parameter estimates may not be capturing what is intended, but instead some inherent friction in the household's decision making process. Examples of these frictions may include: exogenous time constraints necessary to properly consume the offerings of a website, or switching costs across domains. Since our findings speak to the nature of a fundamental feature of any model of the Internet, (i.e. demand), we believe our findings to be important more broadly. For example, search engine competition has motivated some studies on competition for attention (Athey, Calvano, and Gans 2013; Gabaix 2014). There has been some formal statistical work on the competition for attention in the context of conflicts for very specific applications, such as, news aggregators and news sites (Athey and Mobius 2012; Chiou and Tucker 2015), and between different search instruments (Baye et al. 2016). As we describe later, our results suggest a model that stresses several frictions we consider to be important for shaping attention allocation decisions online.

2.2 Dynamics of the Internet Ecosystem: 2008-2013

The era we examine is one characterized by rapid technical advance and widespread adoption of new devices. Continuing the patterns seen since the commercialization of the Internet in the 1990s (Greenstein 2015), new technical invention enabled the opportunity for new types of online activity and new devices. By the beginning of our sample, many online suppliers and startups had begun experimenting with applications that made extensive use of data-intensive video.

The start of our time period, 2007, is near the end of the first diffusion of broadband networks. By 2007, close to sixty-two million U.S. households had adopted broadband access for their household Internet needs, while by 2013 the numbers were seventy-three million. The earlier year also marked a very early point in the deployment of smart phones, streaming services, and social media. The first generation of the iPhone was released in June 2007, and it is widely credited with catalyzing the entry of Android-based phones the following year. By 2013, more than half of U.S. households had a smartphone. Tablets and related devices did not begin to diffuse until 2010, catalyzed, once again, by the release of an Apple product—in this case, the iPad in April, 2010.

The enormous changes in online software were also relevant to our setting and often complemented improved broadband quality: streaming services began to grow. YouTube began to permit video uploads exceeding 10 minutes, and Netflix and Hulu both began offering streaming services at the beginning of our time sample, 2008. Social media was quite young at the beginning of our sample; by 2013, it had become mainstream and widely used. In summary, the supply of sites for users changed dramatically between the two years we examine (2008 and 2013).

3 Data, Attention Measures, and Summary Statistics

3.1 Data

We obtain household machine-level browsing data from ComScore for the years 2008 and 2013. We observe one machine for each household for the entire year, either all of 2008 or all of 2013. Here, the machine should be interpreted as the household's primary home computer. The information collected includes the sites visited on the machine, and how much time was spent at each site. For simplicity, we consider only the first four weeks of a month and do not consider partial fifth weeks, so the maximum number of weeks for a household cannot exceed forty-eight. Importantly, we delete households that have fewer than six months of at least five hours of monthly browsing. We also delete the very few households with more than 10,080 minutes online per week, which was the maximum amount of time allowed and thus the data from these households are the result of a defective tracking device. For 2008, we are left with 40,590 out of 57,708 households, and for 2013 we are left with 32,750 out of 46,926 households. In both years, this amounts to over one million machine-week observations. We observe an average of 42.1 and 41.5 (medians 45 and 44) machine weeks per household (s.d. = 6.9 both years) for 2008 and 2013.

Comscore attempts to obtain a balanced sample of households across years. The demographics we observe include (1) household income categories, (2) educational attainment of the head of the household, (3) household size, (4) age of the head of the household, and (5) an indicator for the presence of children. For income, ComScore's sampling of households is known to target higher-income households, and we observe that those income levels are comparable across the 2008 and 2013 data. Unfortunately for education attainment, the education identifiers are mostly missing in 2008, and only available for roughly half of all households in 2013. Meanwhile for age, there do not appear to be any major differences in the sample composition across years (the 2013 heads of households are mildly younger). In addition, ComScore provides no information on the speed of the broadband connection except to indicate that virtually none of them connect through dial-up.

One concern with the data is our measurement of total attention expenditure. If a Comscore household leaves a browser open, we do not know if the user is calmly consuming its content or whether the user has left the room. Comscore ends such sessions after a period of inactivity, but this is certainly a limitation of the data that biases total attention expenditure and average expenditure per site visit upwards. However, our main finding is that our two key measures of attention allocation decisions (attention concentration and attention expenditure per site visit) are constant across time and not explained by demographics. We do not expect time or specific household demographics to be meaningfully correlated with a tendency to consume calmly or a tendency to leave the room. And, with regard to our calculations of attention expenditure, these calculations are conditional on total attention expenditure and should be unaffected.

3.2 Attention Measures

We measure four aspects of household attention allocation decisions at the weekly level. The first two are standard for the literature and for commercial purposes: the total amount of attention allocated online, and attention market shares by category of site (which essentially generate website rankings). Our additional two measures are novel, and are designed to capture the breadth and depth of a household's online attention. These two new measures are: the concentration of attention across sites and the average expenditure per site visit. The total amount of attention allocated online for a household is simply the sum of all minutes the household allocated across all sites in the given week. We classified the top 20 domains in 2008 and 2013 by categories defined by Webby and easily computed category-level market shares for 2008 and 2013. Regarding our new measures, the household's concentration of attention across sites is measured using a Herfindahl-Hirschman Index where the market share of each site visited is equal to the total number of minutes spent at that site divided by total minutes spent at all sites. Then the Herfindahl-Hirschman Index measure for that household is obtained by summing the squares of those market shares.[3] We measure the average expenditure per site visit by calculating the fraction of site visits that exceed ten minutes. We refer to these last two measures as summaries of "how" households behave online since we consider them to be behavioral measures, which abstract from total time online and content preferences.

