TECHNOLOGY POLICY INSTITUTE

Distinguishing Bandwidth and Latency in Households' Willingness-to-Pay for Broadband Internet Speed

August 2017

Yu-Hsin Liu, Jeffrey Prince, and Scott Wallsten

Distinguishing Bandwidth and Latency in Households' Willingness-to-Pay for Broadband Internet Speed

Yu-Hsin Liu, Jeffrey Prince, and Scott Wallsten*

August 2017

Abstract

We measure households' willingness-to-pay for changes in key home broadband Internet connection features using data from two nationally administered, discrete choice surveys. Both surveys include price, data caps, and download and upload bandwidth, but only one includes latency. Together, these surveys allow us to measure tradeoffs between bandwidth and other connectivity features such as price and data caps, and perhaps most notably, provide the only empirical evidence to date of tradeoffs between bandwidth and latency. We find that households' valuation of bandwidth is highly concave, with relatively little added value beyond 100 Mbps. For example, households are willing to pay about $2.34 per Mbps ($14 total) monthly to increase bandwidth from 4 Mbps to 10 Mbps, $1.57 per Mbps ($24) to increase from 10 to 25 Mbps, and only $0.02 per Mbps ($19) for an increase from 100 Mbps to 1000 Mbps. We also find households willing to pay about $8.66 per month to reduce latency from levels obtained with satellite Internet service to levels more common to wired service. Household valuation of increased data caps is also concave as caps increase from 300 GB to 1000 GB, although consumers place a significant premium on unlimited service. Our findings provide the first relative valuation of bandwidth and latency and suggest that current U.S. policy may be over-penalizing latency relative to reductions in bandwidth and data caps. For example, we find that in its CAF Phase II Auction, the FCC is imposing a bidding penalty for latency that is about five times higher than what our WTP estimates suggest it should be relative to bandwidth offered.

---

* Liu and Prince are at Indiana University, Department of Business Economics and Public Policy, Kelley School of Business. Wallsten is at the Technology Policy Institute. We thank several focus groups for their time and insight and ResearchNow for assistance in administering the surveys. We are responsible for all errors.

1. Introduction

Broadband Internet is typically considered an always-on, high-speed Internet connection. Speed in particular has come to be viewed as especially important, with Internet Service Providers (ISPs) competing based on the speeds they can offer. As of 2015, ISPs were spending tens of billions of dollars annually on broadband investment, much of this going toward speed improvements (USTelecom). In addition, policymakers have become increasingly concerned about whether consumers are getting broadband that is "fast enough." These concerns have led the Federal Communications Commission (FCC) to increase its definition of "broadband" (assessed in terms of bandwidth) several times—from 200 Kbps[1] in either direction (i.e., downstream or upstream) in 1999 to 4 Mbps downstream and 1 Mbps upstream in 2010, to 25 Mbps down and 3 Mbps up in 2015.

FCC changes to the definition of broadband are not meaningless, bureaucratic exercises. They can, for example, affect how the Commission views competition, which affects merger and other decisions. In addition, the FCC allocates several billions of dollars per year through the Universal Service Fund. For some of those funds, the FCC uses auctions that are handicapped based on the speed offerings of the bidders.

Despite its prominence in policy and private investment, little is known about how much consumers value incremental increases in bandwidth (particularly at the high end), and almost nothing is known about how they value other features of broadband speed and quality, such as latency and data caps. While bandwidth is the amount of data that can be transmitted over a connection in a given amount of time, latency is the time it takes for a data packet to make the

[1] For dial-up connection, the abbreviation "Kbps" stands for kilobits per second; for broadband connection, the abbreviation "Mbps" stands for megabits per second.

round trip between the user's computer and another computer, typically a server located somewhere else. Thus, speed as understood colloquially is not just bandwidth, but a combination of bandwidth and latency. It is necessary to understand how consumers value those components separately in order to learn how they value "speed."

The value of upgrading from a dial-up connection (e.g., with download bandwidth of 56 Kbps) to a low-end broadband connection (e.g., with download bandwidth of 1 Mbps) was likely large (Greenstein & McDevitt 2011). However, the value of upgrading from, say, a broadband connection bandwidth of 10 Mbps to 100Mbps may not be. As an example, downloading a 4 megabyte[2] file (the size of an average song) would take approximately 9.5 minutes at 56Kbps, 32 seconds at 1 Mbps, 3 seconds at 10 Mbps, and 0.25 seconds at 100 Mbps. It is easy to imagine that a consumer will value reducing the amount of time required to download from 9.5 minutes to 32 seconds more than she will find in reducing the time from 3 seconds to 0.25 seconds, despite the former time reduction corresponding to a bandwidth increase of less than 1 Mbps and the latter time reduction corresponding to a bandwidth increase of 90 Mbps.

Nevertheless, the pace of increase of top-level fixed-line Internet bandwidth does not appear to have slowed as of this writing, as evidenced by rollout of 10 Gbps in some parts of the United States (Fung, 2015). Nonetheless, as with the PC, one could again reasonably ask whether bandwidth increases continue to generate substantial value for consumers. Even according to the FCC, which has advocated for more bandwidth, common activities generally require relatively little bandwidth (in Mbps): Phone calls (< 0.5), Email (0.5), Watching movies (1.5), Watching HD movies (4.0), and Two-way online gaming (4.0). Hence, the vast majority of online activities appear to be accessible with 5.0 Mbps or less. A consumer who values watching video—and

---

[2] 1 megabyte (MB) = 8 megabits (Mb)

most do—would therefore value the bandwidth increase from a connection too slow for smoothly streaming video to one that allows a good SD video stream, and probably continue to value an increase to 5 Mbps so she can watch HD video. But beyond 5 Mbps, the value this subscriber obtains from further increases will depend on factors like how many connected devices operate at a single time and the types of connections she requires.

Variation in latency – the other key component to Internet speed – differs from bandwidth. For latency, the key variation is not across subscription plans, but across technologies, particularly fixed-line and satellite. Nonetheless, a similar question arises: does this difference in latency meaningfully translate to consumer value? And if so, how does this value compare to the value of bandwidth?

In this paper, we measure households' valuations of Internet speed. To do this, we first identify and define the features of Internet connection that together determine speed: bandwidth and latency. Next, we design and execute discrete choice surveys administered to over 1,000 individuals. These surveys allow for variation in the key features of home Internet connection: bandwidth (download and upload), latency, data cap, and price. A survey allows us to generate a degree of variation in bandwidth and latency that we generally do not see in the marketplace, thereby facilitating empirical analyses.[3] We administer two surveys: one that includes latency as a connection feature and one that omits it. By administering these two separate versions, we can assess whether consumers value bandwidth differently depending on whether they also consider latency.

[3] The availability of high bandwidth is subject to regional infrastructure, so market data generally does not have the desired variation for our purpose. Even if some market data on high bandwidth can be observed, we likely have a significant selection issue (i.e. some areas are selected for building the fiber infrastructure).

We find, similar to previous analyses, that consumers highly value bandwidth enhancements at lower speeds, but the incremental value of bandwidth decreases rapidly (Figure 1). In particular, households on average value increasing bandwidth from 4 Mbps to 10 Mbps at about $14 ($2.34/Mbps), 10 to 25 Mbps at $24 ($1.57/Mbps), 25 to 50 Mbps at $14 ($0.57/Mbps), 50 to 75 Mbps at $8 ($0.32/Mbps), and 75 to 100 Mbps at $4 ($0.16/Mbps). Households were willing to pay only an additional $19 ($0.02/Mbps) for bandwidth increased from 100 Mbps to 1Gbps.

Figure 1: Estimated WTP for Bandwidth Increments Above 4 Mbps

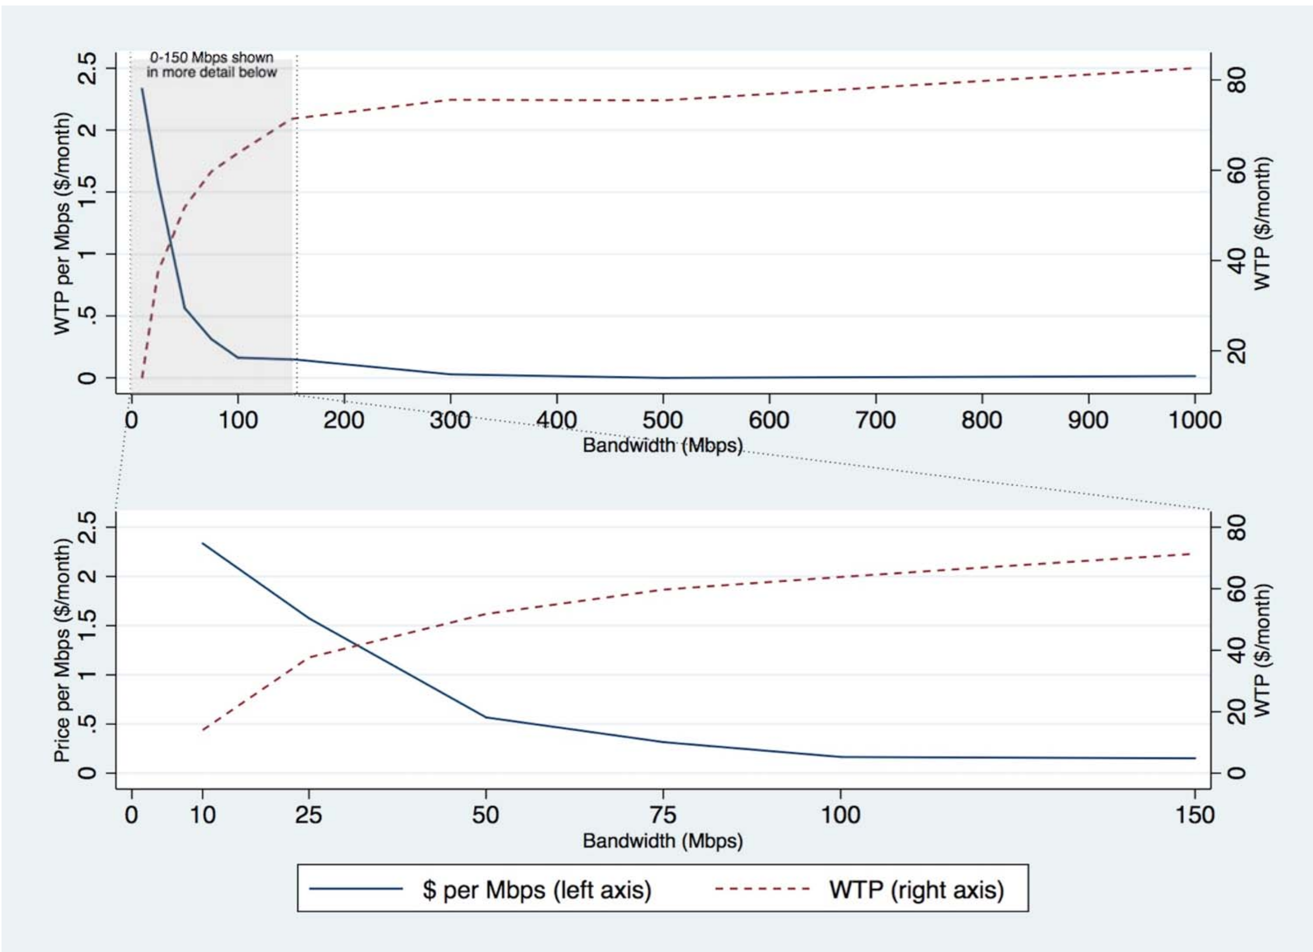

Household valuation of allowed data transmission (data caps) is similarly concave, but consumers appear to place a premium on unlimited data transfer. Households valued an increase in a data cap from 300 GB to 600 GB at $12 ($0.04/GB), another $11 to increase the cap from 600 GB to 1000 GB ($0.03/GB), and an additional $35 for an increase from 1000 GB to unlimited. This premium on unlimited is especially noteworthy, since households use only 190 GB of data per month on average as of 2016 (Telecompetitor, 2016).

