NET Institute*

www.NETinst.org

Working Paper #16-07

September 2016

**Broadband Mergers and Dynamic Bargaining:
An Application to Netflix**

Daniel Goetz
Princeton University

* The Networks, Electronic Commerce, and Telecommunications (“NET”) Institute, <http://www.NETinst.org>, is a non-profit institution devoted to research on network industries, electronic commerce, telecommunications, the Internet, “virtual networks” comprised of computers that share the same technical standard or operating system, and on network issues in general.

Broadband Mergers and Dynamic Bargaining: An Application to Netflix*

Daniel Goetz[†]

September 30, 2016

Abstract

I measure how mergers in the market for broadband internet service affect short-run welfare. Mergers between internet service providers (ISPs) with non-overlapping markets may decrease welfare by increasing ISP bargaining leverage against content providers. However, study of this welfare channel has been stymied by a lack of data on interconnection fees between content and internet service providers. I estimate an industry model of demand, plan choice, pricing and interconnection bargaining using data on plan prices, consumer choice sets and bargaining delays between major U.S ISPs and the leading purveyor of streaming video content, Netflix. Intuitively, if delaying agreement over interconnection degrades quality of service to subscribers, then the opportunity cost of lost subscriptions identifies the fee. To map disagreement times and ISP competition into interconnection fees, I develop a multilateral dynamic bargaining model with asymmetric information. ISPs make take-it-or-leave-it offers to learn about Netflix's benefit from interconnection, while simultaneously competing for subscribers who value Netflix quality of service. I structurally estimate the model and recover fixed interconnection fees ranging from 44 to 69 million USD. I find that a proposed merger between TimeWarner and Comcast that was challenged by the Federal Communications Commission would have slightly raised interconnection fees and bargaining length, reducing consumer welfare by 1.9 percent.

---

*Financial support from the NET Institute (www.netinst.org) is gratefully acknowledged

[†]Fisher Hall, Princeton University. Correspondence to dtgoetz@princeton.edu

1. Introduction

Most consumers in the U.S. can access high speed internet through only one or two internet service providers (ISPs). This bottleneck also affects content creators, who, to ensure their content reaches consumers, must deal with large ISPs that exert significant market power. Regulators face frequent decisions on competition policy, from ISP mergers to the introduction of municipal fiber options, all of which can drastically affect the terms at which content creators must bargain to reach consumers—as well as the options that consumers have to connect to the internet. This paper analyzes the effect of internet competition policy on consumer welfare, and on the division of surplus between U.S. ISPs and an important content creator, Netflix.

I conduct my analysis in the context of a high-profile bargaining event from 2013-2014. While negotiating with ISPs over a new interconnection agreement, Netflix quality of service (QoS) was badly degraded to a number of U.S. ISPs, with QoS returning piecemeal after agreement was concluded at different times with different ISPs. Examining this particular episode of interconnection negotiation is important to welfare and relevant to future policy. By the increasing amount of time the average American spends online watching streaming video, it is evidently high value; moreover, as quickly growing content, streaming video services must frequently renegotiate interconnection terms with ISPs.

Since no mergers or substantial regulatory interventions were undertaken during the negotiation, I specify and estimate a structural model to provide a framework for policy experiments. I model consumer demand for firms and plan choice, firm pricing, and bargaining over the terms of interconnection. I first combine data on market shares and ISP-specific Netflix QoS degradation to estimate the distribution of household preferences for prices, download speeds, and quality Netflix service. I next estimate fixed, ISP-specific payments that Netflix must pay to restore quality of service to that ISP's consumers—the aforementioned interconnection agreements. The key insight in the paper is that the delay in agreement timings encodes information about the eventual payment ISPs receive from Netflix. I use the demand estimates and agreement timings to estimate the parameters of a multilateral dynamic bargaining model of the interconnection market. Holding the demand and bargaining parameters fixed, I can then simulate environments where the set of negotiating ISPs is different, due either to merger or entry.

My main counterfactual is to allow the two largest cable internet providers, Comcast and TimeWarner, to have merged before the bargaining event begins. This counterfactual is of regulatory interest, as the two firms attempted to merge in early 2015 but were blocked by the Federal Communications Commission on the grounds that it would have created a firm with enormous bargaining power in the interconnections market. I find that before the merger, *ex ante* expected delays and interconnection fees among the four largest firms were 3.03 quarters and \$120m USD

on average. After the merger, I estimate that the average delay and fee increase by 3.7 percent among the largest providers. The increased delay reduces *ex ante* expected consumer welfare by roughly 4 percent.

The model has three types of agents: consumers, internet service providers, and the upstream content creator—Netflix. I estimate consumer preferences for internet access and for Netflix using individual level data on plan choices conditional on ISP choice, combined with nationwide ISP internet subscription shares and Netflix-ISP pairwise QoS degradation. Individuals first choose among the set of horizontally differentiated ISPs available to them, then choose one of multiple vertically differentiated plans available from that ISP. While ISPs charge higher monthly subscription fees for plans with more features, this pricing and feature schedule is mostly constant nationwide and over time, making it difficult to estimate elasticities. I use fine geographic variation in the set of available ISPs and plans to identify the distribution of preferences for price, download speed, access technology, and Netflix quality of service in a mixed-logit framework. I find that during periods of disagreement, ISP profits decrease by about 0.8 percent on average, and that the price elasticity is inelastic. I therefore make the assumption that marginal costs are zero to recover profits from the demand curve.

I combine the model of firm profits generated by the demand curve with a model of dynamic bargaining, where lost profits due to disagreement identify the eventual interconnection fee. I develop what is, to my knowledge, the first estimable dynamic multilateral bargaining game in the structural IO literature. The core of the model features one sided asymmetric information to generate delays: at the beginning of the game, Netflix draws ISP-specific surpluses from interconnection which it will realize on agreement with that ISP. Conditional on observable covariates—e.g., the ISP’s network footprint and technology—these draws are independent, and observed only by Netflix. Each period, every ISP that has not yet reached agreement with Netflix makes offers to screen Netflix’s ISP-specific benefit from agreement. Consumer substitution between ISPs leads to profit interactions, which costs the ISP subscriptions due to degraded quality of service if the ISP adopts too fine a screening strategy.