3.3 Summary Statistics

In Tables 1 and 2, we present summary statistics corresponding to our household demographics and attention measures, respectively. If a household is online in a given week, on average it allocates roughly fifteen hours of attention online per week in 2008 and fourteen hours online in 2013. Household attention concentration measured through a Herfindahl-Hirschman Index is roughly constant across years at 2,900 and the fraction of site visits where attention expenditure exceeds 10 minutes is also constant at 75%. As we will discuss in greater depth, the similarities in "how" households allocate attention across years extends beyond the simple means. We discuss category market shares in greater detail later, but the interested reader can skip ahead and view Figure 2; category market shares change drastically among the top domains between 2008 and 2013.

[Tables 1 & 2 about here]

4 Results

4.1 Main Findings

[3] The Herfindahl-Hirschman Index is a standard tool used in economics for measuring market concentration. Here, we adapt the concept to measure how concentrated a user's time is across different websites. For example, if a user spends 20 minutes per week on Facebook and 5 minutes per week on Twitter, then the websites' "market shares" are 80% (20/25) and 20% (5/25), respectively. The Herfindahl-Hirschman Index is then obtained by squaring those shares and adding them together: HHI = 80^2 + 20^2 = 6800.

Before analyzing our attention metrics, we first measure how total attention varies with standard demographics and shares of content categories. Table 3 illustrates the relationship between total online attention per week and standard demographics. Here we find what has become a well-established result in the literature (Goldfarb & Prince, 2008): total attention declines with income. We further find that this result is strikingly consistent over time, even when considering a quite volatile time period. To illustrate, Figure 1 shows total attention allocation by income using our data, for both 2008 and 2013, alongside the same calculation by Goldfarb and Prince (2008). In both of the years in our sample, we observe that total attention allocation is decreasing in income (conditional on broadband adoption), the same as found by Goldfarb and Prince (2008) using survey data.

[Table 3 about here]

[Figure 1 about here]

Turning to shares of content categories, Figure 2 shows the proportion of time allocated to different major content categories. Here we see significant shifts in attention across categories, with movement away from chat and news and toward social media and video.

[Figure 2 about here]

A useful merit of our analysis for our attention measures is that most of our results can be illustrated graphically. While a comparison of the summary statistics of "how" households allocate attention online found that the means are economically identical across years, here we illustrate that the similarity extends beyond means to the full joint density[4]. Through visual inspection, it is not only clear that the entire joint density of our measures of how households allocate attention online is constant across years, we also note from Table 4 that it is constant across different types of households within each year. The most that any of the demographics can explain in terms of the difference in means across demographic groups within a year is that the most educated households have a Herfindahl-Hirschman Index which is lower by roughly 3% compared to the least educated households.

[4] Kolgomorov-Smirnov tests for equality of the distribution function reject our hypothesis that the distributions are statistically identical, however, there is no meaningful difference (i.e. in percentage terms the differences are negligible): visual inspection of the joint densities confirms that. That there is a statistically significant difference between the distributions is not surprising given the number of observations in each year is roughly 1.5 million.

[Figure 3 about here]

Figure 3 illustrates the unconditional joint densities of attention concentration (on the horizontal axis) and the fraction of site visits where attention expenditure exceeded 10 minutes (on the vertical axis) for 2008 and 2013. The manner in which the joint density is illustrated, and subsequently in Figures 4 and 5, is via a "heat map": regions that are particularly red are regions where many households' attention measures are located. In contrast, regions that are blue have many fewer households whose attention measures are located there.

It is quite apparent that the joint densities across the two years are virtually identical. The joint densities are unconditional on total time online or any demographic. As a robustness exercise, we also break down these joint densities into different quartiles based on total attention spent online and by the extent of multitasking (simultaneous browsing) in Figures 4 and 5, respectively. While the amount of total attention spent online and the extent of multitasking certainly affect the shape of the joint density in any given year, that shape remains constant across years for all quartiles. We consider it quite remarkable then that, conditional on a level of total attention spent online or a level of the extent to which a household multitasks, our measures of how a household behaves online in terms of its attention concentration and attention expenditure per site are virtually identical and independent of household demographics. Demographics such as income level are already known in the existing literature to affect total attention spent online (discussed above), but conditional on total attention spent online there is virtually no additional explanatory power provided by demographics. To be precise, Table 4 shows that the largest effect of any demographic on mean attention concentration is roughly 3% and at most 2% on the fraction of attention expenditure per site exceeding 10 minutes.

[Figures 4 and 5 about here]

[Table 4 here]

4.2 Robustness

One limitation of our data is that we only observe browsing behavior on the household primary home computer and not on alternative devices. In 2008, the browsing behavior is likely reflective of most of the browsing occurring in the home, but is surely less reflective of total household browsing in 2013 because of the introduction of new devices during that five-year period. Indeed, we observe less total time allocated to the primary home machine in 2013, which is surely more than made up for through increased device usage.[5] However, rather than limiting the strength of our conclusion that how households allocate attention online has remained unchanged from 2008 to 2013, we believe it reinforces it: *despite* households adopting new devices from 2008 to 2013, household concentration of attention across sites and average expenditure per site visit have remained unchanged. Ordinarily one would have expected there to be some sort of selection effect drawing certain types of site visits, such as those requiring low average expenditure per visit, away from the home computer towards handheld devices. Selecting short visits away from the home computer towards handheld devices would increase the average expenditure for site visits remaining on the home computer.