Meanwhile, households on average value a decrease in latency from 300-600+ms to under 10ms (approximately equivalent to moving from a satellite to a wired connection of the same bandwidth) at about $8 per month.

Our results across the two surveys (with and without latency) indicate that failing to account for latency increases measured Willingness-To-Pay (WTP) for download bandwidth, but has no effect on measured WTP for upload bandwidth or data cap. Hence, consumers may interpret download bandwidth as a proxy for the combined quality of download bandwidth and latency when latency is not expressed.

Our findings also allow us to compare household valuation of the features of broadband quality. For example, consumers value increasing bandwidth from 10 to 25 Mbps about three times more than they value reducing latency from satellite to wired levels. In other words, while households value latency improvements, they are generally viewed as notably less valuable than bandwidth and data cap improvements, at least across observed market levels (e.g., 1,000 Mbps download bandwidth, 100 Mbps upload bandwidth and unlimited data cap).

We also find sensible differences in valuation across subgroups divided by demographics and usage. For example, as expected, respondents engaged in tasks such as gaming and file transfers place greater value on speed than those who don't.

Our findings have several important implications. First, they provide valuable information for broadband market players (e.g., telcos, satellite, and cable companies) about the value – and ultimately the likely profits – of investments that generate high speeds, as well as which types of improvements generate the most value. In doing so, they provide the first relative valuation of bandwidth and latency. They also suggest that an omitted feature, at least in this setting, can lead to biased estimates, particularly when it is a component of a broader product attribute (speed). The bias appears to be contained to only a subset of the other features comprising "speed" (download bandwidth).

Our findings also help to inform policies aimed at generating higher available speeds by identifying the extent of at least some of the benefits, which can be weighed against the costs. They can be particularly informative for auction design when allocating broadband subsidies, where bids are handicapped based on speed and data-cap provision. As a specific example, we find that Federal Communications Commission (FCC) Phase II auctions implicitly valued latency relative to bandwidth and data caps about five times what our valuation estimates indicate. Hence, current U.S. policy may be over-penalizing latency relative to reductions in bandwidth and data caps.

1.1 Related Literature

Our paper relates to a substantial stream of literature on Internet demand, which has been measured in different ways, using both market data and survey/experimental data. Studies using market data include Goolsbee and Klenow (2006), who use the variation of opportunity cost, measured by individual forgone wages, to estimate Internet demand in terms of time use (hours). Nevo et al. (2016) use the variation of shadow price, observed from the usage-based broadband plan in multiple billing cycles, to estimate Internet demand in terms of data use. One of the two key speed factors, bandwidth, is partially captured in Nevo et al. (2016), which generates utility only through the reduction of the marginal cost on consuming more data. Ahlfeldt et al. (2016) use market data on housing prices to capture the part of consumer surplus generated via broadband access. Their results show that broadband access increases property price significantly, while the marginal effect diminishes to zero when the actual bandwidth reaches approximately 5 Mbps.[4] In our paper, we assess whether household willingness-to-pay (WTP) for bandwidth is substantial up to 1 Gbps, importantly covering speeds surrounding 25 Mbps – the FCC's current definition for "broadband."

A major challenge of the aforementioned studies using market data is their limited ability to pin down how much households value speed increments, let alone increments in separate components of speed. This limitation in empirical measurements is largely due to limitations in market data; it has been difficult or even impossible to get data on broadband sales and prices at varying speeds that has sufficient and exogenous variation (Rosston et al., 2010). The limitations of market data are especially apparent when measuring the WTP for very-high bandwidth, as the variation in availability creates a selection problem. Given these challenges with market data, we chose to build on the stream of literature utilizing surveys and experiments.

[4] Actual speed of 5 Mbps roughly corresponds to advertised speed of 20 Mbps as mentioned in the paper, indicating a significant gap between actual and promised services in UK during the data period.

Prior literature using surveys and experiments to measure the value of Internet speed generally focused on measuring bandwidth. In a late 90's experiment, Varian (2002) allowed the research team to vary the metered price exogenously for bandwidths ranging from 8-128 Kbps (6 levels). The experiment revealed that participants have a low WTP for improved speeds within this range. Varian concluded that "the problem with broadband is not *access* but *applications*." Rappoport et al. (2002) found a contradictory result using survey data in 2000: Broadband users no longer perceived dial-up as a substitute, as evidenced by the nearly zero cross-price elasticity. In a later study, Dutz et al. (2009) further found that from 2005-2008, the substitution pattern between dial-up and broadband becomes weaker, and the own-price elasticity of broadband becomes less elastic, indicating overall increases in the residential value of broadband. Additional studies assessed broader Internet performance factors (e.g. connection reliability, always-on feature) and the preferences of heterogeneous users (e.g. rural/urban, high/low technical ability, experienced/inexperienced Internet users) (Savage and Waldman 2005, 2009; Rosston et al., 2010).

A notable limitation of these studies stems from ambiguity regarding the meaning of Internet speed. Most of them regard bandwidth as equivalent to speed, but this assumption is not entirely adequate. To users, Internet speed means how quickly they can complete a given task. From this perspective, bandwidth is just one factor that determines Internet speed, as tasks involve more than just downloading and uploading files. Higher bandwidth may not necessarily lead to perceivable improvements in efficiency in many cases. In the most recent survey study of which we are aware, Rosston et al. (2010) describes the top speed level as "Blazing fast downloads and uploads. It is really great for gaming, watching high-definition movies, and instantly transferring large files." For highly dynamic activities, such as gaming or real-time communication, it is latency, not bandwidth that determines “speed” for a typical American household (Grigorik, 2013a). The WTP being captured using this description is thus a mixture of the value of both bandwidth and latency. Latency is also primarily responsible for web page loading time and a primary driver of speed for several other activities, due to the significant improvements in bandwidth in the recent decade. To the best of our knowledge, our paper is the first to economically evaluate the importance of latency.

2. Internet Speed and Policy

Internet speed has become an integral part of setting policy. The FCC has established definitions, in Mbps, of the minimum bandwidth to be considered broadband, and another definition used to be eligible for the $4.5 billion annually from the universal service program’s Connect America Fund. Yet, as discussed above and detailed below, speed is a function of both bandwidth and latency. The Connect America Fund Phase II Auction recognizes the dual nature of speed by taking into account proposed bandwidth and latency, but without any analytics supporting the Auction’s weighting of the two components.

This section first discusses the elements that define speed and then related policy issues.

2.1. Determinants of Speed

In the physical world, speed is the rate at which an object covers distance (e.g. miles per hour). Internet speed, however, has a less straightforward definition because the object being measured is not clearly defined. A common description of Internet speed is the download / upload (henceforth omitted) bandwidth to measure the efficiency of data transmission, with more bandwidth generally equated with faster Internet. Bandwidth is the maximum throughput of the

Internet path, usually measured in Mbps, so more bandwidth generally means faster download speeds.[5]

Download time, however, does not fully capture the duration of transmitting data from sender to receiver. Even before the download process begins, multiple two-way communication processes between sender and receiver confirm the necessary information for the task. These processes generate latency, which is the amount of time it takes for a data packet to make a round trip between sender and receiver, usually measured in milliseconds.[6] Unlike download time, which is proportional to data size, no matter how small a transmitting file is, latency is a fixed time cost. It thus is important to Internet speed, especially for high-frequency or high-interaction tasks. For example, the FCC (2016a) notes that latency is important especially for real-time, interactive tasks, such as voice / video over IP, online gaming, distance learning, VPN platform, and telemedicine.[7]

In contrast to bandwidth, the FCC has no benchmark for latency. However, in 2011, the FCC launched the Measuring Broadband America (MBA) program, which began collecting the relevant data and reporting latency for each ISP in the MBA program. The data show that latency varies by technology. In general, fiber has the best performance, followed by cable and DSL, though the difference among these terrestrial technologies is not significant. Satellite internet

[5] The FCC (2016b) assesses the difference between the actual download / upload bandwidth and advertised bandwidth, and finds that the average ratio (across ISPs and over time) of actual speed to advertised bandwidth is 105.6%, meaning that overall the advertised bandwidth is close to the true bandwidth, while in some cases the advertised bandwidth could be too conservative.

[6] A millisecond is a thousandth of a second. Probably a more familiar description of latency for online gamers is “ping,” often reported in the standard Internet speed test in addition to download and upload speed.

[7] Page 27-30. Note that “service consistency” is another important metric that the FCC recognizes yet does not provide a benchmark. We do not evaluate this feature as we focus on the dimensions of Internet speed, and there is no clear definition on consistency. Excluding this feature should not affect the accuracy of the WTP estimation on bandwidth and latency; see section 3.1 for details.

service, however, has about 20 times the latency of terrestrial ISPs due to the distance between satellites and the ground.

In some situations, such as downloading large files or streaming video, bandwidth is the crucial component of speed, while in others, such as gaming, latency is likely to be the key component. From a user's perspective, speed means the time it takes to complete a meaningful task, such as downloading a file, loading a webpage, making an online voice call, watching streaming video, etc. It is this concept of speed that may contribute to the willingness-to-pay for a given Internet connection. Therefore, both bandwidth and latency can play important roles in determining WTP.[8]

In sum, a full and precise characterization of Internet speed must include bandwidth and latency. The distinction between bandwidth and latency is particularly important because it can lead to different policy implications. If shortening latency generates higher value, policymakers may wish to shift, or at least add, focus to this component of speed. This may be especially true when bandwidth develops further and highly-interactive applications become increasingly popular. Beyond just informing policy, our estimation of the value of latency may also provide reference to ISPs on their quality improvement plans.