My structural assumptions allow me to numerically solve for optimal state-contingent firm strategies, from which I can recover the probability of observing the joint agreement timings that appear in the data. This forms the basis of a maximum likelihood estimation strategy to recover the parameters of the distribution of ISP specific cost-savings Netflix stands to gain from agreement. I argue that since I only observe disagreement duration, and not actual fees, including profit interactions is crucial for identification. Intuitively, two firms that face distributions with the same variance but different means will adopt screening strategies with the exact same equilibrium distribution of agreement timings, since only relative payoffs matter. With flow profits included, if an ISP that faces a large dip in profits if agreement is not reached still adopts a fine screening strategy, then the mean of its distribution of payoffs must be high.

I evaluate the heterogeneous impact of the merger using the recovered joint distribution of agreement timings and offers. The average reduction in bargaining time hides significant heterogeneity: the merged firm actually concludes bargaining roughly 11 percent quicker in expectation than Comcast or TimeWarner, while AT&T's bargaining duration lasts 14 percent longer. These results come from the fact that while the merged firm has more to gain from screening than Comcast or TimeWarner, they also have more profit to lose from disagreement. AT&T evidently has less to lose, and is therefore able to stretch out its screening for longer. I find that the merged firm recovers a fee 9.0 percent higher than either Comcast or TimeWarner, while AT&T's finer screening allows it to enjoy a 4.1 percent higher fee. Finally, The model also predicts exactly what offers would have been made to ensure observing the agreement times we did; these range from 44 million USD for AT&T to 69 million USD for Comcast.

As a fairly complicated dynamic structural IO model with a specific application, this paper sits at the intersection of several literatures. First, a recent literature including Grennan (2013) and Gowrinsankaran, Nevo and Town (2015) has fruitfully built up the Nash-in-Nash framework as the workhorse model of multilateral bargaining. As it is static and efficient, this model cannot accommodate dynamic elasticities of substitution that arise from switching costs, or rationalize bargaining delays. Crawford and Yurukoglu (2012) and Dafny, Ho and Lee (2016) also use the Nash-in-Nash framework to estimate a multilateral bargaining game, but they add the feature that consumers can substitute between providers. Substitution by downstream consumers is a key source of incentives in upstream bargaining, and I extend their insights to a dynamic setting.

There is an extremely sparse literature on estimating dynamic games with incomplete information to which this paper contributes. Most recently, Fershtman and Pakes (2012) provide a framework for estimation in the form of an "experience based equilibrium" or EBE. Unlike my paper which solves for rational equilibrium by making strong structural assumptions, Fershtman and Pakes develop a new equilibrium concept that has firms paying only limited attention to states. Ambrus, Chaney and Salitsky (2016) develop and estimate the bilateral version of the one-sided screening framework. While one could fit the model of Ambrus et al. to the durations data, it would ignore the real interactions between ISPs as they compete for customers, and would remove the bite of the switching cost reduction counterfactual.

The supply side also bears a resemblance to a recent literature on multilateral optimal stopping problems. Honoré and de Paula (2010) estimate structural parameters for two individuals whose optimal stopping times are interrelated in a durations model, while Björkegren (2015) estimates joint durations for arbitrary individuals on a network timing their stopping decisions. This paper utilizes insights from lattice theory, similarly to Björkegren as well as Jia (2008) to estimate the supply side of the model.

Finally, the paper contributes to a literature on network neutrality that has, until now, been mostly theoretical in its methods. Roughly speaking, network neutrality is the policy that ISPs

should not be allowed to explicitly discriminate against content by selectively throttling transmission, and this principle has recently been passed by the U.S. Federal Communications Commission and upheld in U.S. courts. Lee and Wu (2009), Economides and Hermalin (2012) Gans (2015), and Peitz and Schuett (2016) have all analyzed whether a policy of network neutrality is welfare improving, and what types of distributional impacts different neutrality policies will have. I add to this discussion the idea that competition policy can also have an impact on short-run neutrality; in particular, on the type of incidental violations of neutrality due to bargaining that regulations against explicit discrimination might not be suitable for preventing.

The rest of the paper is organized as follows: Section 2. motivates the problem and empirically documents the bargaining delays, while Section 3. presents the model. Section 4. describes the data, and Section 5. provides some reduced form evidence on the validity of the framework. Section 6. outlines the estimation procedure and describes the obstacles to identification, and Section 7. gives the results and counterfactuals.

2. Background

2.1. Structure of the internet

The internet is best understood as a two-sided market. On one side are consumers, who purchase access to the internet in order to enjoy online services and content such as email and streaming video. On the other side are the providers of content, such as Netflix and Google, who charge consumers either directly, via subscription fees, or indirectly, via advertisements, for accessing the content. In the middle are two types of firms that intermediate the relationship between consumers and content providers.

Consumers interact with "last-mile" or "edge" internet service providers like Comcast and Verizon. These providers typically offer access to the internet for consumers and small businesses/content creators for a monthly fee. Access is tiered: higher bandwidth, which allows for a faster connection and more simultaneous connections by members of a household, can be had for a higher fee. These firms are called "last-mile" as they own, and have exclusive right to sell service on, the physical infrastructure (such as Cable or Fiber Optic lines) that brings the internet from large interconnection points to individuals' homes.

At interconnection points, last-mile ISPs exchange data with "transit" ISPs such as Level3 and Cogent. Large content providers like Netflix must buy access from these transit providers. In addition, they can pay to upload content to so-called "content delivery networks" (CDNs)—caches of servers distributed around the country that ensure no consumer is far from a content source—that also buy access from transit ISPs on the content providers' behalf. Like consumers, content providers often pay a fixed access fee for a certain guaranteed bandwidth from transit

providers. For the content providers, buying internet access through transit ISPs allows their content to reach last-mile ISPs, and the last-mile ISPs' subscriber base.

For most of the 1990s and 2000s, transit ISPs and last-mile ISPs sent data back and forth to each others' networks at no charge—an arrangement known as "peering."¹ Since the internet was designed to facilitate flows of information, not lock them behind payments, and since a roughly equivalent magnitude of data was flowing in both directions, this arrangement could easily be rationalized.² It also incentivized the largest content providers like Google and Yahoo to build out their own transit networks, incurring a one time fixed cost to reap lower variable costs of transmission.