A possible concern with our finding is that the Comscore dataset is not reflective of the general population, posing a problem for external validity. In response, we note two things from our main findings. First, as we show above in Table 3 and Figure 1, our data replicate a well-known result from the existing literature: that conditional on broadband adoption, total attention allocation to the internet is decreasing in income. Second, one may worry that our result is simply caused by households in the Comscore dataset going to the same sites in 2013 as they did in 2008, and therefore stability in our breadth and depth measures would be driven by the fact that households are just consuming the same sites (and in the same way), repeatedly. However, as shown in Figure 2, there were very significant changes in *where* attention was allocated between those two years, with households substituting away from Chat and News towards Social Media and Video,

[5] Several third party data sources confirm this trend. For example, eMarketer data (https://www.emarketer.com/Chart/Monthly-Time-Spent-Online-Among-US-Internet-Users-by-Device-Dec-2013-Dec-2016-billions-of-minutes/205244), and a different set of data from Comscore (https://www.marketingcharts.com/industries/retail-and-e-commerce-27327).

despite no change in "how" attention was allocated in terms of attention concentration and attention expenditure per visit. In Table 5 we show the Top 20 domains in 2008 and 2013. Together, Figure 2 and Table 5 indicate that Comscore households have changed where they are allocating their attention, and that Comscore households appear to visit the same well known sites that the general population visits and which appear in other calculations of the top sites visited in the US using different data[6].

[Table 5 about here]

5. Implications

Our findings have important implications for several different areas within the economics of digitization, including those pertaining to broadband and competition policy. We will generally stress cautionary implications, because that is what our findings tell us at a broad level. Online website choice does not resemble competition between product rivals with priced products. When users allocate their time, they behave differently.

Regarding broadband policy, our findings are informative for efforts toward universal service through the Connect America Fund. It has been a longstanding principal of universal service policy to reduce prices for the underserved. For example, if non-adopters or underserved users are less costly to serve (in terms of any costs related to internet *usage*, as opposed to connection per se), universal service and subsidies could be tailored to these households. Our findings suggest total usage rises with lower income and that this relationship is highly stable, casting doubt on the idea that households that are often among the underserved would generate lower costs with respect to usage. Further, the very limited connection between age and income (among other demographics) with expenditure per site visit suggests little (usage) cost difference for service to the elderly, assuming presuming total time online is linked to the data intensity of usage (e.g. video). That undermines a common approach, which varies prices for the underserved with age,.

[6] For example, see Nielsens Top 10 list on the most visited sites (https://marketingland.com/google-is-most-visited-site-of-2013-despite-big-drops-in-desktop-traffic-nielsen-68235).

Turning to implications for competition policy, competition is usually modeled and measured through cross-price elasticities. In the world of online attention, the price of access to content (i.e., websites) is not relevant for most users. Hence, prices cannot be used to construct cross-price elasticities among online competitors. Instead, the vast majority of sites charge consumers the same price: a unit of time. Since certain sites inherently require more or less time to consume, in a world with exogenous slots of time, competition may be driven as much by substitution across sites that fit a slot adequately as they are by product characteristics. In other words, if a user finds himself with 5 minutes to spend online, that user may substitute across vast categories of sites which are different in the service they provide but which are similar in that they all consume approximately 5 minutes per visit. Competition for online attention may therefore operate as much along the lines of substitutability in terms of minimum attention expenditure required to consume the site.

Relatedly, our persistence finding suggests the existence of heretofore unaccounted-for frictions in online time allocation. The existence of such frictions may cause calculations based on frictionless models to overestimate the consequences of changes in merging firms. Examples of such frictions may be: households having exogenous slots of time to fill (and that the length of those slots may determine which website are consumable in that time frame), or switching costs across websites (which might include searching for additional websites).

That also has direct implications for advertising. If advertisers seek to maximize the number of unique exposures per dollar spent on advertising, rather than the absolute number of exposures, our results suggest that knowing the demographics of the visitors to a site will not be useful in determining the extent of duplicate exposures to advertisements. In contrast, knowing the visitors to a site have low breadth of attention across all sites and have high depth can be particularly useful to an advertiser interested in minimizing the number of duplicate exposures. Sites that attract low breadth and high depth visitors may add value to advertisers as much as sites that offer advertisers access to specialized, targeted audiences. Consequently, information about the breadth and depth of a site's visitors can be quite useful to advertisers, even independent of known demographics.

Lastly, our findings point towards an additional new approach to privacy policy.

As privacy laws threaten to limit the ability to track online behavior directly, general information about breadth and depth for a website's visitors may be particularly valuable.

6. Conclusion

This paper uses extensive microdata on user online activity to characterize household allocation of attention in the absence of prices. We characterize household heterogeneity in allocation of attention and discover a remarkable stability in "how" households allocate their scarce attention online: not only is the mean of attention concentration and expenditure per site visit constant across 2008 and 2013, but so are the entire joint densities of those two measures. This stability persists despite large changes in where households are going online in addition to a number of major changes to the Internet over that time period: the introduction of handheld devices to access the Internet (such as the iPhone), faster broadband, and vastly greater video offerings. We emphasize that this stability persists on the primary home computer *despite* new devices being introduced during the time period we examine, and which were certainly adopted by many of the households in our sample. Perhaps even more surprisingly, demographics almost entirely fail to predict the attention concentration or expenditure per site of the household. We believe it is difficult to reconcile these patterns with frictionless models of time allocation. To explain these patterns seems to require appealing to a set of frictions that is constant across time and demographics. We believe a primary exogenous friction is unchanging household habits, such as family and work habits that are largely uncorrelated with household demographics, which shape the availability of time remaining for online consumption.

These findings can serve as an important guide for future modeling of demand for online attention as competition for that attention. We observe substitution across categories of sites and to new devices to access the Internet, yet no change in the findings or second moments of attention concentration or expenditure per site visit. What model of demand is consistent with these properties? While we do not propose a theoretical framework generating these attention allocation patterns, our empirics are the first to identify these fundamental patterns of online attention and which future models should be capable of producing.