[8] In extreme cases, the value of improving one quality may dwarf the value of improving the other. For example, Grigorik (2013a) documents the Hibernia Express project that aims at decreasing transatlantic latency between financial markets in North America and Europe. The estimated cost for the project, which will reduce latency for traders by about 5 milliseconds, is over $400 million, or about $80 million per millisecond decrease.

2.2. Policy Issues

The FCC, other policymakers, and advocates have focused almost exclusively on bandwidth, steadily increasing the minimum bandwidth necessary to be considered "broadband."[9] Latency, however, has not received much government attention, nor do ISPs advertise it like they do bandwidth. This incomplete characterization of speed raises three policy concerns.

First, focusing on bandwidth alone may cause policymakers to focus on a metric that does not maximize consumer welfare. Second, when latency is included as a relevant criterion for policy, the FCC has no empirical evidence on how to incorporate it and how it should be weighed relative to bandwidth. Third, we have little evidence on how much consumers value increases in bandwidth and none on how much they value decreases in latency.

As an example of the importance of this research, consider the Connect America Fund Phase II Auction. This reverse auction is intended to distribute almost $2 billion to areas currently unserved by terrestrial broadband. Providers will bid for subsidies to provide broadband service to designated census blocks. Winners will not be chosen based on the bid alone, however. The auction process will create a score for each bidder's bid that is a function of the bandwidth, latency, and data caps the provider intends to offer. In particular, it will use the following formula to determine each bid's score: [10]

---

[9] When the FCC first started tracking broadband adoption, any connection that offered speeds exceeding 200 Kbps downstream was considered broadband, and later increased that to 768 Kbps downstream. In 2010 it changed the definition to 4 Mbps down, and in 2015 to 25 Mbps down and 3 Mbps up. The definition for being eligible for universal service funds through the Connect America Fund has also changed over time, with the current standard being at least 10 Mbps down and 1 Mbps up.

[10] Ibid., para. 15.

$$score = 100\frac{subsidy\ requested}{reserve\ price} + performance\ weight + latency\ weight$$

where the performance weight is defined by the speed and data cap, as listed in Table 1 and the latency weight as shown in Table 2.

Table 1: FCC CAF Phase II Auction Performance Weights[11]

Performance Tier	Bandwidth	Usage Allowance	Weight
Minimum	≥10/1 Mbps	≥ 150 GB	65
Baseline	≥25/3 Mbps	≥ 150 GB or U.S. median, whichever is higher	45
Above Baseline	≥100/20 Mbps	2 TB	15
Gigabit	≥ 1 Gbps/500 Mbps	2 TB	0

Table 2: Auction Latency Weights[12]

Latency	Requirement	Weight
Low	≤ 100 ms	0
High	≤ 750 ms & "mean opinion score" of ≥ 4	25

Thus, the lower the score, the more competitive the bid. While the Commission did not base the weights on any particular empirical analysis, the performance weights are consistent with the general finding that consumers see diminishing marginal returns to increased speeds, as illustrated in Figure 2.

The latency penalty, which effectively applies only to satellite providers, is consistent with consumers valuing lower latency over higher latency. The 25 point penalty, however, implies that the FCC believes, for example, that consumers value a 10/1 connection with low latency more than a 25/3 connection with high latency. This may be true, but no data exist to support or contradict it.

---

[11] Federal Communications Commission, "In the Matter of Connect America Fund ETC Annual Reports and Certifications," para. 17.
[12] Ibid.

Figure 2: Score as a Function of Download Speed

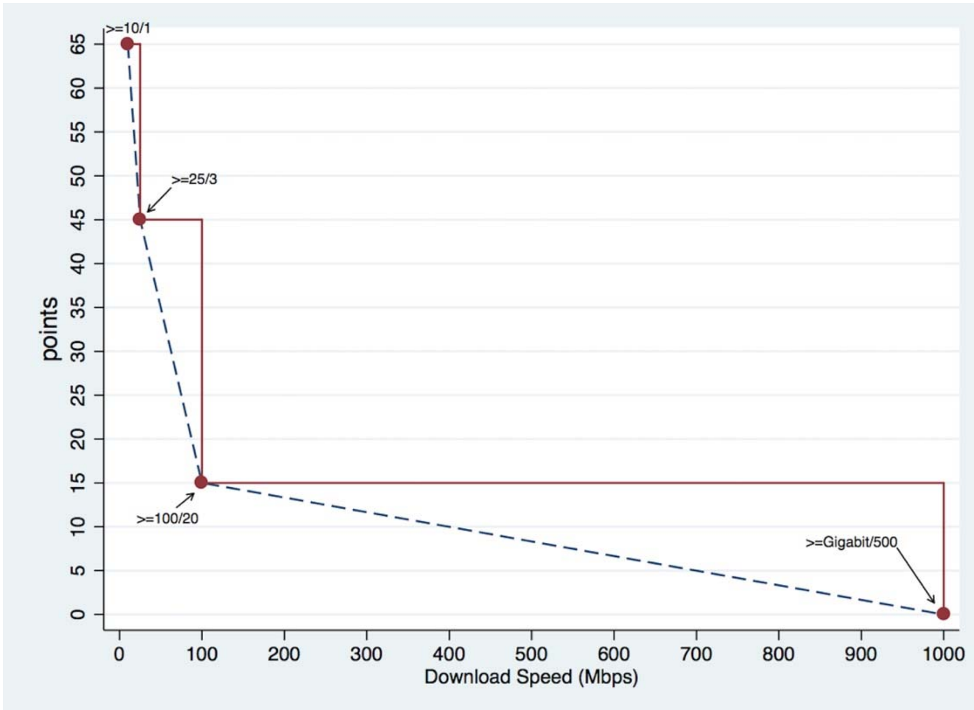

Given that the scoring formula is likely to affect the type of technology that receives subsidies in the auction, to the extent that consumer welfare is a consideration, it is crucial that the score incorporates consumer value of incremental changes in bandwidth and latency, as well as a proper relative valuation.

3. Survey Design

To estimate the WTP of the Internet speed components, we collect and analyze data from a nationwide online survey that employs repeated discrete choice experiments (DCEs). Prior work has shown that DCEs mitigate the reporting inaccuracy of stated-preference data (Carare et al., 2015). Even if hypothetical bias may potentially overestimate demand, the WTP estimation for changes in feature levels is statistically unbiased (Ding et al, 2005; Miller et al., 2011).[13] A reliable DCE method, however, requires a careful design to elicit respondents' truthful answers, as if they are making a choice in the real market (Ben-Akiva et al., 2016). We thus structure the survey in three parts. We first collect relevant demographic information in order to conduct comparative analyses and to ensure a representative sample based on region, age, race, and sex, based on recent Census figures. Demographics include geographic region, age, race, sex, education, employment status, and household income.

The second part of the survey collects information regarding each respondent's current Internet plan and usage. Only respondents who had a subscription and were the primary decision-maker in the household were included in the final sample. Before gathering this information, we provide respondents with cognitive buildup by describing each of the home Internet features about which we will inquire in this and the third part of the survey. That is, by recalling their daily Internet experience, we expect that this section will facilitate their later choice tasks, leading to more accurate answers. The cognitive buildup is in the Appendix, and was carefully vetted by several focus groups.[14] Following the cognitive buildup, we ask respondents about Internet-connected devices in the home, various uses for home Internet, and information on their current home Internet plan. For the home plan, we gather information on provider, bundling, price, download bandwidth, upload bandwidth, latency, and data cap.

[13] Specifically, Miller et al. (2011) attribute the upward bias to the non-incentive-aligned feature of the choice experiment. Participants do not need to actually pay for their choice in the hypothetical experiment and hence understate the possibility of choosing "none." The result is the biased demand intercept, but not the slope parameters, so WTP estimations remain valid in their test. In the context of Internet service, an incentive-aligned design is not possible due to high product cost, e.g. we cannot realistically offer a fiber-level service if no such infrastructure exists. On the other hand, we suspect that the Internet has become a necessity for most households, even at a reasonably high price. The tendency to choose "none" is likely to be low, especially given we are surveying current subscribers.

[14] The highway analogy we use to explain bandwidth and latency was universally described as helpful and intuitive in understanding these speed components.

Based on interactions with focus groups, we recognized many respondents may not be aware of some of their current plan features, particularly latency, data cap, and upload bandwidth. Anticipating this challenge, we constructed an "approximate" current plan as follows. For latency, we approximate the average level of the latency for each respondent by combining data from SamKnows, which the FCC uses to measure several connection details including latency, with respondents' answers on their geographical information, ISP, and Internet plan. We estimate unknown upload bandwidth by using the download bandwidth if respondents have fiber Internet, and by using one-eighth of the download bandwidth figure for other technologies. We estimate an unknown data cap by assuming it is unlimited, as 92.4% of respondents indicate that they have never been charged or received notice for exceeding data limits. Finally, the monthly fee for Internet may be bundled with other services, such as TV, phone, and WiFi. In this case, we infer the Internet standalone price as a portion of the bundling price from a rule we constructed based on observed market pricing (See Table 3 below).

Table 3: Internet Pricing Rule for Bundlers

Download bandwidth	Bundled with phone	Bundled with TV	Bundled with phone + TV
Less than 50 mbps	80%	45%	40%
50-99 mbps	85%	55%	50%
100-299 mbps	90%	60%	55%
300 mbps and above	90%	65%	60%

The final part of the survey consists of repeated choice experiments. Here, we mimic the real market choice situation while exogenously varying our variables of interest – particularly, price, bandwidth, and latency. In the DCEs, individuals make a series of choices over hypothetical Internet plans (alternatives), defined by a set of attributes. Since our primary goal is to estimate the WTP for bandwidth and latency, the core attributes are price, download bandwidth, upload bandwidth, and latency. Additionally, we include data cap to account for its effect on the value of bandwidth (Nevo et al., 2016). We vary and describe their levels based on market practices, as shown in Table 4.

In principle, we could include other common attributes, such as installation service quality or reliability, to more completely describe the service attributes. However, we do not include these additional characteristics for several reasons. First, their omission actually better approximates market choices. Features such as installation service quality and reliability are relatively difficult to measure and are seldom, if ever, highlighted in plan descriptions. Second, it is not clear that omitting these additional features raises any endogeneity concerns. Specifically, we may worry that the level of one included feature may imply a certain level for another omitted feature, thus confounding our ability to accurately measure the value of the included feature. Such implied relationships seem unlikely for the features we omit, and was confirmed by our focus groups. Last, we did not want to risk cognitive overload for our respondents, the risk of which increases as we include more varying attributes.