Starting in the late 2000s streaming video began to account for a large fraction of transmitted data. Transit ISPs serving online video distributors (OVDs) could no longer reasonably expect peering treatment from last-mile ISPs, as consumers watched far more video than they uploaded, breaking the symmetry that made peering viable. As the largest OVD, Netflix not only paid the transit ISPs for connection, but also pursued a policy of paying the fees that last-mile ISPs charged transit ISPs.³

In mid-2012, Netflix decided to roll out their own CDN called Open Connect. Doing so involved incurring a large fixed development and deployment cost⁴ but would save them money in two ways. First, it would allow them to save on the variable cost of using third party CDNs. Second, by locating the servers inside last-mile ISPs' own networks, Netflix would no longer need to ensure that it paid for sufficient bandwidth from transit ISPs to accommodate demand at peak times. With Open Connect servers located in, for instance, Comcast's network, Netflix could update the servers slowly and during off-peak times when Comcast consumers were not streaming, and therefore save on transit costs.⁵ Open Connect would allow Netflix to deliver service reliably and at lower cost.

Although Netflix's standing offer to install Open Connect servers in the largest U.S. ISPs was rejected by those ISPs, transmission of Netflix content to the end users remained reliable until mid-2013. At that point, Netflix's transit providers, the last-mile ISPs, or both the transit and last-mile ISPs halted the routine upgrades and maintenance that had ensured reliable transmission, and quality of service plummeted for the largest ISPs.⁶ This sharp drop in throughput is

---

1

²For instance, since the internet was intended to facilitate information flows, costs to develop systems that could establish "tolls" on data may have been prohibitive. As a two sided market, in the early days of the internet it may also have made more sense to subsidize content to make higher margins on consumer access fees.

³Paragraph 12, Statement of Ken Florance, Vice President of Content Delivery at Netflix since 2012.

⁴Netflix Petition to Deny, pg. 49, paragraph 1. Fixed investment in R&D and deployment on the order of \$100 000 000.

⁵Netflix petition to deny, pg. 49, paragraph 2. "Open Connect...uses a 'proactive caching' method to conduct daily content updates during periods when networks are least used, such as early in the morning, to avoid congesting the network."

⁶Harvard Business School case N9-616-007.

illustrated in Figure 2 and Figure 3: Cogent, Netflix's transit provider, sees a sudden and massive drop in its average bitrate to last-mile ISPs Comcast and Verizon but not to the smaller last-mile ISP Cablevision. Internap, a transit provider which did not carry Netflix traffic, saw no reduction in throughput with any last-mile ISP.

As Netflix's throughput data in Figure 1 indicates, the slowdown did not resolve for many months. Moreover, from Netflix's own quarterly reports it is clear that agreement with different ISPs was reached at different times.⁷ The resolution of the slowdown is also explicitly tied to the payment by Netflix of collocation fees for Open Connect servers in official Netflix testimony brought before the Federal Communications Commission.⁸ The evidence makes it clear that this degradation in the quality of service of Netflix content was business related, and that the event can be usefully modeled as a prolonged, dynamic negotiation between Netflix and the largest U.S. ISPs.⁹

2.2. Network neutrality

Early architects of the internet promoted a design such that all content could be equally accessible by all consumers, regardless of where the content originates or which last-mile ISP the consumer subscribes to. This equal treatment of content is the principle of "network neutrality", which was finally codified by the FCC in 2015. The principle does not rule out paywalls erected by content providers themselves, but does prohibit last-mile ISPs from acting as gatekeepers for internet content.

In the Netflix event, *ex post* testimony makes it difficult to determine whether ISPs were acting as gatekeepers and throttling Netflix content in exchange for payment. Instead, there were numerous opportunities for both last-mile ISPs and transit ISPs (acting on Netflix's behalf) to upgrade transmission capacity, but both parties failed to do so.¹⁰ The inability to assign responsibility to last-mile ISPs even after the fact raises two questions: first, whether regulations that prohibit "gatekeeper" behaviour can be reasonably enforced, either at all or in a timely enough fashion to protect consumers; second, whether prohibiting "gatekeeper" behaviour is the key to avoiding quality of service degradation to consumers.

Rather than focus on the legal framework for interaction between last-mile ISPs, transit ISPs and content, I take the approach that renegotiation of the terms of interconnection is inevitable, and leads to unpreventable QoS degradation. I focus instead on policies that are designed to directly affect consumers—some of which, like the introduction of a publicly-owned fiber optic

---

⁷April 2014 Netflix Q-10.

⁸Netflix petition to deny, pg. 49, paragraph 2.

⁹Petition to deny, pg. 52, paragraph 2. Comcast, Verizon, Time Warner and AT&T "presumably made the business decision that the present discounted value of benefits from degrading the quality of the Netflix video stream to [their] subscribers was greater than the present discounted value of the costs."

¹⁰From consultations with an economist who was with the FCC at the time the slowdown was occurring.

alternative, have already begun to be rolled out. By integrating upstream negotiations and consumer demand in a single framework, I will be able to evaluate the effectiveness of "consumer-facing" pro-competitive policies on the length of negotiations.

3. Model

There are two components to the model. On the demand side, consumers choose internet plans (price and download speed combinations) in a two step procedure. They first select a provider from a market-specific choice set of ISPs, then choose a plan from that ISP's menu. Consumers value price and the download speed according to their own heterogeneous observed and unobserved characteristic. I do not model consumers' choice of whether to purchase Netflix or not, but allow disutility for reductions in service quality to appear in the preference for download speed.

On the supply side, ISPs make a take-it-or-leave-it offer every period to Netflix over the fee Netflix will pay for collocation. ISPs then set prices after observing the vector of agreement. The evolution of plan characteristics, coverage and mergers are taken as given. A component of the marginal benefit to Netflix of agreement with each ISP is known by Netflix, but unobserved by the ISP; the ISP therefore faces a trade off between offering a low fee that is accepted quickly with high probability, or demanding higher fees that may not be accepted, pushing agreement into the discounted future. In equilibrium the ISP will make offers with a positive probability of rejection, implying that there is a delay in reaching agreement. Moreover, whether an ISP's offers are accepted or rejected will affect the flow profits of competing ISPs because consumers value quality of service and are willing to substitute between ISPs. This substitution leads to strategic interactions that imply bunching of agreement timings in equilibrium.