References

Allen, R. 2015. Mobile Internet Trends Mary Meeker 2015 1. Smart Insights. http://www.smartinsights.com/internet-marketing-statistics/insights-from-kpcb-us-and-global-internet-trends-2015-report/attachment/mobile-internet-trends-mary-meeker-2015-1/

Anderson, Simon, 2012. Advertising on the Internet. In *the Oxford Handbook of the Digital Economy*, eds. Martin Peitz and Joel Waldfogel, Oxford University Press.

Athey, S., Calvano, E. and Gans, J. 2013. The Impact of the Internet on Advertising Markets for News Media, Stanford working paper.

Athey, S. and Mobius, M. 2012. The Impact of News Aggregators on Internet News Consumption, Stanford working paper.

Baye, M., De los Santos, B., and Wildenbeest, M. 2016. What's in a Name? Measuring Prominence and Its Impact on Organic Traffic from Search Engines. *Information Economics and Policy* 34: 44-57.

Brynjolfsson, Erik, and JooHee Oh, 2012, The Attention Economy: Measuring the Value of Free Digital Services on the Internet. In *33rd International Conference on Information Systems (ICIS 2012)*. Orlando, Florida: December.

Chiou, Lesley, and Catherine Tucker, 2015, Digital Content Aggregation Platforms: The Case of the News Media. Working paper.

Gabaix, Xavier. 2014. A Sparsity-Based Model of Bounded Rationality. *Quarterly Journal of Economics* 129: 1661-1710.

Goldfarb, A. and Prince J. 2008. Internet Adoption and Usage Patterns are Different: Implications for the Digital Divide. *Information Economics and Policy* 20: 2-15.

Goolsbee, Austan, and Peter Klenow, 2006, Valuing Consumer Products by the Time Spent Using Them: An Application to the Internet. American Economic Review, Papers and Proceeedings. May. pp 108-113.

Greenstein, Shane. 2015. How the Internet Became Commercial. Privatization, Innovation, and the Birth of a New Network. Princeton Press; Princeton, NJ

Lieber, E., and C. Syverson, 2012, "Online versus Offline Competition," In *the Oxford Handbook of the Digital Economy*, eds. Martin Peitz and Joel Waldfogel, Oxford University Press.

Nevo, Aviv, John L, Turner, and Jonathan W. Williams, 2015, Usage Based Pricing and Demand for Residential Broadband, Econometrica, Forthcoming.

Peitz, M., and J. Waldfogel, 2012, *The Oxford Handbook on the Digital Economy*, Oxford University Press.

Wallsten, S. 2013. What are We Not Doing When We're Online? NBER Working paper.

Webster, James, 2014. The Marketplace of attention: How Audiences Take Shape in a Digital Age, MIT Press.

Figures

Figure 1

Total Time Online by Income (2008, 2013)

Comscore & Goldfarb and Prince (2008)

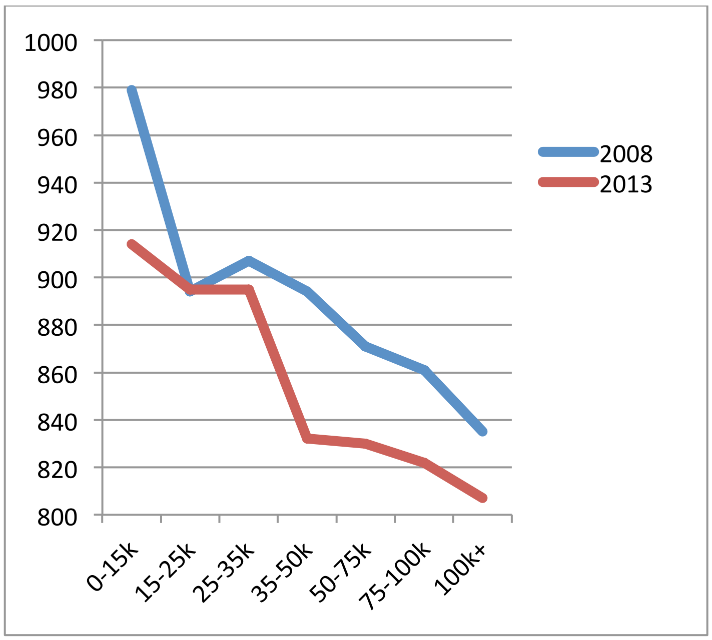

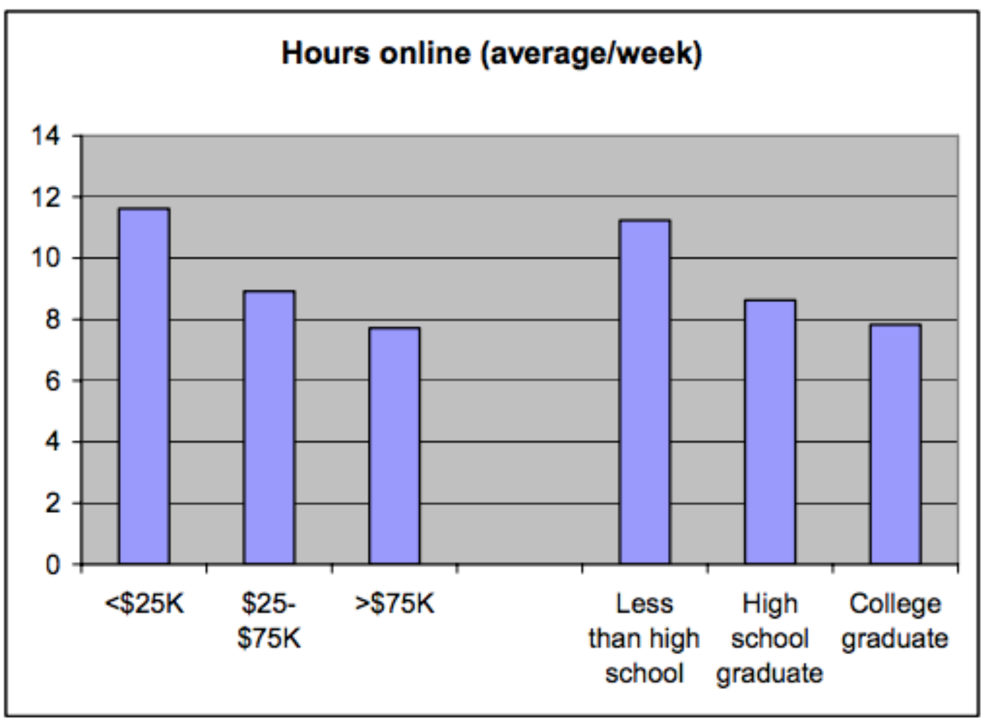

Figure 2

Changes in Share of Attention across the Top 1000 Sites by Category (2008, 2013)