Table 4: Attributes, Descriptions, and Levels

Attributes	Descriptions and Levels	Levels
Price	The monthly fee the household pays for home Internet Service	20 levels, ranging from $20 to $115 per month in increments of $5
Download Bandwidth	Download bandwidth here means the maximum Mbps your home Internet service can provide to your devices - All levels are good for basic email, web browsing, online shopping, music downloads, file sharing, and SD video streaming for a single device. - With 10 Mbps and above, the Internet is good for HD video streaming and online gaming for a single device - With 25 Mbps and above, the Internet is good for 4K UHD video streaming for a single device - To download a HD movie (5 GB), it takes 6.8 mins for a 100 Mbps Internet	10 levels: 4 Mbps, 10 Mbps, 25 Mbps, 50 Mbps, 75 Mbps, 100 Mbps, 150 Mbps, 300 Mbps, 500 Mbps, 1000 Mbps
Upload Bandwidth	Upload bandwidth here means the maximum Mbps your home Internet service can move from your devices to the Internet. The upload bandwidth is relevant for uploading large files, pictures, or videos to file-sharing site (e.g. Dropbox), social media (e.g. YouTube), or another user (e.g. via Skype).	4 levels: 1 Mbps, 3 Mbps, 25 Mbps, 100 Mbps
Latency	Shorter latency improves the responsiveness of web loading, online gaming, video streaming, and online video / voice call.	6 levels: Less than 10 ms, 10-30 ms, 30-60 ms, 60-150 ms, 150-300 ms, 300-600+ ms
Data Cap	The data cap per month for the given monthly fee. Additional charge will apply for additional data usage.	4 levels: 300 GB, 600 GB, 1000GB, unlimited

Each respondent is presented with eight different choice questions, as is common for such surveys.[15] In addition, to mitigate any endogeneity concern, we explicitly state that any omitted feature should be assumed to be identical across all alternatives. In other words, any omitted attributes are controlled for, i.e., held fixed, when making the comparison.

In addition to the 5-attribute survey described above (the "latency survey"), we also administered a 4-attribute version, with latency excluded (the "no-latency survey"). As latency is often omitted when describing Internet speed, we may expect that the general public tends to attribute the utility brought by latency to bandwidth. Hence, without explicitly controlling latency, changes in bandwidth may also encapsulate implicit changes in latency, causing overestimation of WTP for it. Put another way, such a false belief forces a positive correlation between bandwidth and "unobservables" (i.e., all other factors affecting utility for broadband, including latency), even though the variation of bandwidth is designed to be exogenous. The "no latency" survey also omits latency in the cognitive buildup and as a requested feature of the current Internet plan, but otherwise is identical to the "latency" survey. We provide a sample question from the latency survey in the Appendix (the no-latency questions look identical but for the removal of the latency feature).

We conclude this subsection with a brief description of our process for arriving at an optimal design, i.e., the construction of the levels for each attribute presented to each respondent for each choice. For a statistically optimal design, we rely on D-optimality (Zwerina et al. 2010), which we implement in SAS. We use a fractional factorial design to capture the main effects.[16]

---

[15] Discussion with focus groups indicated fatigue was not a factor for this number of choices.
[16] We specify that the interaction effects among bandwidth, latency, and data cap must not be confounding with other two-way interaction effects. We use SAS %mktruns and %mktex to produce candidate runs given our target sample size. We avoid dominated alternatives (i.e. better quality but lower price) and some too-unrealistic

Our relative D efficiency is 79% (75%) for the latency (no-latency) survey. The chosen design generates 120 (40) choice questions that are grouped into 15 (5) sets of 8 questions for the latency (no-latency) survey, and we define each set as a version. We randomly vary the order of the 8 questions as well as the order of the alternatives for each choice. We then randomly distribute the versions across respondents. Given that we have 978 (433) US households for the latency (no-latency) survey, the effective sample size is 7,824 (3,464).

4. Data

Our data come from ResearchNow's (RN) standing Internet panel. As we aim for 1,050 and 450 completed surveys for the latency and no-latency surveys respectively, RN makes sure that the target sample sizes are satisfied and the sample composition of responders is similar to the US census, particularly on geographical region, sex, ethnicity, and age. A qualified record requires the household respondent to be at least 18 years old, have a home Internet plan, be the primary decision-maker for the home Internet plan, know both the monthly fee and the download bandwidth for his/her home Internet plan, and take at least 4 minutes for completing the online survey.[17] Monetary incentives are provided to the respondents who complete the survey. The success rates are 31% and 47% for the latency and no-latency surveys respectively.[18] We then exclude some misleading responses due to potential misunderstanding (e.g. reported current

---

alternatives (e.g. 20 dollars a month for 1000 Mbps download bandwidth) by using the SAS %macro. We then evaluate and select the design by using SAS %choiceff.

[17] Our focus group result indicates that 10-15 mins is common in a rather interactive responding environment.

[18] 33% is disqualified and 36% is overquotaed for the latency survey.. 42% is disqualified and 11% overquotaed for the no-latency survey. Overquota refers to over sampling for some demographic groups, relative to the US population composition. The reason for overquota is because some groups (in our case, Hispanic) are relatively difficult to reach, especially when we require a large sample size. To reach them RN repeats the sampling procedure multiple times until the targeted quota of all demographic groups is satisfied. The overquota sample is then discarded in the analysis.

price / bandwidth is zero, data cap is smaller than 100 GB). This further selection ensures our data quality and only modestly reduces our sample size.

4.1. Summary Statistics

Our final dataset includes 978 U.S. households for the latency survey and 433 U.S. households for the no-latency survey. The demographic distributions are similar across surveys, which is crucial for ensuring that we can compare results across surveys meaningfully. The surveys are also reasonably representative of the US population, with certain expected exceptions given that we only retained respondents who were the primary decision-makers in the purchase of their home Internet. Differences with Census are evident for age, education, employment status, and income. Hence, our sample appears to be reasonably representative of the population of decision-making, home-Internet plan purchasers. The full demographic summary statistics are in Table 5.

Table 5: Respondent Demographics, Summary Statistics[19]

	US Census	Latency Survey	No-latency Survey
Region			
Northeast	17.4%	20.4%	17.3%
Midwest	21.0%	21.6%	21.7%
South	37.9%	35.1%	37.8%
West	23.7%	22.9%	23.3%
Age			
18-24 years	9.6%*	12.1%	9.0%
25-34 years	18.2%	18.5%	19.8%
35-44 years	17.3%	19.3%	23.7%
45-54 years	18.8%	20.2%	19.6%
55-64 years	16.8%	17.2%	16.6%
65 years or above	19.1%	12.7%	11.3%
Race			
African American	12.3%	11.5%	12.4%
Asian and Pacific Islander	5.3%	5.6%	5.5%
Caucasian	62.3%	67.6%	66.4%
Hispanic	17.1%	12.6%	11.8%
Native American	0.7%	0.6%	1.2%
Others	2.3%	2.1%	2.8%
Sex			
Female	50.8%	50.9%	50.7%
Male	49.2%	49.1%	49.3%
Education			
< High School	13.3%	0.9%	3.2%
High School Diploma	27.8%	12.6%	13.6%
Some college	29.2%	27.3%	30.4%
Bachelor Degree	18.5%	35.0%	34.1%
Graduate Degree	11.2%	24.2%	18.7%
Employment			
In labor force	63.7%	74.8%	73.8%
Not in labor force	36.3%	24.4%	24.9%
Preferred Not to Say		0.9%	1.4%
Household Income			
< $15,000	12.5%	5.6%	2.1%
$15,000 - $24,999	10.6%	6.7%	5.3%
$25,000 - $49,999	23.5%	19.7%	22.6%
$50,000 - $74,999	17.8%	19.3%	18.9%
$75,000 - $99,999	12.1%	14.9%	15.7%
$100,000 - $149,999	13.1%	18.4%	17.7%
> $150,000	10.4%	9.4%	10.4%
Preferred Not to Say		6.0%	7.4%

[19] Rounding error may be present, *20-24 years. Source: US Census American FactFinder 2015, 2016

Table 6 contains information about households' device ownership and Internet usage. The rank order of device ownership (in terms of average units) is: mobile phones, laptops, tablets, streaming TV device, and desktop computer. About 10 percent of the respondents indicated that their household has an additional connected device(s) not included in our list, with the most common being a game console.

Table 6: Household Current Devices and Internet Usage, Summary Statistics[20]

	Latency Survey (obs. = 978)					**No-latency Survey (obs. = 433)**				
	Mean	SD	Min	Med	Max	Mean	SD	Min	Med	Max
Device*										
Mobile	2.06	1.23	0	2	5	2.13	1.14	0	2	5
Laptop	1.55	1.08	0	1	5	1.36	1.00	0	1	5
Tablet	1.32	1.14	0	1	5	1.72	1.04	0	2	5
Stream TV	0.97	1.20	0	1	5	0.98	1.09	0	1	5
Desktop	0.86	0.93	0	1	5	0.78	0.87	0	1	5
Usage**										
Email	0.97	0.16	0	1	1	0.95	0.21	0	1	1
Shopping	0.86	0.35	0	1	1	0.84	0.37	0	1	1
File sharing	0.83	0.37	0	1	1	0.84	0.37	0	1	1
Social media	0.81	0.39	0	1	1	0.82	0.38	0	1	1
Banking	0.80	0.40	0	1	1	0.78	0.40	0	1	1
Video streaming	0.73	0.44	0	1	1	0.71	0.45	0	1	1
News	0.70	0.46	0	1	1	0.69	0.46	0	1	1
Music streaming	0.56	0.50	0	1	1	0.56	0.50	0	1	1
Video call	0.49	0.50	0	0	1	0.55	0.50	0	0	1
Health info	0.45	0.50	0	0	1	0.48	0.50	0	0	1
Gaming	0.40	0.49	0	0	1	0.47	0.50	0	0	1
Gov. Info	0.37	0.48	0	0	1	0.40	0.49	0	0	1
Class	0.26	0.44	0	0	1	0.26	0.44	0	0	1
Telework	0.15	0.36	0	0	1	0.15	0.36	0	0	1

[20] *5 means 5 or more. **0 no and 1 yes

Regarding Internet usage, note that the survey asked respondents whether they do these activities, but not how much time they spend on each. Hence, we have information on the extensive, but not intensive, margin for each activity. Email is by far the most commonly used application, with near universality. Email is followed by shopping, file sharing, social media, banking, and video streaming. About fifty percent of households use video calling and music streaming services and forty percent use the Internet for gaming. The least popular uses of connected devices include getting government information (37%), doing school work (26%), or teleworking (15%).

Table 7 provides summary statistics for respondents' currently owned Internet plans. In terms of Internet technology, cable accounts for 58 percent of households' connections in the sample, followed by DSL (20%), fiber (19%), and satellite (1%). The remaining was unknown by the respondents. Fiber has the best average performance on all non-price measures, especially upload bandwidth. The tendency of switching to the best hypothetical choice (from the choice experiment) is also lower for households who currently have fiber than for households who currently have other technologies. Finally, we note that the overall likelihood of switching to the alternative plan is 0.46, suggesting that our hypothetical alternatives are generally comparable to existing Internet plans.