3.1. Demand

3.1.1. Flow utility

Every period, consumers choose among all the plans offered by the available ISPs competing in their local market. Letting f index firms (ISPs) and m markets, consumer i in market m at time t first chooses among $f \in \mathcal{F}_{mt}$. Each firm offers a menu of multiple vertically differentiated plans in market m , $j \in \mathcal{J}_{f,mt}$, which may vary by market. Conditional on choosing a firm, a consumer chooses among the available j offered by f in m at time t .

A firm's plans are differentiated at the national level on two dimensions: the monthly fee p_{jft} , and the download speed q_{jft} .¹¹ Consumers also care about whether the ISP has sufficient

---

¹¹In the 2013 CPS Internet Use Supplement, over 80% of respondents indicate speed, price, or reliability as the most important feature of plan service, with the upload speed and data cap accounting for less than 15% of the remainder.

throughput with Netflix to serve streaming content. a_{ft} is a dummy indicating whether the ISP f and Netflix have agreed or not, which I previously argued maps one to one with whether there is sufficient throughput for streaming service. Each ISP f also has an unobserved dimension of heterogeneity, ξ_{ft} . As in similar papers on demand for telecommunications services, this heterogeneity may reflect quality of customer service and availability of bundled services (e.g. telephone, cable).

Consumers have heterogeneous preferences for ISP plan characteristics. In particular, depending on observed demographics and unobserved characteristics, a household may value differently the prices, download speed, and technology of their provider. I assume that switching between vertically differentiated products within a provider, as well as switching between providers or to the outside option, are all costless to the consumer.

Let f_{imt} denote individual i 's choice of firm f in market m at time t . Then flow utility is given as

$$U_i(f_{imt}) = \begin{cases} \alpha_{i,g(f)} + \gamma(1 - a_{ft}) + \max_{j \in \mathcal{J}_{f_{mt}}} \{\alpha_{ip}p_{jft} + \alpha_{iq}q_{jft}\} + \xi_{ft} + \epsilon_{ift} & \text{if } f_{imt} \in \mathcal{F}_{mt} \\ \epsilon_{i0t} & \text{else} \end{cases}$$

The first term is individual i 's technology-specific fixed effect, where $g(f)$ denotes the technology type of firm f . The available technologies are DSL, Cable, Fiber Optic, and Satellite. The second term is the component of f -specific welfare that depends on whether or not agreement has been reached between ISP f and Netflix. Intuitively I expect $\gamma < 0$, and show in reduced form evidence that market shares do decrease as a result of disagreement. The third term is i 's optimal flow utility from the plan specific price and download speed characteristics.

I model consumer heterogeneity $(\alpha_{ig}, \alpha_{ip}, \alpha_{iq})$ as being comprised of observed and unobserved characteristics. In particular,

$$\begin{aligned} \alpha_{i,g(f)} &= \mathbf{x}'_i \boldsymbol{\alpha}_g^o + \nu_{ig} \\ \alpha_{ip} &= -\exp\left(-(\mathbf{x}'_i \boldsymbol{\alpha}_p^o + \nu_{ip})\right) \\ \alpha_{iq} &= \mathbf{x}'_i \boldsymbol{\alpha}_q^o + \nu_{iq} \end{aligned} \tag{1}$$

Where $\mathbf{x}_i \equiv (x_{i1}, \dots, x_{iR})$ denote consumer characteristics and $\{\nu_{ig}, \nu_{ip}, \nu_{iq}\}$ denote unobserved (by the econometrician) consumer tastes. As usual, ϵ_{ift} is assumed to be distributed according to the type I extreme value distribution. I borrow the functional form for the heterogeneous price coefficient from Berry, Levinsohn and Pakes (2004).

3.1.2. Market shares

As usual, the conditional choice probabilities have an analytic form due to the assumption of the type I extreme value error:

$$\Pr(f_{imt} = f) = \frac{\exp(\delta_{ft} + \kappa_{ifmt})}{1 + \sum_{f' \in \mathcal{F}_{mt}} \exp(\delta_{f't} + \kappa_{if'mt})}$$

Where δ_{ft} incorporates the mean utility from choosing f in market m at time t , and κ_{ifmt} incorporates individual and market specific deviations from the mean.

I then aggregate in both the i and m dimensions. For the i dimension, I must integrate with respect to the conditional distribution of consumer heterogeneity ω_i with a given market m . To aggregate across markets, I use a time-varying weight that corresponds to the share of census block groups (10 digit FIPS code) for which that particular combination of firms are the only residential providers.

$$s_{ft} = \sum_{mt} w_m \int s_{fmt}(\omega_i) dG_{\omega, s_{t-1}}(\omega_i | \theta)$$

The above shares are only across ISPs; I can also generate the distribution of consumer plan purchases within each firm for given parameter values (and δ). In this model, consumer heterogeneity in $(\alpha_{ip}, \alpha_{iq}, \eta_i)$ is the only reason why within-firm shares are non-degenerate. Thus, matching conditional purchase moments will be informative about these consumer heterogeneity parameters. The share of individuals purchasing plan j from firm f at time t is given by:

$$s_{j|ft} = \sum_{m|f \in m} \tilde{w}_m \frac{\int \mathbf{1} \left[j = \underset{j'}{\operatorname{argmax}} \{ \alpha_{ip} p_{jft} + \alpha_{iq} q_{jft} \} \right] s_{ifmt} dG(\omega_i, s_{ifmt-1})}{\int s_{ifmt} dG(\omega_i, s_{ifmt-1})}$$

Where $\tilde{w}_m = w_m / \left(\sum_{m'|f \in m'} w_{m'} \right)$

3.2. Supply

I write the supply model as a game played between downstream internet service providers like Comcast and Verizon, and the upstream content provider, Netflix. In most of the empirical bargaining literature the bargaining problem is a static game, where an upstream-downstream pair chooses the marginal input price to maximize joint surplus—conditional on the marginal input prices of all other pairs. However, the primary purpose of this supply model is to explain the pattern of bargaining delays in the data, necessitating a dynamic approach.

Simply introducing dynamics to the static model is not enough to induce bargaining delay:

starting with Rubinstein (1982), a long line of micro theory papers on bilateral bargaining show that complete information and a positive discount rate is usually enough to guarantee instantaneous agreement.¹² To generate delays I assume one-sided incomplete information: Netflix has information on its fixed operating cost savings from direct interconnection that the ISPs do not. ISPs make offers of fees every period to screen Netflix, in each period trading off low fees that will be accepted quickly against high fees that may be rejected because Netflix's actual valuation is too low.