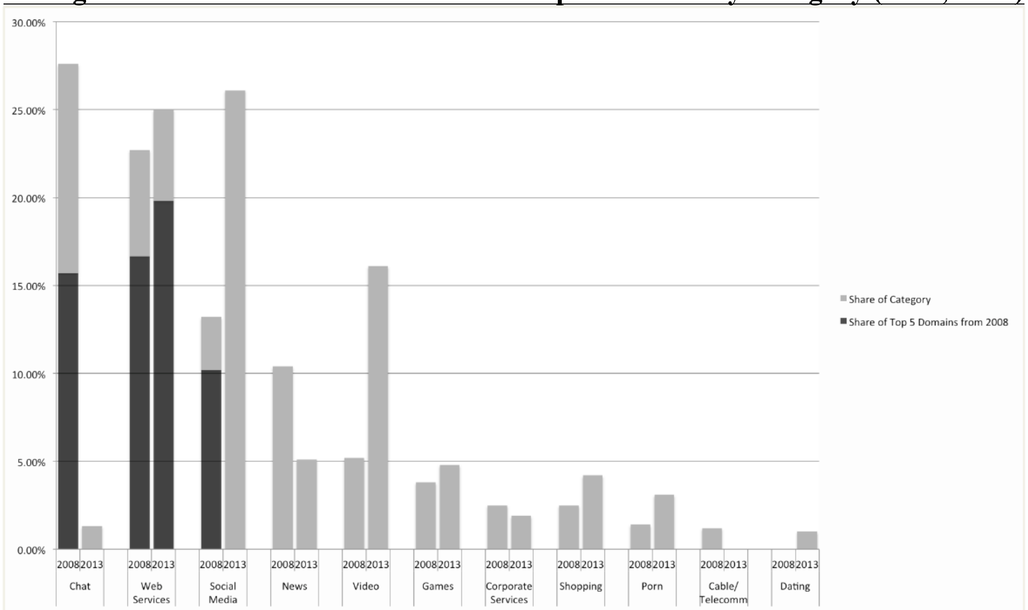

Figure 3

Unconditional Distribution of Online Attention (2008 vs. 2013)