Table 7: Household Current Internet Plan, Summary Statistics[21]

	Mean	SD	Min	Median	Max	Obs
All						
Price ($)	64.84	32.74	10	60	290	1411
Down BW (Mbps)	149.56	198.49	1	75	1000	1411
Up BW (Mbps)	57.58	136.73	1	16	1000	1411
Latency	2.29	0.69	1	2	6	978
Unlimited Data	0.95	0.22	0	1	1	1411
Switch	0.46	0.50	0	0	1	1411
DSL						
Price ($)	58.86	34.78	10	50	200	286
Down BW (Mbps)	129.40	205.78	1	60	1000	286
Up BW (Mbps)	33.79	85.49	1	9	1000	286
Latency	2.99	0.49	1	3	6	191
Unlimited Data	0.97	0.17	0	1	1	286
Switch	0.47	0.50	0	0	1	286
Cable						
Price ($)	65.99	30.36	10	60	210	817
Down BW (Mbps)	151.96	184.52	1	75	1000	817
Up BW (Mbps)	34.08	83.63	1	12	1000	817
Latency	2.08	0.50	1	2	6	597
Unlimited Data	0.94	0.23	0	1	1	817
Switch	0.49	0.50	0	0	1	817
Fiber						
Price ($)	66.24	35.58	10	60	290	262
Down BW (Mbps)	164.28	228.10	1	75	1000	262
Up BW (Mbps)	160.115	233.44	1	75	1000	262
Latency	2.04	0.37	1	2	4	160
Unlimited Data	0.95	0.26	0	1	1	262
Switch	0.44	0.50	0	0	1	262
Satellite						
Price ($)	65.00	19.86	45	60	110	10
Down BW (Mbps)	151.00	303.53	5	50	1000	10
Up BW (Mbps)	109.00	313.14	1	10	1000	10
Latency	6.00	0.00	6	6	6	7
Unlimited Data	0.90	0.32	0	1	1	10
Switch	0.50	0.53	0	1	1	10

[21] latency 1: less than 10 ms, 2: 10-30 ms, 3: 30-60 ms, 4: 60-150 ms, 5:150-300 ms, 6: 300-600+ ms

5. Econometric Methods

We use the conditional logistic regression model (McFadden, 1974; Greene, 2012) to estimate utility parameters and ultimately calculate the WTP. Let $\boldsymbol{x}_{ijk}$ be a vector of attributes for alternative *j* in choice question *k* that individual *i* faces. A linear random utility model can be written as:

$$u_{ijk} = \boldsymbol{x'}_{ijk}\boldsymbol{\beta} + \varepsilon_{ijk}$$

We interpret the errors (ε_{ijk}) as individual idiosyncratic preference and assume that it is independently and identically distributed with type I extreme value distributions. With this assumption the probability for individual *i* to choose alternative *j* among 4 alternatives in question *k* is then

$$\text{Prob}(Y_{ik} = \text{j}) = \frac{\exp(\boldsymbol{x'}_{ijk}\boldsymbol{\beta})}{\sum_{j=1}^{4} \exp(\boldsymbol{x'}_{ijk}\boldsymbol{\beta})}$$

Since we observe individual choices in each question, we are able to generate the likelihood function based on these probabilities. We then optimize the likelihood function with respect to $\boldsymbol{\beta}$ and obtain the estimated utility parameters for each attribute, clustering our errors on individuals.

The calculation of WTP for attributes relies on $\boldsymbol{\beta}$. In our case, the attributes include price, download bandwidth, upload bandwidth, latency, and data cap. For illustration we partition $\boldsymbol{x'}_{ijk}$ into $[p_{ijk}, db_{ijk}, ub_{ijk}, l_{ijk}. dc_{ijk}]$ and the corresponding $\boldsymbol{\beta'}$ is $[\beta_p, \beta_{db}, \beta_{ub}, \beta_l, \beta_{dc}]$. Using this formulation, the point estimation of WTP for higher download bandwidth, for example, can be monetized using the estimated β_p and β_{db} in the following formula:

$$WTP(bandwidth) = -\frac{\beta_{db}}{\beta_p}$$

Note that bandwidth is an ordinal variable with multiple levels, and in our illustration we put an unnecessary restriction that each one-level increase generates the same change in utility, no matter the starting point. We relax this restriction in estimation by translating ordinal variables (bandwidths, latency, and data cap) into multiple dummies, allowing us to have different WTP estimates for changes that start at different levels. For example, we allow the download bandwidth to vary in twenty levels, ranging from 4 Mbps to 1000 Mbps. Hence, setting the baseline bandwidth to be 4 Mbps, we interpret the coefficients of the other nineteen dummies as the utility increase of a bandwidth improvement from 4 Mbps. This allows us to examine the shape of WTP for the bandwidth improvement. In our experiment design, we make sure that all the dummies can be identified. Finally, we estimate the variance of WTP by using a linear transformation of the variance-covariance matrix of $\boldsymbol{\beta}$, also known as the delta method.

As discussed earlier, a key merit of using a survey is the ability to generate sufficient variation in our variables of interest and clean identification of the underlying parameters. The use of a hypothetical environment, however, may also induce unrealistic responses that generate bias. To minimize this possibility, we carefully designed our survey to elicit respondents' preferences and mimic the real market situation, as described in Section 3. However, we cannot require participants to break their current Internet contracts and actually pay for the plan they choose in the survey.

To mitigate the potential hypothetical bias and further boost our effective sample size, we ask respondents to compare their best choice from the hypothetical question to their status quo plan. Similar hybridization of the survey and market data is suggested in Savage and Waldman

(2009) and Rosston et al. (2010). The simple way to incorporate the additional market data is to treat the status quo questions as another eight independent draws, and thus we simply append the likelihood function with these additional choices. The concern, however, is that the status quo is the same across all the eight status quo questions, forcing the logit error term to be correlated. To address the concern, we model this correlation pattern using an individual-level status-quo fixed effect, which we assume to be normally distributed across individuals. Hence, we include a status quo dummy variable in our linear utility model, with a random coefficient that is assumed to be random across individuals (not choices). Given this added structure, our utility equation becomes:

$$u_{ijk} = \boldsymbol{x'}_{ijk}\boldsymbol{\beta} + \gamma_{\boldsymbol{i}} SQ_{ijk} + \varepsilon_{ijk}$$

where γ is normally distributed with mean μ_{sq} and standard deviation σ_{sq}, both to be estimated. We use this appended utility model to construct probabilities and our likelihood function, which we estimate using maximum simulated likelihood[22].

Another concern when including the status quo question is the range of respondents' reported feature levels. As several of our Internet features, including price, download and upload bandwidths, and data cap are represented using dummy variables for discrete levels, the reported figures may not always match the predetermined dummies. Additionally, we cannot create and estimate coefficients for all the possible levels that respondents may report. One solution is to map the reported answer to the nearest predetermined dummy level; however, doing so would either overstate or understate the variation, creating undesired noise. To optimally use the information in the open-ended setting, we use linear interpolation between the two levels that

[22] The key assumption is that the remaining error term is noise that is redrawn (independently) on every choice occasion involving the status quo, just as in the hypothetical questions.

contain the reported value. For reported values that are smaller than our lowest dummy level or larger than our highest dummy level, we simply map to the nearest level. The precise coding rule for the interpolation is as the following:

$$\text{for } x \in [a, b], I_a = \frac{b - x}{b - a}, I_b = \frac{x - a}{b - a}$$

where *x* is the reported value (e.g. 12 Mbps), *a* and *b* are the predetermined dummy levels (e.g. 10 Mbps and 25 Mbps), I_a and I_b are the corresponding dummy variables in the utility function. For example, if *x = 12*, and the nearest dummy variables are I_{10} and I_{25}, we would have I_{10} = (2/15) and I_{25} = (13/15), and all other dummies concerning levels of *x* equal to zero. Essentially, we are assuming that the treatment effect of the continuous value *x* is the linear combination of the treatment effects of levels *a* and *b*.

6. Results

Table 8 contains our parameter estimates for both the latency and non-latency surveys. In columns (1) and (4), we use the responses of the hypothetical choice experiment only. As each respondent makes eight choices, the effective sample size is eight times of the number of respondents. Columns (2), (3), (5), and (6) include the follow-up status quo questions.[23] By including the status quo questions, we further use the revealed preference data and double the effective sample size, as suggested in Rosston et al. (2010). As can be seen across the columns within the two survey types (latency and no-latency), the parameter estimates are qualitatively identical and quantitatively similar across the three specifications. In what follows, we use the

[23] We include two alternatives in each status quo question. One is the chosen plan from the precedent hypothetical choice question, and the other one is the reported status quo plan, presented as the respondents viewed it.

parameter estimates from the model that uses all the data and allows for heterogeneous preferences for the status quo – the mixed logit model (columns (3) and (6)).

Table 8: Parameter Estimation Results[24]

	latency			**Non-latency**		
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
Variables	Clogit \wo status quo	Clogit \w status quo	mixlogit	Clogit \wo status quo	Clogit \w status quo	mixlogit
Price	-.0212***	-.0195***	-.0222***	-.0239***	-.0235***	-.0258***
Download						
10 Mbps	.3667***	.3523***	.3116***	.4728***	.3257**	.4556***
25 Mbps	.8972***	.7443***	.8371***	1.5495***	1.1279***	1.3796***
50 Mbps	1.1939***	1.0691***	1.1522***	1.4869***	1.2729***	1.4409***
75 Mbps	1.3725***	1.1976***	1.3279***	1.7516***	1.4990***	1.6806***
100 Mbps	1.4591***	1.2748***	1.4194***	1.9078***	1.6348***	1.8064***
150 Mbps	1.6271***	1.4207***	1.5874***	2.0491***	1.7550***	1.9780***
300 Mbps	1.6783***	1.5605***	1.6815***	2.1653***	1.9868***	2.1320***
500 Mbps	1.6450***	1.5679***	1.6787***	2.3329***	2.0809***	2.2961***
1000 Mbps	1.7695***	1.7282***	1.8371***	2.2577***	2.0829***	2.2780***
Upload						
3 Mbps	.2527***	.1881***	.2228***	.2150***	.1817***	.2324***
25 Mbps	.4158***	.3746***	.4130***	.3586***	.4013***	.4529***
100 Mbps	.5356***	.5145***	.5441***	.6297***	.7567***	.7680***
Latency						
10-30 ms	.1203***	.0505	.0841*			
30-60 ms	.0530	.0036	-.0081			
60-150 ms	-.0381	-.0663	-.0896*			
150-300 ms	-.0767	-.1245***	-.1435***			
300-600+ ms	-.1626***	-.1580***	-.1926***			
Data cap						
600 GB	.2969***	.2645***	.2716***	.2982***	.2753***	.2643***
1000 GB	.5289***	.4852***	.5243***	.6485***	.6069***	.6373***
Unlimited	1.3072***	1.1512***	1.2945***	1.4117***	1.2645***	1.3829***
Status Quo Dummy Mean		-.1893***	-.2329***		.0007	.0246
Status Quo Dummy SD			1.6151***			1.5868***

Note: *10% significance, **5% significance, ***1% significance

[24] Baseline download bandwidth 4 Mbps, baseline upload bandwidth 1 Mbps, baseline latency less than 10 ms, and baseline data cap 300 GB.