3.2.1. Actions and timing

Before specifying the maximization problems of the ISPs and Netflix, it is helpful to define the action space of each player and the timing of moves.

Bargaining is exogenously and simultaneously initiated with all ISPs by Netflix at time t_0 . I assume that ISPs take the evolution of plan characteristics, coverage and marginal costs as given, so that for $t < t_0$ they only have to solve repeated static optimization problems with respect to price. At each $t \geq t_0$, an ISP can make a lump-sum transfer offer $\tau_{ft} \in \mathbb{R}$ to Netflix as long as one of its previous offers has not yet been accepted.

At each $t \geq t_0$, Netflix observes the array of offers made by firms that have not yet made an acceptable offer, then decides whether to accept or reject each one. Netflix's action space is the vector $\mathbf{a}_t \equiv (a_{1t}, a_{2t}, \dots, a_{Ft})$, where $a_{ft} = 1$ indicates acceptance of f 's time t offer and $a_{ft} = 0$ indicates rejection. Once an offer has been accepted, bargaining is concluded between Netflix and that ISP. Letting T_f denote the time f 's offer is accepted, $a_{ft} = 1$ for all $t \geq T_f$ and define $\bar{\tau}_f = \tau_{fT_f}$ to be f 's acceptable offer.

At the beginning of each period, all remaining ISPs make offers simultaneously. That is, ISPs experience both incomplete information vis-à-vis Netflix's marginal valuation, and imperfect information regarding each others' actions. After offers are made, Netflix decides which to accept and which to reject. Firms then set prices competitively, and firms whose offers are rejected move into the next period and make new offers.

3.2.2. ISPs' problem

Each ISP seeks to maximize expected profits by choosing a best response sequence of offers. It is best responding both to other ISPs' offers, as well as Netflix's optimal strategy of accepting/rejecting offers. At time t , firm f 's *ex ante* period profits depend on its own strategy τ_{ft} , its competitors' actions τ_{-ft} , and Netflix's strategy $\mathbf{A}_t$. Conditional on these quantities and its

---

¹²If parties are behavioural, as in the optimism model of Yildiz (2004), bargaining delay can still occur even with complete information.

information set $\mathcal{I}_{ft}$, profits are written as:

$$\pi_{ft}(\tau_{ft}, \tau_{-ft}, \mathbf{A}_t, \mathcal{I}_{ft}, \theta) = E_{\mathbf{a}_t} \left[\sum_m \sum_j M_{mt} s_{j|f_{mt}}(\mathbf{a}_t) \cdot s_{f|mt}(\mathbf{a}_t, \theta) \cdot (p_{jft} - mc_{jft}) + a_{ft} \tau_{ft} \mid \tau_{-ft}, \mathbf{A}_t, \mathcal{I}_{ft} \right]$$

Where recall that $\mathbf{a}_t \equiv (a_{1t}, \dots, a_{Ft})$ is the complete realized action vector for Netflix, which is a random variable from the perspective of any given f . That is, because f is uncertain of Netflix's marginal benefit to agreement, it only knows the probability of acceptance (conditional on Netflix's optimal strategy $\mathbf{a}_t^*$).

The first component of profits is the profit due to subscriptions: I sum across markets and f 's offered plans j . Notice that because consumers value the real download speed, which is affected by whether agreement is concluded or not, shares depends directly on the entire $\mathbf{a}_t$ vector. If all of f 's competitors have concluded agreement while f has not, then f will suffer as consumers substitute towards firms with higher quality of service.

The second component of profits $a_{ft} \tau_{ft}$ is equal to the per-period lump sum transfer f stands to receive if its offer is accepted. At time t , I assume that the information set $\mathcal{I}_{ft}$ of ISP f contains the complete history of acceptances and rejections, $\{\mathbf{a}_{t_0}, \dots, \mathbf{a}_{t-1}\}$.

Given period profits, an ISP f 's dynamic problem starting from t_0 is as follows:

$$\max_{\tau_f} E_{\mathbf{a}_{t_0}, \mathbf{a}_{t_0+1}, \dots} \left[\sum_{t \geq t_0} \beta^{t-t_0} \pi_{ft}(\tau_{ft}, \tau_{-ft}, \mathbf{A}_t, \mathcal{I}_{ft}, \theta) \mid \tau_{-f}, \mathbf{A}, \mathcal{I}_{ft_0} \right]$$

Where $\mathbf{A}$ denotes Netflix's complete state contingent strategy and τ_{-f} denotes ISP f 's opponents' complete state contingent strategies.

3.2.3. Netflix's problem

After exogenously initiating bargaining at t_0 , Netflix chooses its strategy of acceptances and rejections to maximize its own profits. Netflix period profits are given as:

$$\pi_{Nt}(A_t, \tau_t) = M_{Nt}(p_{Nt} - mc_{Nt}) + \sum_f (\mu_f - \tau_{ft}) a_{ft}$$

Where M_{Nt} is the exogenously evolving number of Netflix subscribers.¹³

μ_f is the ISP-specific benefit to Netflix from concluding agreement with ISP f . This object is

¹³An alternative approach would be to nest the demand for Netflix in the consumer demand system. However, there are no conditional choice moments available for Netflix demand. Moreover, when looking at aggregate subscribers there does not seem to be a decrease during the period of disagreement. The model, identification and the solution method are all robust to the inclusion of substitution away from Netflix.

the per-period savings in the fixed cost of operation Netflix stands to gain by concluding agreements with the ISPs. It is ISP-specific, and I assume it is comprised of two components: a publicly observable component of the savings in fixed costs, which can depend on firm specific factors that are unique to the firm's infrastructural network, and a private component that is known only to Netflix, and also varies by ISP. I postpone further parameterization of μ_f to the estimation and identification section.

Given period profits, Netflix's dynamic problem starting from t_0 is as follows:

$$\max_{A_1, A_2, \dots} \sum_{t \geq t_0} \beta_N^{t-t_0} \pi_{Nt}(A_t, \tau_t)$$

Notice that I subscript Netflix's discount rate. Although I postpone the full discussion until later, for now I note solving the full two-sided dynamic model is infeasible: I can either assume $\beta_N = 0$ and solve the model, or assume $\beta_N = \beta$ but restrict Netflix's strategy space.