2008	2013

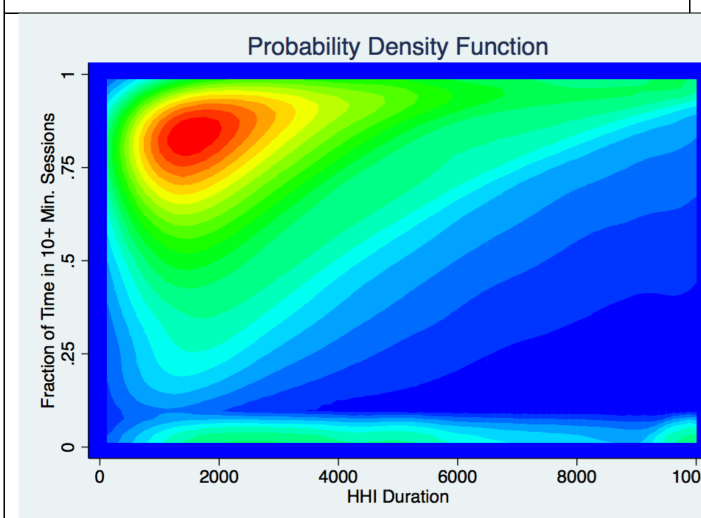

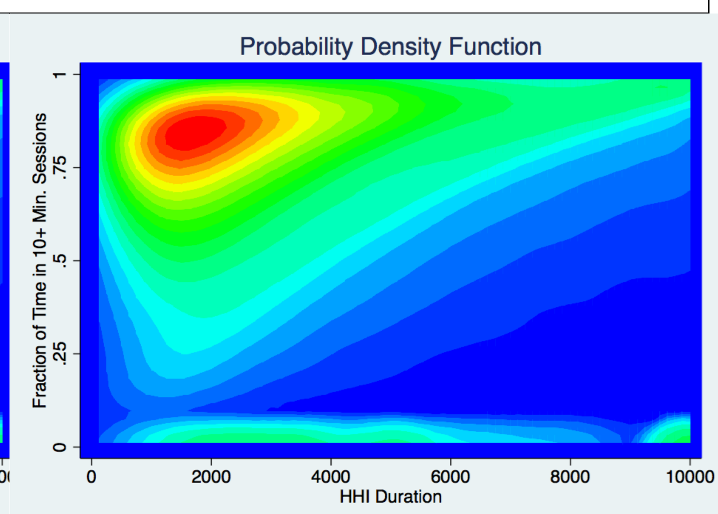

Figure 4

<u>**Distribution of Online Attention (2008 vs. 2013)**</u>

Broken Down by Total Attention

Quartile	**2008**	**2013**
1st	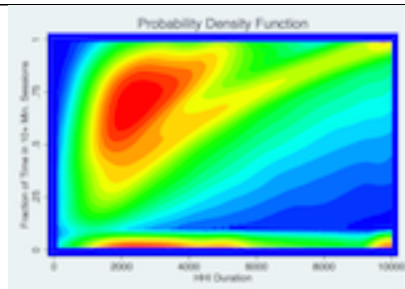	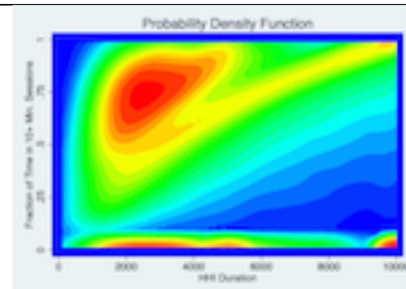
2nd	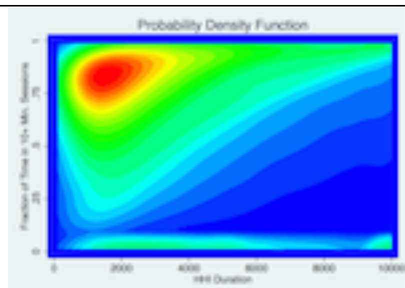	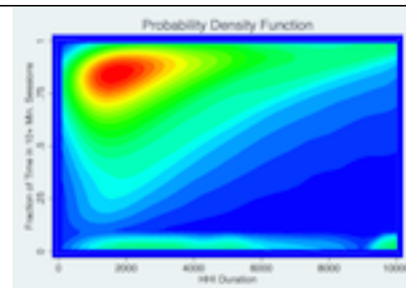
3rd	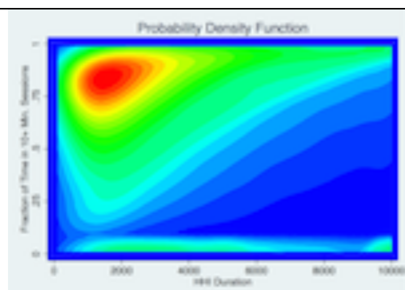	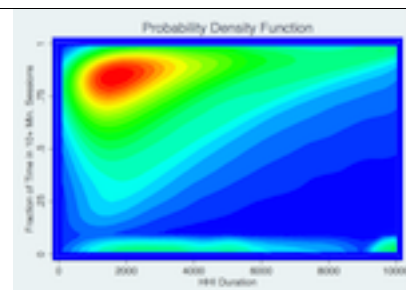
4th	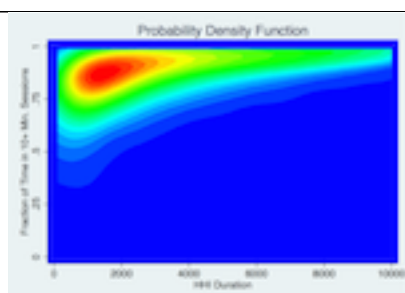	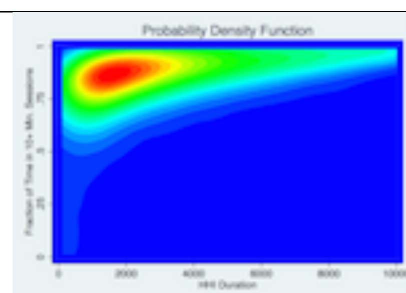

Figure 5

<u>Distribution of Online Attention (2008 vs. 2013)</u>

Broken Down by Extent of Multitasking

Quartile	**2008**	**2013**
1st	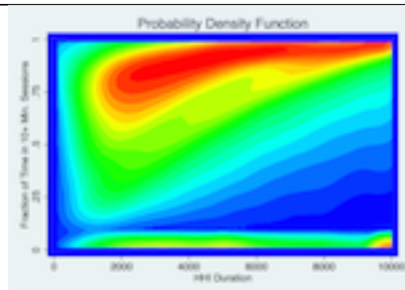	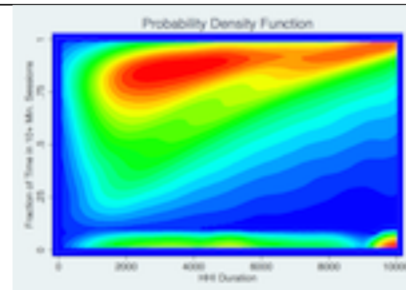
2nd	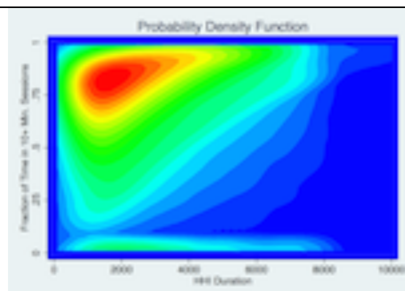	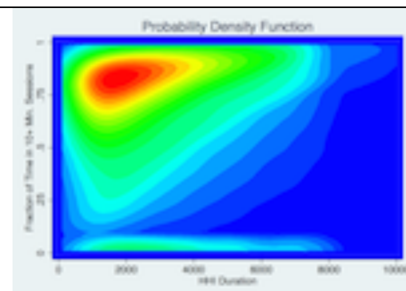
3rd	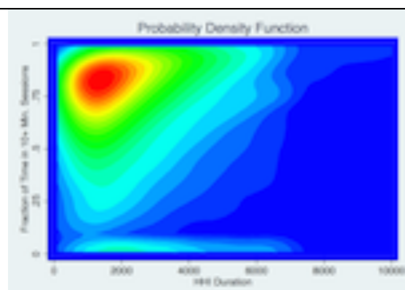	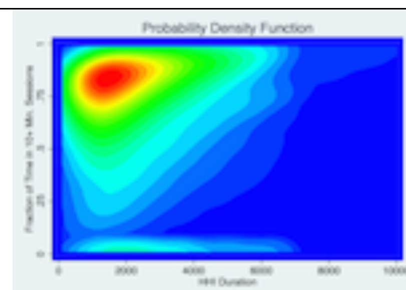
4th	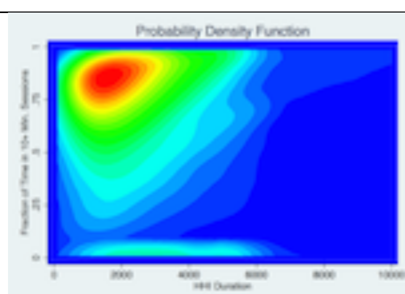	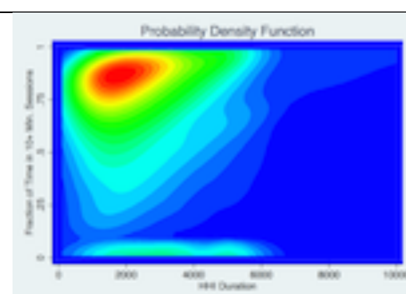