Table 9 contains our WTP estimates for various feature levels for both the latency and no-latency surveys. We find that, first, the marginal WTPs for both download and upload bandwidths decrease significantly as bandwidth increases. The WTP of a representative household for an improvement on the download bandwidth from 4 Mbps to 25 Mbps ($37.63) is just slightly less than the WTP improvement from 25 Mbps to 1000 Mbps ($44.96).

Table 9: WTP Estimation

	Latency Survey			**Non-latency Survey**		
Variables	beta	WTP	s.e.	beta	WTP	s.e.
Price	-.0222***			-.0258***		
Download						
10 Mbps	.3116***	$14.01***	$4.11	.4556***	$17.63***	$5.40
25 Mbps	.8371***	$37.63***	$3.82	1.3796***	$53.39***	$4.90
50 Mbps	1.1522***	$51.80***	$3.71	1.4409***	$55.76***	$5.03
75 Mbps	1.3279***	$59.70***	$3.79	1.6806***	$65.04***	$4.94
100 Mbps	1.4194***	$63.82***	$3.80	1.8064***	$69.90***	$4.98
150 Mbps	1.5874***	$71.37***	$3.77	1.9780***	$76.55***	$4.94
300 Mbps	1.6815***	$75.60***	$3.79	2.1320***	$82.51***	$5.04
500 Mbps	1.6787***	$75.47***	$3.87	2.2961***	$88.86***	$4.97
1000 Mbps	1.8371***	$82.59***	$3.83	2.2780***	$88.15***	$5.15
Upload						
3 Mbps	.2228***	$10.01***	$1.55	.2324***	$8.99***	$1.94
25 Mbps	.4130***	$18.57***	$1.57	.4529***	$17.53***	$2.04
100 Mbps	.5441***	$24.46***	$1.78	.7680***	$29.72***	$2.35
Latency						
10-30 ms	.0841*	$3.78**	$1.90			
30-60 ms	-.0081	-$0.36	$1.90			
60-150 ms	-.0896*	-$4.03**	$1.99			
150-300 ms	-.1435***	-$6.45***	$1.94			
300-600+ ms	-.1926***	-$8.66***	$2.00			
Data cap						
600 GB	.2716***	$12.21***	$1.76	.2643***	$10.23***	$2.54
1000 GB	.5243***	$23.57***	$1.74	.6373***	$24.66***	$2.58
Unlimited	1.2945***	$58.20***	$1.87	1.3829***	$53.51***	$2.65

Note: *10% significance, **5% significance, ***1% significance

Using these figures, we plot WTP for download bandwidth (relative to 4 Mbps) in Figure 3: Estimated WTP for Increase in Bandwidth (from 4 Mbps), which shows a steep rise for small improvements, followed by significant flattening around 100 Mbps. WTP for upload bandwidth shows a similar pattern. For example, the WTP of a representative household for an improvement in upload bandwidth from 1 Mbps to 3 Mbps ($10.01) is just slightly less than the improvement from 3 Mbps to 100 Mbps ($14.45).

Figure 3: Estimated WTP for Increase in Bandwidth (from 4 Mbps)

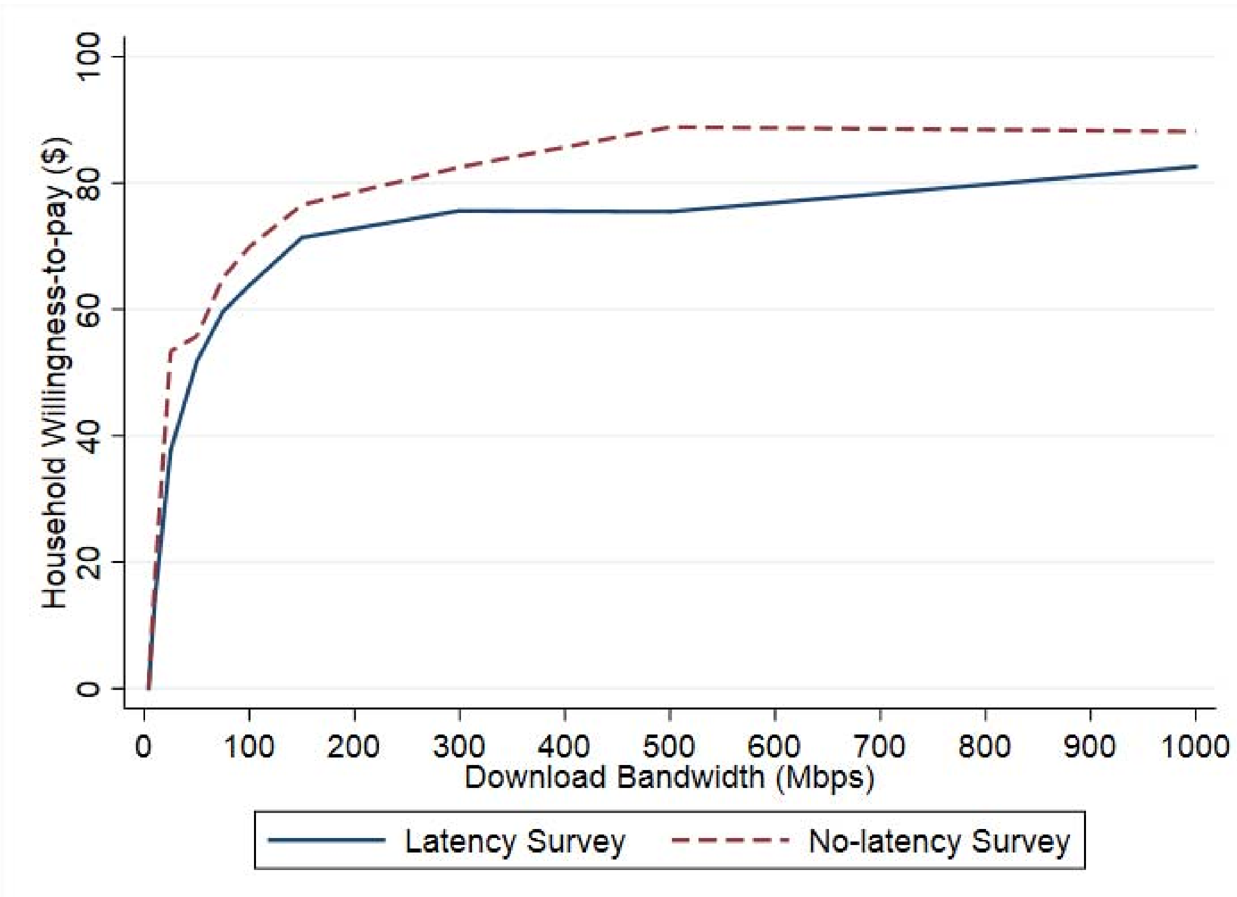

Regarding latency, households are willing to pay to avoid high latency levels. Users expressed a WTP of $8.66 to go from 300-600+ ms down to the baseline of less than 10 ms, and $4.03 to go from 60-150 ms down to the baseline. Lastly, we see a similar pattern for data caps as we do with bandwidth—decreasing marginal value of higher caps—but with a premium on unlimited data. In particular, households are WTP \$12.21 to increase their data cap from the baseline level of 300 GB to 600 GB (about \$0.04/GB), \$11.36 to increase from 600 GB to 1,000 GB (about \$0.03/GB), and another \$34.63 to increase the cap from 1,000GB to unlimited.

Table 9: WTP Estimation also shows that failing to account for latency increases measured WTP for download bandwidth, but has no effect on measured WTP for upload bandwidth or data cap. Figure 2 illustrates the difference by comparing the download bandwidth WTP measures between the two surveys. This finding provides evidence that consumers may interpret download bandwidth as a proxy for the combined quality of download bandwidth and latency[25]. It also suggests that an omitted feature, at least in this setting, can lead to biased estimates, particularly when it is a component of a broader product attribute (speed). However, the bias appears to have been contained to only a subset of the other features comprising "speed" (download bandwidth).

With these figures in hand, we can compare our relative WTP estimates to the FCC's point tradeoffs for its CAF Phase II auction (see Tables 1 and 2). Tables 10 and 11 make these comparisons. In Table 10, we compare the point savings from improving performance tiers (based on download/upload bandwidth and data cap) in the FCC auctions to the increased WTP by the average household from such improvements. If the FCC's point system matched our WTP estimates, relatively speaking, then the percent change in points from tier to tier would equal the percent change in WTP from tier to tier. While they are reasonably similar, we cannot compare the "minimum" to "baseline" change. The FCC's point system implies that moving from

[25] We reject a joint test of equality across the two surveys for the WTP estimates for download bandwidth at the 1% level (one-tail p-value = 0.0003).

"baseline" to "above baseline" is worth about 50 percent more than moving from "minimum" to "baseline." Our WTP estimates suggest that moving from "baseline to "above baseline" is worth about 70 percent more than moving from "minimum" to "baseline." Similarly, the point system implies that moving from "above baseline" to "gigabit" is worth about half as much as moving from "baseline" to "above baseline." Our WTP estimates suggest that moving from "above baseline" to "gigabit" is worth about 64 percent as much as moving from "baseline" to "above baseline."

Table 10: Comparison of Auction and Survey Valuation of Performance Tiers

Performance Tier	Bandwidth	Usage Allowance	Points	Marginal Point Reduction in FCC Auction	Implied value of change relative to previous change	Marginal Increase in measured WTP[26]	Implied value of change relative to previous change
Minimum	≥10/1 Mbps	≥ 150 GB	65	.		.	
Baseline	≥25/3 Mbps	≥ 150 GB or U.S. median, whichever is higher	45	20		\$33.37	
Above Baseline	≥100/20 Mbps	2 TB	15	30	1.5	\$56.73	1.7
Gigabit	≥ 1 Gbps/500 Mbps	2 TB	0	15	0.5	\$36.32	0.6

In Table 11, we compare the relative gains from ascending performance tiers to the gains from improvement from high to low latency. We measure gains as points saved in the auction and as increased WTP by the average household and then compare score improvements. Here, we have a particularly striking finding: The relative gain in WTP from ascending a performance

[26] Baseline calculated as: (difference in WTP between 25 and 10 Mbps download) + (difference in WTP between 1 and 3 Mbps upload). Above baseline calculated as: (difference in WTP between 100 and 25 Mbps download) + (17/22)*(difference in WTP between 25 and 3 Mbps upload) + (difference in WTP between 1000 GB and 300 GB data cap). . . . Gigabit calculated as: (difference in WTP between 1000 and 100 Mbps download) + (difference in WTP between 100 and 25 Mbps upload) + (5/22)*(difference in WTP between 25 and 3 Mbps upload).

tier compared to improving latency is approximately **5 times** the relative gain in auction points when making the same comparison. This means the auction handicaps high-latency bidders much more than our WTP estimates would suggest is consistent with consumer value. Again, we note that we are likely underestimating this imbalance, as our calculations use the WTP for latency improvement from 300-600+ ms to less than 60 ms, which is likely more than the improvement implied by the auction.