3.2.4. Value function formulation

To make progress on solving the supply side model will require formulating the ISPs' and Netflix's profit maximization problems as value functions. The difficulty of a value function formulation comes from the fact that the current vector of commonly observed agreement plus a firm's own information about μ_f are not sufficient as state variables when there are more than two ISPs.

For instance, with three firms, suppose the state is such that one firm has agreed but the other two have not. Without knowing how long ago the first firm agreed, the second and third firms cannot infer each others' current strategies, which are necessary to determine the transition probabilities to next period's agreement vector. At any given time, if the first firm agreed very early, then the rate of screening by the other two firms would have been different than if the first firm agreed only recently. Each firm must be able to accurately infer what the other knows about its distribution to predict the probability of agreement in this period.

The history of agreement times is commonly observed, and could work as a state variable. However, for any given history, strategies of its opponents, and Netflix optimal strategy, f is able to back out the sequence of optimal offers of all $-f$. This sequence combined with Netflix's optimal strategy gives the points in the distributions of μ_{-f} such that Netflix's true valuation is beneath that point. Call this point b_{-f} . Then b_{-f} is a sufficient statistic to predict $-f$'s optimal strategies going forward.

The value function for ISP f is:

$$V_f(\mathbf{b}, t) = \max_{\tau_{ft}} \{ \pi_{ft} + \beta^{t-t_0} E_{\mathbf{b}_t} [V_f(\mathbf{b}', t+1)] \}$$

For Netflix the problem is deterministic. Their value function is simply

Where τ_t is the vector of offers in the current period, which Netflix fully observes.

3.2.5. Asymmetric information

In the model, asymmetric information is one-sided: ISPs do not observe the marginal benefit to Netflix of agreeing, but Netflix observes the differences in flow payoffs to ISPs. There are two reasons why I make this assumption. The first is that empirically, information about individual ISP profits before, during, and after the slowdown was readily available in the form of quarterly earnings reports. The majority of ISPs in negotiation were publicly traded, and hence Netflix could easily observe how ISPs profits, subscription rates, etc. were being affected by the slowdown. In contrast, although Netflix profits and subscriber rates were also available, their numbers vis-à-vis each ISP were not, and so ISPs were in the dark about how much exactly Netflix valued direct interconnection with each ISP.¹⁴ Essentially, I assume here that ISPs do not observe more than I do as the econometrician when it comes to the game they play against Netflix.

The second reason for assuming one-sided asymmetric information is for tractability: the dynamic screening model has a unique equilibrium in the bilateral case with declining price offers over time. Other bargaining models with delay such as Admati and Perry (REStud, 1987) do not necessarily have unique equilibria, even in the bilateral case. Generating multiple equilibria from the same set of parameter values is not a threat to identification; however, it creates the usual problems in estimation, and would preclude me from using the maximum likelihood approach I do here.

4. Data

4.1. Demand data

The primary data on market shares are gathered from ISPs' quarterly and yearly earnings reports (10-Q and 10-K) which are available for all publicly traded companies in the U.S. Total internet subscriber numbers are given every quarter; combined with auxiliary data on market sizes detailed below, these numbers will be used to construct market shares. These reports also contain ancillary data on mergers and purchases/swaps of infrastructure between ISPs, which must be

---

¹⁴While it is plausible that ISPs can monitor their customers to determine how many are Netflix consumers, this would require a detailed monitoring program that began *before* the disagreement began, while Netflix only need consult publicly available 10-Q and 10-K forms. Thanks to a discussant at the CEA Annual General Meeting for pointing this out.

addressed to avoid bias in estimation. Some ISPs are privately held—e.g. Cox and RCN—in which case I use estimates of the subscription base from Leichtman Research Group. I plot the raw quarterly subscriber numbers in Figure 4.

Although subscriber data is at the national level, most consumers have access to only one or two fixed high-speed internet options. That is, a consumer in the market for fixed internet access is constrained to choose an ISP that has already invested to connect that consumer's housing unit to the ISP's network. Data on what choices are available to consumers comes from the National Broadband Map (NBBM), a government initiative with data available from 2010-2014 which collects information at half-yearly intervals on ISP connections at the census block level. For each census block, ISPs report whether they provide service to that block and their maximum advertised speed. I aggregate all blocks that face that same choice set over time into a market, using the sum of all housing units in the component blocks (reported in the 2010 census) to give the market size. I also assume that all consumers have access to satellite internet as part of their choice set. Table 1 gives the fraction (and cumulative fraction) of the top twenty largest markets by ISP choice set, excluding satellite (which is assumed to be part of every choice set) except where it is the only choice.

Prices, data caps and contract lengths are gathered from several sources, including the FCC Urban Rate Survey and Open Connectivity Database.¹⁵ Where prices are missing, I collect them by hand from stored ISP frontpages on the Internet Archive Project. When the Internet Archive is unable to recover the prices—for instance, due to prices being hidden behind a localization layer—I comb ISP-specific consumer reviews on DSLreports.com. The median cost of a megabit of download speed for the ISPs in my sample is plotted in Figure 7.

In the model a consumer's utility from picking an ISP is affected by whether or not that ISP is currently experiencing slower throughput of Netflix content. Netflix reports the throughput in megabits per second of its content to all major U.S. broadband providers every month; the raw data is provided in Figure 1. For each ISP, I regress throughput on a linear time trend, and code the "slowdown" dummy as 1 if throughput falls below 80% of its predicted value. Alternatively, instead of a dummy, I could divide the realized throughput by the predicted trend to generate a continuous measure. I use the dummy as it will make solving and estimating the model simpler. With the dummy, firms only need to keep track of whether their competitors are experiencing a slowdown or not, and not the intensity of that slowdown. The recovered durations of slowdowns are reported in Table 3. I argue in the next section that durations correspond to periods of disagreement.

Identifying observed and unobserved taste heterogeneity parameters in the model is difficult without more micro moments. I supplement the market-level demand data with two additional datasets on consumer choices.

---

¹⁵No relationship with Netflix's Open Connect CDN.