Table 1

Summary Statistics for Household Demographics

Variable	2008 N = 40,590		2013 N =32,750	
	Mean	Std. Dev.	Mean	Std. Dev.
Income < $15k	0.14	0.34	0.12	0.33
Income $15k-$25k	0.08	0.27	0.10	0.30
Income $25k-$35k	0.09	0.29	0.11	0.31
Income $35-$50k	0.11	0.31	0.15	0.35
Income $50-$75k	0.23	0.42	0.21	0.40
Income $75-$100k	0.16	0.36	0.13	0.34
Income $100k+	0.20	0.40	0.19	0.39
Age of Head of Household 18-20	0.00	0.07	0.05	0.21
Age of Head of Household 21-24	0.02	0.14	0.07	0.26
Age of Head of Household 25-29	0.05	0.22	0.08	0.27
Age of Head of Household 30-34	0.07	0.26	0.10	0.30
Age of Head of Household 35-39	0.11	0.31	0.08	0.28
Age of Head of Household 40-44	0.15	0.35	0.10	0.31
Age of Head of Household 45-49	0.17	0.38	0.12	0.33
Age of Head of Household 50-54	0.15	0.35	0.12	0.33
Age of Head of Household 55-59	0.10	0.30	0.09	0.29
Age of Head of Household 60-64	0.07	0.25	0.07	0.25
Age of Head of Household 65+	0.10	0.30	0.12	0.32
HH size = 1	0.07	0.25	0.12	0.32
HH size = 2	0.34	0.47	0.25	0.43
HH size = 3	0.25	0.43	0.21	0.40
HH size = 4	0.18	0.39	0.19	0.39
HH size = 5	0.11	0.31	0.16	0.37
HH size = 6+	0.05	0.22	0.07	0.27
Education < High School	0.00	0.01	0	0
Education High School	0.00	0.06	0.03	0.17
Education Some College	0.00	0.06	0.19	0.40
Education Associate Degree	0.00	0.02	0.16	0.37
Education Bachelor's Degree	0.00	0.06	0.11	0.32
Education Graduate Degree	0.00	0.04	0.01	0.08
Education Unknown	.99	0.11	0.49	.50
Children Dummy	.68	.47	.73	.44

Table 2

Summary Statistics for Attention Measures

	Year = 2008 *N =1,721,820*			
Variable	**Mean**	**S.D.**	**Min**	**Max**
Total attention per week (mins)	884	1281	1	10080
Attention concentration (HHI)	2868	2026	33	10000
Fraction of Avg Expenditure > 10	0.75	0.23	0	1
	Year = 2013 *N = 1,360,683*			
Total attention per week (mins)	849	1091	1	10078
Attention concentration (HHI)	2968	2061	1.51	10000
Fraction of Avg Expenditure > 10	.76	.22	0	1

Table 3

Linear Regression - Time Per Week on Demographics

	2008	**2013**
Covariate	**Minutes per Week**	**Minutes per Week**
Income \$15k-\$25k	-80*** (-3.83)	-19 (-0.95)
Income \$25-\$35k	-73*** (-3.57)	-19 (-0.96)
Income \$35k-\$50k	-91*** (-4.73)	-79*** (-4.49)
Income \$50k-\$75k	-118*** (-7.16)	-85*** (-5.08)
Income \$75k-\$100k	-131*** (-7.46)	-95*** (-5.25)
Income \$100k+	-166*** (-9.90)	-124*** (-7.14)
Education High School	262 (1.84)	-
Education Some College	289 (1.97)	18 (0.64)
Education Associate Degree	189 (1.12)	13 (0.46)
Education Bachelor's Degree	348 (2.34)	80** (2.72)
Education Graduate Degree	248 (1.63)	131 (1.91)
HH Size = 2	-8 (-0.38)	-35* (-2.03)
HH Size = 3	10 (0.44)	-35 (-1.86)
HH Size = 4	27 (1.14)	-10 (-0.48)
HH Size = 5	75** (2.86)	1 (0.05)
HH Size = 6	114*** (3.69)	-21 (-0.87)
Age of Head of Household 21-24	-387*** (-4.20)	9 (0.34)

Age of Head of Household 25-29	-434*** (-4.88)	-16 (-0.62)
Age of Head of Household 30-34	-478*** (-5.42)	-36 (-1.47)
Age of Head of Household 35-39	-402*** (-4.58)	-21 (-0.84)
Age of Head of Household 40-44	-361*** (-4.11)	-18 (-0.71)
Age of Head of Household 45-49	-382*** (-4.36)	41 (1.69)
Age of Head of Household 50-54	-408*** (-4.66)	53* (2.12)
Age of Head of Household 55-59	-502*** (-5.71)	14 (0.54)
Age of Head of Household 60-64	-531*** (-6.01)	11 (0.40)
Age of Head of Household 65+	-551*** (-6.28)	15 (0.59)
Children	3 (0.25)	132*** (10.46)
Constant	959*** (6.12)	800*** (21.53)
R-Squared	0.01	0.01
N	1,710,147	1,359,331