Table 11: Comparison of Value for Improvements in Performance Tiers vs Latency

Performance Tier	Reduction in Score Based on Latency Improvement / Reduction in Score Based on Tier Improvement	Increase in WTP from latency improvement / increase in WTP from Bandwidth/Usage	FCC Latency Weight Relative to WTP
Minimum			
Baseline	1.25	0.26	4.8
Above Baseline	0.83	0.15	5.5
Gigabit	1.67	0.24	7.0

As different speed features can be particularly important for some specific users, we allow for heterogeneity in valuations across user demographics. To test for demographic differences, we create group dummy variables and interact them with all the Internet features, running separate analyses for each group. Specifically, we create dummy variables: young (under 45 years old), highly educated (bachelor degree and above), high income ($75,000 and above), and child (at least one child in household). Table 12 reports WTP differences by group; that is, the amount more (or less) that the average household in each group is willing to pay for a feature level, compared to its counterpart (e.g., the difference in WTP between a household with children and one without). We find that WTP differs little by age group, although younger respondents care less about premium upload speeds (> 100 Mbps) and do place somewhat higher value on increased data caps. In contrast, respondents with more education (bachelor and above) tend to generally care more about download bandwidth, and wealthier households ($75,000 and above) are willing to pay more for premium levels (above 100 Mbps) of download bandwidth. Households with child(ren) are willing to pay significantly more for unlimited data ($12.34 more than households without children) but they care less about the upload bandwidth. WTP for reduced latency is not meaningfully different across these demographic groups.

Table 12: WTP Differences by Groups

(using latency survey)

	Demographics				**Behaviors**			
Variables	**young**	**highly educated**	**high income**	**child**	**video**	**gaming**	**file**	**music**
Download								
10 Mbps	$2.45	$24.90***	-$5.65	-$15.98	-$0.24	$9.51	$15.80	-$4.14
25 Mbps	-$1.52	$11.70	$5.15	-$6.61	$17.27*	$24.18***	$20.44**	$16.18**
50 Mbps	-$2.27	$24.98***	$12.33*	-$18.77	$11.03	$21.69***	$33.15***	$9.87
75 Mbps	-$0.38	$22.08***	$12.27	-$8.62	$24.89***	$20.57***	$30.55***	$14.11*
100 Mbps	-$3.27	$17.73***	$16.59**	-$11.65	$27.43***	$25.63***	$30.99***	$18.82**
150 Mbps	$2.21	$16.71***	$16.04**	-$15.14*	$23.88***	$33.16***	$32.04***	$18.72***
300 Mbps	$4.19	$10.73	$10.99	-$6.95	$28.25***	$33.27***	$33.07***	$24.66***
500 Mbps	$4.51	$21.35***	$16.99**	-$13.22	$27.01***	$33.07***	$33.73***	$21.86***
1000 Mbps	-$0.45	$18.83***	$16.97**	-$5.58	$30.82***	$38.47***	$28.01***	$25.91***
Upload								
3 Mbps	$2.14	$1.01	$0.98	-$4.72	$6.91*	$0.31	$4.41	-$0.72
25 Mbps	-$1.10	-$3.07	$8.18***	-$13.85***	$6.37	$4.54	$9.73**	-$0.97
100 Mbps	-$8.01**	-$5.83	$5.25	-$10.76***	$3.88	-$2.49	$9.69**	-$5.345
Latency								
10-30 ms	-$1.17	$10.75***	$1.56	-$2.20	-$2.35	-$8.18**	-$1.48	-$4.68
30-60 ms	$2.61	$6.72*	-$1.70	-$5.98	-$5.52	-$12.74***	-$12.44**	-$2.51
60-150 ms	-$7.43*	$8.67**	-$2.92	-$6.38	-$15.41***	-$17.84***	-$9.76*	-$11.19***
150-300 ms	$0.16	$4.39	$0.07	$0.61	-$13.05***	-$16.97***	-$14.81***	-$10.39***
300-600+ ms	-$5.86	$0.39	-$7.35*	$2.81	-$16.88***	-$15.45***	-$18.17***	-$14.84***
Data cap								
600 GB	$2.98	$0.60	$6.14*	-$2.18	$8.53*	$7.70**	$12.98***	$7.32**
1000 GB	$9.48***	$6.73*	$7.47**	$7.70*	$19.97***	$9.86***	$25.49***	$11.87***
Unlimited	$9.71***	-$3.30	$1.72	$12.34***	$21.86***	$18.94***	$16.82***	$16.80***

Note: *10% significance, **5% significance, ***1% significance

We also test for differences in WTP based on user behaviors. We first identify four major bandwidth and/or latency demanding activities, and then group respondents based on their survey answers. To do this, we again create dummy variables – this time based on usage – and interact them with all the Internet features, again running separate analyses for each group. Specifically, we create dummy variables: video (use either stream video or video call or both), gaming (plays games online), file (transfer files online), and music (stream music online). Households that use video, play games, transfer files, or stream music tend to have higher WTP for more download bandwidth and higher data caps, and a greater aversion to high latency. For those who transfer files (83% of the households), upload bandwidth is relatively important. Moreover, gamers (40% of the households) show significantly more interest in low-level latency, which is consistent with what we observed from several gaming discussion forums.

7. Conclusions

In this paper, we measured households' valuations for home Internet speed in terms of bandwidth (download and upload) and latency, as well as valuations for data caps. Using carefully designed discrete choice surveys we find that households that highly value bandwidth enhancements at lower speeds, but valuation is highly concave, meaning that the incremental value of bandwidth decreases rapidly. Households also moderately value reduced latency, being willing to pay $8.66 per month to reduce latency from satellite to wired-terrestrial levels. We also find sensible patterns of differences in valuation for home Internet features across subgroups divided by demographics and usage. Lastly, our results across the two alternative surveys – one including latency and one excluding it – show that failing to account for latency increases measured WTP for download bandwidth, but has no effect on measured WTP for upload bandwidth or data cap.

Our findings mark the first relative valuation of bandwidth and latency. In doing so, they provide valuable information for broadband ISPs about the likely profitability of investments that generate high speeds, as well as which types of speed improvements may be most profitable. Our results across the latency and no-latency surveys suggests that omission of a feature can lead to at least some biased estimates, particularly when it is a component of a broader product attribute (speed).

Our findings also have important policy ramifications. In general, measuring some of the benefits of speed can help inform any policy intended to generate higher available speeds; these can be weighed against the costs. A specific application is the design of auctions meant to allocate broadband subsidies, with bids that are handicapped based on speed and data-cap provision. Our results suggest that FCC Phase II auctions have an implied value of latency relative to bandwidth and data caps that is at least five times what our valuation estimates indicate. Consequently, current U.S. policy may be over-penalizing latency relative to reductions in bandwidth and data caps.

References

Ahlfeldt, G, Koutroumpis, P., and T. Valletti. 2016. "Speed 2.0: Evaluating Access to Universal Digital Highways." CEPR Discussion Paper 11046

Akamai, 2015. "State of the Internet." 2015 Q4 Report.

Belshe, M. 2010. "More Bandwidth Doesn't Matter (much)." https://www.belshe.com/2010/05/24/more-bandwidth-doesnt-matter-much/.

Ben-Akiva, M., McFadden, D., and K. Train. 2016. "Foundation of Stated Preference Elicitation: Consumer Behavior and Choice-based Conjoint Analysis." Working paper.

Bode, K. 2015. "AT&T Takes Billions in FCC Subsidies for Broadband Expansion." DSL Reports, http://www.dslreports.com/shownews/ATT-Takes-Billions-in-FCC-Subsidies-For-Broadband-Expansion-134949.

Broida, R., 2012. "How Important is PC Processor Speed, Really?" CNET, http://www.cnet.com/news/how-important-is-pc-processor-speed-really/.

Carare, O., McGovern, C., Noriega, R. and J. Schwarz. 2015. "The willingness to pay for broadband of non-adoptersin the U.S.: Estimates from a multi-state survey." *Information Economics and Policy*, 30, 19-35

Caussade, S., de Dios Ortúzar, J., Rizzi, L. I., and D. Hensher, 2005. "Assessing the Influence of Design Dimensions on Stated Choice Experiment Estimates." Transportation Research Part B, 39(7):621-640.

Czajkowski, M., Giergiczny, M., and E. Zawojska. 2015. "Does the Number of Discrete Choice Alternatives Matter for Respondents' Stated Preferences? The Case of Tap Water Quality Improvement." Working Paper.

Ding, M., Grewal, R., and J. Liechty. 2005. "Incentive-Aligned Conjoint Analysis," *Journal of Marketing Research*, 42, 67–82

Dutz, M., Orszag, J.M., and R. Willig. 2009. "The Substantial Consumer Benefits of Broadband Connectivity for US Households." Mimeo

FCC, 2015, "2015 Measuring Broadband America Fixed Report," published Dec. 30, 2015

FCC, 2016a, "2016 Broadband Progress Report," published Jan. 28, 2016

FCC, 2016b, Broadband Speed Guide, https://www.fcc.gov/reports-research/guides/broadband-speed-guide, accessed 5/11/2016.

FCC, 2016c, Connect America Fund website, https://www.fcc.gov/general/connect-america-fund-caf, accessed March 25, 2016.

Forman, C., Goldfarb, A., and S. Greenstein. 2012. "The Internet and Local Wages: A Puzzle." *American Economic Review*, 102(1), 556-575.

Fung, B. 2015. "This City Just Rolled Out America's Fastest Internet Service Yet." The Washington Post, https://www.washingtonpost.com/news/the-switch/wp/2015/09/09/this-city-just-rolled-out-americas-fastest-internet-service-yet/.

Goolsbee, A., and Klenow, P. J. 2006." Valuing Consumer Products by the Time Spent Using Them: An Application to the Internet." *American Economic Review*, 96(2), 108-113.

Greene, W. H. 2012. "Econometric Analysis 7th ed." Upper Saddle River, NJ: Prentice Hall.

Greenstein, S. and McDevitt, R. C. 2011. "The broadband bonus: Estimating broadband Internet's economic value." *Telecommunications Policy*, 35(7), 617-632.