First, the FCC's Measuring Broadband America (MBBA) program consists of very high frequency testing data for a sample of roughly 10 000 households. I observe household ISPs, their tested download speeds, and which census block group they are in, but I do not observe their demographics. From the tested download speeds I back out which plan within an ISP's menu each household subscribes to. ISP participation in the program is voluntary, so this data is not nationally representative of subscriber shares across ISPs. However, it is representative across speed tiers within ISPs; the data therefore provides moments to recover parameters governing taste heterogeneity.¹⁶

Second, the CPS in 2011, 2013, and 2015 and the ACS in 2013 and 2014 both have internet use supplements, giving 5 observations total. Individuals report what mode of access they used to connect to the internet—satellite, DSL, Cable, Fiber, wireless, or dialup—and detailed demographics. I use this data to construct moments on observed consumer heterogeneity in the taste for technologies parameter. Data is only reported at the CBSA and PUMA level, so I must make assumptions on the choice sets individuals face. Means of demographics conditional on plan choice are reported in Table 2

4.2. Supply data

The most important variable in the supply side analysis is the duration of bargaining delay. Since this data is not directly reported, I construct it from several primary data sources. Constructing this data will involve identifying not only how long bargaining lasted at negotiating ISPs, but also identifying which ISPs were involved in bargaining at all.

The starting point for the supply data are the ISP-specific slowdown durations from Table 3. While I will take slightly modified versions of these durations as indicating when negotiations were taking place, the slowdowns themselves do not prove negotiations were happening: for that I rely on written evidence.

Some proof that the slowdown was indeed a bargaining event, at least for the largest ISPs, comes from the public (partially redacted) version of Netflix's "Petition to Deny".¹⁷ The Petition states that by January, 2014, Netflix and Comcast had reached a deal¹⁸; that agreement happened at the same time as restoration of Netflix service to Comcast shown in Figure 1 provides confirmation that, at least for Comcast, the restoration of service marked the end of negotiations.

¹⁶Households are periodically added in to the program; I only use units that were present in the beginning of the sample to avoid attributing jumps in within-ISP shares due to sample entry to movements in the covariates. Figure 5 gives the number of units in the sample over time, which declines steadily due to attrition.

¹⁷The Petition to Deny is a legal document filed to the Federal Communications Commission, in this case to argue against the application for merger of Comcast and Time Warner that was announced in mid 2014. In making its case for why these entities should not merge, Netflix detailed the difficulties it had during the installation of its Open Connect servers in 2013. Perhaps because its goal in the Petition to Deny was to argue the dangers of too much market power, only the four largest ISPs are explicitly named in the document.

¹⁸Petition to Deny, pg. 57, paragraph 2 — pg. 58 paragraph 2.

The document also notes that, by the end of 2014, "none of the U.S.'s four major ISPs [had] agreed to partner with Open Connect without payment", implying that the parties were indeed negotiating over explicit transfers from Netflix to the ISPs.¹⁹ That service was also restored to all ISPs at varying times—but all before the end of 2014—is suggestive that periods of negotiation at each ISP were concordant with periods of slower service.

Additional proof comes from Netflix's 2014 first quarter letter to its investors, issued on April 21, 2014. Netflix notes "now nearly all cable Internet households receive great quality Internet video", implying that Time Warner Cable also concluded negotiations by the first quarter of 2014.²⁰ From the same document, Netflix mentions the extremely poor QoS that AT&T U-Verse customers receive, and argues that "[it] is free and easy for AT&T to interconnect directly with Netflix and quickly improve their customers' experience, should AT&T so desire.", implying that the slowdown could be alleviated as soon as Netflix and AT&T could settle on a price. From this report it is clear that AT&T actually took longer to resolve negotiations, so that the slowdown truly indicates negotiation time and not just greater technological difficulties implementing interconnection with Open Connect.

One can reasonably conclude from the above evidence that periods of degraded QoS match periods of negotiation between the largest ISPs and Netflix. For smaller ISPs this may not necessarily be true. Large last-mile ISPs—especially AT&T — sometimes act as transit ISPs for smaller last-mile ISPs, implying that the duration of the QoS reduction of a small last-mile ISP and its larger business partner may be linked, even if the small ISP concludes a deal earlier. Since I have no data on transit agreements, I assume that the duration of slowdown is a valid measure of negotiation length even for small ISPs. I postpone a discussion of the potential bias this may cause, especially for counterfactuals, to the estimation section.

5. Reduced form evidence

The structural model posits that bargaining delays result from asymmetric information, and that bunching in agreement timings can be explaining by strategic competition between ISPs due to marginal consumers' willingness to pursue higher QoS. Before estimating the model and performing the counterfactual, it is worth evaluating whether or not there is reduced form evidence for these mechanisms in the data.

I first analyze whether consumers' observed ISP choices suggest that they value Netflix streaming quality. I find that consumers do substitute away from ISPs that are in a state of disagreement with Netflix. This suggests that not only do consumers value Netflix QoS, they value it enough (and are informed enough) to leave ISPs (or at least, not sign up with ISPs) whose streaming

---

¹⁹Petition to deny, pg. 49, paragraph 2.

²⁰Netflix 2014Q1 letter to investors, pg. 5 paragraph 3.

quality is degraded.

I also examine whether consumers change their speed upgrading behaviour at affected ISPs during periods of disagreement. Consumers frequently upgrade their speed without switching ISPs and since higher speeds are offered to the same household at the same marginal cost, higher speeds are more profitable for the ISP to offer. Therefore, an alternative theory of the bargaining delay might involve ISPs reducing streaming quality to force marginal consumers to move to higher, more profitable speed tiers. I show using individual level data that consumers do not disproportionately increase their rate of upgrading at ISPs affected by the disagreement, falsifying this hypothesis.

Finally, I explore whether ISPs that share more markets are more likely to have simultaneous agreement timings. Although there is very limited data, the evidence is suggestive that increased market overlap leads to an increased likelihood of simultaneous agreement. I also justify the use of the one-sided asymmetric information screening model by looking at institutional details of the industry.

5.1. Consumer substitution

Consumer substitution between ISPs may induce profit complementarities that can help the model match simultaneous agreement timings.²¹ To test whether this substitution holds and is empirically relevant, I run the following regression using the ISPs' reported subscriber data, and a dummy variable indicating whether an ISP is currently in disagreement with Netflix:

$$\Delta \log subscribers_{it} = \beta_1 disagree_{it} + \beta_2 \Delta \log p_{it} + \gamma_i + \gamma_t + \epsilon_{it}$$

If the QoS degradation is having a persistent negative effect on the growth rate of subscribers, one would expect β_1 to be negative and significant. The coefficient on $\Delta \log p_{it}$ should also be negative, although potential endogeneity with the supply curve or unobserved demand shocks may lead it to be positive.