Table 4

SUR – Fraction of Sessions > Ten Minutes and Time HHI Across Sites

	2008	2008	2013	2013
Covariate	**HHI**	**Fraction > 10**	**HHI**	**Fraction > 10**
Income $15k-$25k	10 (1.37)	-0.00*** (-3.84)	22** (2.98)	0.00* (2.45)
Income $25-$35k	7 (0.99)	-0.01*** (-11.54)	1 (0.10)	-0.00 (-0.04)
Income $35k-$50k	-8. (-1.32)	-0.01*** (-14.78)	11 (1.57)	-0.00*** (-4.20)
Income $50k-$75k	-30*** (-5.52)	-0.01*** (-19.44)	16* (2.51)	-0.00*** (-4.10)
Income $75k-$100k	-1 (-0.26)	-0.01*** (-23.77)	-28*** (-3.94)	-0.00 (-1.76)
Income $100k+	-43*** (-7.61)	-0.02*** (-28.05)	-14* (-2.12)	-0.00*** (-6.68)
Education High School	624*** (4.30)	0.09*** (6.17)	-	-
Education Some College	530*** (3.65)	0.07*** (5.01)	-12 (-1.08)	-0.01*** (-10.18)
Education Associate Degree	403* (2.49)	0.10*** (6.05)	-65*** (-5.85)	-0.01*** (-11.85)
Education Bachelor's Degree	299* (2.05)	0.09*** (5.95)	-99*** (-8.60)	-0.01*** (-9.63)
Education Graduate Degree	309* (2.10)	0.10*** (6.33)	-126*** (-5.32)	-0.02*** (-6.70)
HH Size = 2	-44*** (-6.54)	-0.00 (-0.59)	-20** (-2.84)	-0.00 (-0.29)
HH Size = 3	-58*** (-7.21)	-0.00 (-0.30)	-18* (-2.34)	-0.00 (-0.71)
HH Size = 4	-71*** (-8.68)	0.00 (0.53)	-18* (-2.19)	0.00 (1.34)
HH Size = 5	-103*** (-11.75)	0.00** (2.94)	-36*** (-4.31)	-0.00 (-0.52)
HH Size = 6	-235** (-22.92)	0.00*** (4.31)	-50*** (-5.16)	-0.00 (-1.59)
Age of Head of Household 21-24	87** (3.25)	-0.00* (-2.57)	-20 (-1.85)	-0.00*** (-3.62)
Age of Head of Household 25-29	50* (2.00)	-0.01* (-2.41)	-33** (-3.15)	-0.01*** (-7.44)
Age of Head of Household 30-34	100*** (4.03)	-0.00 (-1.06)	-0 (-0.02)	-0.00 (-0.78)
Age of Head of Household 35-39	105*** (4.27)	0.00 (0.90)	-8 (-0.77)	-0.00* (-2.54)
Age of Head of Household 40-44	185*** (7.51)	0.00 (1.52)	51*** (5.12)	-0.00*** (-4.26)
Age of Head of Household 45-49	232*** (9.43)	0.00 (0.92)	-0 (-0.04)	-0.00*** (-4.38)
Age of Head of Household 50-54	233*** (9.47)	-0.00 (-0.81)	-48*** (-4.87)	-0.01*** (-6.22)
Age of Head of	199*** (8.04)	-0.01*** (-3.47)	20* (1.98)	-0.01*** (-6.16)

Household 55-59				
Age of Head of Household 60-64	304*** (12.18)	-0.01* (-2.49)	16 (1.52)	-0.01*** (-4.81)
Age of Head of Household 65+	360*** (14.56)	-0.01** (-2.78)	53*** (5.41)	-0.01*** (-7.30)
Children	-59*** (-12.78)	-0.00 (-1.66)	-142*** (-27.01)	-0.00 (-1.05)
Minutes per Week	-0 (-0.37)	0.00*** (531.17)	-0*** (-181.12)	0.00*** (438.72)
Constant	2652*** (18.26)	0.62*** (41.25)	3346*** (228.47)	0.71*** (473.26)
N	1,710,147	1,710,147	1,359,331	1,359,331
R-Squared	0.00	0.14	0.03	0.13

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note that across years the education dummies are relative to no high school in 2008 and relative to high school in 2013. Std. errors not clustered.

Table 5

The Top 20 Sites of 2008 and 2013 (by Total Time Allocated)

Ranking #	**2008 Top 20 Sites**	**Category**	**2013 Top 20 Sites**	**Category**
1	myspace.com	Social Media	facebook.com	Social Media
2	yahoo.com	News	youtube.com	Video
3	yahoomessenger.exe	Chat	google.com	Web Services
4	aim6.exe	Chat	yahoo.com	News
5	google.com	Web Services	tumblr.com	Personal Blog
6	msnmsgr.exe	Chat	msn.com	News
7	youtube.com	Video	aol.com	News
8	msn.com	News	craigslist.org	Shopping
9	aol.com	News	bing.com	Web Services
10	aim.exe	Chat	ebay.com	Shopping
11	facebook.com	Social Media	amazon.com	Shopping
12	live.com	News	twitter.com	Social Media
13	msn.com-prop	Chat	yahoomessenger.exe	Chat
14	myspaceim.exe	Chat	go.com	Sports
15	ebay.com	Shopping	wikipedia.org	Web Services
16	waol.exe	Chat	live.com	News
17	starware.com	Corporate Services	skype.exe	Chat
18	pogo.com	Games	reddit.com	Social Media
19	craigslist.org	Shopping	outlook.com	Web Services
20	go.com	Sports	netflix.com	Video