Grigorik, I. 2013a. "Chapter 1. Primer on Latency and Bandwidth." *High Performance Browser Networking*, O'Reilly Media, Ebook ISBN 978-1-4493-4471-9

Grigorik, I. 2013b. "Chapter 10. Primer on Web Performance." *High Performance Browser Networking*, O'Reilly Media, Ebook ISBN 978-1-4493-4471-9

McFadden, D. 1974. "Conditional Logit Analysis of Qualitative Choice Behavior." In P. Zarembka, ed., *Frontiers in Econometrics*. New York: Academic Press.

Miller, K.M., Hofstetter, R., Krohmer, H., and Z.J.Zhang. 2011. "How Should Consumers' Willingness to Pay Be Measured? An Empirical Comparison of State-of-the-Art Approaches." *Journal of Marketing Research*, 48, 172-184

Netflix Internet Connection Speed Recommendations, https://help.netflix.com/en/node/306, accessed 5/11/2016.

Nevo, A., Turner, J., and J. Williams. 2016. "Usage-Based Pricing and Demand for Residential Broadband." *Econometrica*, 84(2), 411-443.

Rappoport, P., Kridel, D., Taylor, L., Duffy-Deno, K., and J. Alleman. 2002. "Residential Demand for Access to the Internet." In Madden, G. (ed.) *The International Handbook of Telecommunications Economics: Volume II*, Edward Elgar Publishers, Cheltenham.

Rosston, G., Savage, S., and D. Waldman. 2010. "Household Demand for Broadband Internet in 2010." *The B.E. Journal of Economic Analysis and Policy*, 10(1), Article 79.

Savage, S. and D. Waldman. 2005. "Broadband Internet Access, Awareness, and Use: Analysis of United States Household Data." *Telecommunications Policy*, 29, 615-633.

Savage, S., and D. Waldman. 2008. “Learning and Fatigue During Choice Experiment: A Comparison of Online and Mail Survey Modes.” Journal of Applied Econometrics, 23, 351-371.

Savage, S., and D. Waldman. 2009. “Ability, Location and Household Demand for Internet Bandwidth.” *International Journal of Industrial Organization*, 27, 166-174.

Telecompetitor.com, 2016, http://www.telecompetitor.com/igr-average-monthly-broadband-usage-is-190-gigabytes-monthly-per-household/, accessed August, 2017.

USTelecom, 2017, https://www.ustelecom.org/broadband-industry/broadband-industry-stats/investment, accessed August, 2017.

Varian, H., 2002. “The Demand for Bandwidth: Evidence from the INDEX Project.” Mimeo

Zwerina, K., Huber, J., and W.F. Kuhfeld. 2010. “A General Method for Constructing Efficient Choice Designs,” *Marketing Research Method in SAS*, SAS Institute Inc., Cary, NC, USA

Appendix: Survey Instrument

Read Me First: Description of Internet Features

In this section, we describe four main features of a home Internet plan: **price**, **bandwidth**, **latency**, and **data cap**. We also explain how the last three features may affect your experience of using the Internet. You may find this information helpful when answering questions in section II. You can always refer to this "Read Me" section while answering the survey.

- **Price**: The total monthly fee you pay for home Internet service, before tax.

- To understand **bandwidth** and **latency**, imagine that the Internet is a highway connecting two cities. In this analogy,
    - **Bandwidth** is the number of lanes in that highway.
    - **Latency** is the time it takes to travel from one city to the other via that highway, and therefore depends on highway conditions such as congestion, traffic signs, speed limit, and distance.

  Sometimes a highway needs to allow large amounts of traffic to move smoothly, such as during rush hour or to and from a football game. Similarly, sometimes an Internet connection needs to move large files, like videos, between two points. Other times the highway must be capable of transporting smaller loads, but more quickly, such as an ambulance rushing to pick someone up and deliver him back to the hospital. Similarly, some Internet games require nearly-instant reactions to a button-push. In this case, the effective speed limit and route distance may become the bottleneck. (Some example tasks are listed in *table 1* and *table 2*)

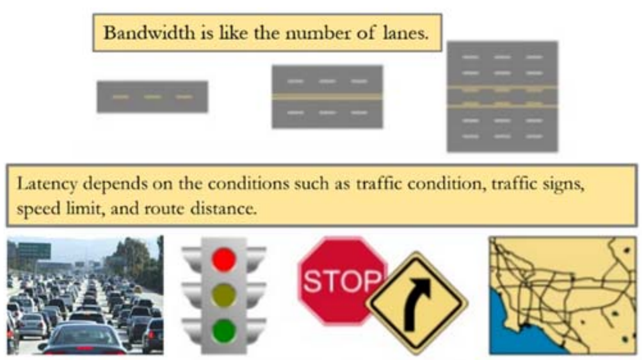

Source: www.slideshare.net

- **Bandwidth** is the maximum throughput in a given amount of time (usually expressed as megabits per second, Mbps) provided by your home Internet connection.

    - **Download bandwidth** is the maximum Mbps your home Internet service can provide to your devices. All the download bandwidth levels included in the survey can smoothly perform basic tasks, such as **e-mail**, **web-browsing or searching**, **online shopping**, **streaming music**, making **voice calls**, and **basic online gaming** over multiple devices.

  - **Upload bandwidth** is the maximum Mbps your home Internet service can move from your devices to the Internet. The upload bandwidth is relevant for uploading large files, pictures, or videos to file-sharing site (e.g. Dropbox), social media (e.g. YouTube), or another user (e.g. via Skype).

Latency is the round-trip time (typically expressed in millisecond, ms), for a data packet to travel from a consumer's home to the closest server and back. Shorter latency improves Internet responsiveness on web page loading and some interactive tasks, including voice/video and online multiplayer games.

Data Cap is the amount of data (typically expressed in gigabytes, GBs) per month the home Internet provider allows for the given monthly fee. Providers typically charge additional fees for data use above the cap.

In the next two pages, we translate some common technical standards to user experiences. We will present this information to you again within the survey.

Table 1. Some Bandwidth-demanded Tasks

Average Bandwidth Levels	Can Stream Video / Video over IP on a single device (Standard Definition, SD)	Can Stream Video / Video over IP on a single device (High-Definition, HD 1080P)	Quality possible for Streaming Video / Video over IP on 2-4 devices concurrently (Ultra HD 2160P)	Time to download / upload 5GB file (e.g., 2 hour HD movie)
4 Mbps	Yes			167 mins
10 Mbps	Yes	Yes	SD	67 mins
25 Mbps	Yes	Yes	HD	27 mins
50 Mbps	Yes	Yes	HD	13 mins
75 Mbps	Yes	Yes	HD	9 mins
100 Mbps	Yes	Yes	Ultra HD	7 mins
150 Mbps	Yes	Yes	Ultra HD	4 mins
300 Mbps	Yes	Yes	Ultra HD	2 mins

500 Mbps	Yes	Yes	Ultra HD	1 mins
1000 Mbps	Yes	Yes	Ultra HD	0.7 mins

Reference sources: Comcast, Netflix, Verizon

Notes for upload bandwidth:

1. Upload capability is irrelevant for streaming video because no content is provided from the user end.
2. In general, though not always, the estimated upload bandwidth of home Internet plans is either:

- Identical to download speed (symmetric). This is the case with fiber technology.
- Around 1/8 of your download speed. This is the case for non-fiber technologies, such as cable. For example, if your download bandwidth is 50 Mbps, the corresponding upload bandwidth is probably around 6 Mbps.

Table 2. Some Latency-sensitive Tasks

Average Latency Levels	Online Gaming	Voice over IP	Web-page loading time for an average website (Belshe 2010)
Less than 10 ms	Unnoticeable delay	Unnoticeable delay	0.8 secs
10 - 30 ms	Unnoticeable delay	Unnoticeable delay	1.1 secs
30 - 60 ms	Minor delay	Unnoticeable delay	1.5 secs
60 - 150 ms	Some delay	Some delay	2.7 secs
150 - 300 ms	Significant delay	Significant delay	4.8 secs
300 - 600 ms	More Significant delay	More Significant delay	8.9 secs

Reference sources: Belshe 2010, pingtest.net

Table 3. Hours of Video Streaming Possible Under Different Data Cap Levels

Monthly Data Cap	Stream video for (SD)	Stream video for (HD)
300 GB	300 hrs	100 hrs
600 GB	600 hrs	200 hrs
1000 GB	1000 hrs	333 hrs

Reference source: Netflix

<u>Sample question from survey including latency:</u>

Choice Question 1

1.1. Please choose the best Internet plan as if you are deciding which to purchase for your home, assuming that all features other than the five listed here are identical (e.g. brand, customer service, contract length, etc).

	Internet A	Internet B	Internet C	Internet D
Price	$ 45	$ 70	$ 80	$ 35
Download Bandwidth	25 Mbps	75 Mbps	1000 Mbps	4 Mbps
Upload Bandwidth	3 Mbps	3 Mbps	25 Mbps	1 Mbps
Latency	60 to 150 ms	300 to 600 ms	30 to 60 ms	Less than 10 ms
Data Cap	600 GB	300 GB	1000 GB	Unlimited
<u>Pick One</u>	○	○	○	○

Note: the over usage rate for data is $10 for an additional 50 GB

1.2. Would you switch to your choice in 1.1 above from the actual Internet plan you currently have assuming that all features other than the five listed here are identical (e.g. brand, customer service, contract length, etc)?

	Your Chosen Plan	Your Current Plan
Price		Reported / inferred price
Download Bandwidth		Reported
Upload Bandwidth		Reported / inferred
Latency		Reported / predicted
Data Cap		Reported / unlimted
<u>Pick One</u>	○	○

Table: Latency Estimation

Brand \ Technology	DSL	Cable	Fiber	Satellite	Don't Know
AT&T	30 - 60 ms		10 - 30 ms		10 - 30 ms
Cablevision		10 - 30 ms			10 - 30 ms
CenturyLink	30 - 60 ms				30 - 60 ms
Charter		10 - 30 ms			10 - 30 ms
Comcast		10 - 30 ms			10 - 30 ms
Cox		10 - 30 ms			10 - 30 ms
Frontier	30 - 60 ms		10 - 30 ms		30 - 60 ms
Google Fiber			10 - 30 ms		10 - 30 ms
Hughes				300 - 600+ ms	300 - 600+ ms
Mediacom		10 - 30 ms			10 - 30 ms
TWC		10 - 30 ms			10 - 30 ms
Verizon	30 - 60 ms		10 - 30 ms		10 - 30 ms
ViaSat/Exede				300 - 600+ ms	300 - 600+ ms
Windstream	30 - 60 ms				30 - 60 ms
Don't know or others	30 - 60 ms	10 - 30 ms	10 - 30 ms	300 - 600+ ms	60 - 150 ms