Results are reported in Table 4. The first four columns are simple OLS regressions with $gamma_i = \gamma$, $\gamma_t = 0$. Column (1) includes all data, including time periods when a merger happens. Column (2) removes mergers; as these are large, noisy events, removing them increases significance even as it reduces the point estimates. Columns (3) and (4) repeat these exercises, but for a different formulation of disagreement times, increasing the duration of disagreement of the first group of ISPs from 2 to 3 quarters.

The fixed effect regressions (5) and (6) add in ISP and time specific dummies to control for firm level trends in subscribers as well as country-wide shocks, and (7) and (8) also control for ISPs' median price. All coefficients are negative, although estimates using the second measure

---

²¹A second mechanism may be through correlated μ_f , which I will pursue in future work.

for disagreement are not significant. Price has the expected negative coefficient, although it is not significant, potentially due to endogeneity with unobserved demand shocks.

Firms whose QoS was degraded experienced a roughly 1.5% decrease in subscriber growth compared to firms that reached agreement; the results suggest an economically significant role for subscriber substitution between ISPs in response to QoS problems with Netflix.

5.2. Consumer speed upgrading

The model postulates asymmetric information as the cause of the bargaining delay; another explanation might be that firms, recognizing that switching costs between ISPs are high, may try to force consumers to upgrade their way out of the QoS reduction. Assuming that ISPs are rational and able to anticipate consumer behaviour, if I find that consumers of affected firms upgrade speed at a higher rate than unaffected firms, it is evidence for this alternative story.

I use a monthly panel of consumer ISP choices constructed from testing data associated with the FCC's Measuring Broadband America program. An example upgrade path over time is presented in Figure 6 (the blue line denotes when the consumer was with Verizon, and the red, Comcast) and sample statistics from the raw data are presented in ??.

I run the following LPM specifications:

$$upgrade_{it} = \beta_0 + \beta_1 slow_{it} + \beta_2 \log(d_{it-1}) + \beta_3 slow_{it} \times \log(d_{it-1}) + \beta_4 t + \alpha_i + \epsilon_{it}$$

The specification estimates the probability of upgrading speed as a function of whether or not your ISP experienced a quality degradation ($slow_{it}$), your speed tier last month (d_{it-1}), and an interaction. As mentioned, this specification is designed to falsify the theory that, due to high switching costs, consumers simply stayed with their ISP and upgraded their way out of the slowdown. This theory would imply that profits from consumer subscriptions actually increase during a slowdown (assuming bandwidth is provided at zero marginal cost), which would suggest an alternative reason for inducing bargaining delays.

Results are reported in Table 5. The coefficient on $slow_{it}$ is negative and significant, implying that far from increasing, upgrading actually decreased for ISPs experiencing a slowdown.

5.3. Strategic overlap

Section 5.1. suggests that consumers leave ISPs that are heavily affected by the quality of service degradation due to prolonged bargaining. Do ISPs recognize this substitution, and respond to it in a strategic way? That is, is the probability that two ISPs conclude agreement at the same time increasing if they compete in more markets?

To analyze this question in a reduced form way, I estimate the following specification using

data from June 2013 to June 2014 inclusive:

$$agree_{it} = \beta_0 + \beta_1 comp\_agree_{it} + \gamma X_{it} + \epsilon_{it}$$

$agree_{it}$ takes a value of 1 if ISP i has concluded negotiations with Netflix at any time $\tau \leq t$. Construction of this variable is discussed in the data section, ???. $comp\_agree_{it}$ takes a value of 1 if i 's "primary competitor" has concluded negotiations at any time $\tau \leq t$, and X_{it} is a (limited) selection of covariates.

$comp\_agree_{it}$ is constructed from the FCC's National Broadband Map. The broadband map gives housing unit pass-by rates for each ISP at the census block level. i 's primary competitor is defined to be the ISP present more often than any other in the markets i serves.

Because there are currently only 11 firms with positive bargaining delays, I restrict the number of additional covariates to 1 in each specification.²² I use the average offered speed of the ISP, an estimate of the fraction of the ISP's consumers that use streaming video, and an estimate of the fraction of the ISP's consumers that have bundled TV service. The first covariate is taken from data I have gathered on ISP pricing, and the latter covariates are constructed using the CPS Internet Use Survey supplementary data. Results are presented in Table 6

The $prim\_comp_{it}$ variable is positive and significant in 2 out of the 3 specifications. At first glance, this provides suggestive evidence that interactions between ISPs matter for agreement timings. However, this reduced form exercise cannot distinguish between whether Netflix's unobserved marginal values of agreement are more positively correlated among ISPs that compete in more markets, or whether ISPs' own strategic interactions motivate the simultaneity.

6. Estimation and Identification

6.1. Demand

I use the standard Berry, Levinsohn and Pakes (1995) approach to estimate demand parameters, with the addition of micro moments as in Berry, Levinsohn and Pakes (2004) to help identify heterogeneity in parameters across the population and provision for choice sets that vary across individuals as in Goeree (2008). I first recover estimates of the unobserved firm (and time) specific heterogeneity ξ_{ft} as a function of demand parameters, denoted θ_d . Next, for choices of instruments Z_{ft} , I assume that at the true θ_d^* , $E[Z_{ft}\xi_{ft}(\theta_d^*)] = 0$. I then minimize a GMM objective criterion based on the empirical counterparts to the above moments, along with additional micro moments that I detail below.

This approach does not nonparametrically identify the distributions of unobserved con-

---

²²With the addition of firms that agree immediately, the sample size may grow.

sumer heterogeneity. I assume that unobserved consumer heterogeneity is distributed multivariate normal with a diagonal covariance matrix. That is, within a market, unobserved technology preferences, and unobserved preferences for price and quality are distributed as:

$$\nu_{ig} \sim \alpha_g^u \mathcal{N}(0, I_1) \quad \forall g \quad (\nu_{ip}, \nu_{iq}) \sim (\alpha_p^u, \alpha_q^u)' \mathcal{N}(\mathbf{0}, I_2)$$

Where g indexes the available technologies (Cable, DSL, Fiber, and Satellite.)

From Equation 1, the demand parameters θ_d to estimate are therefore:

$$\begin{aligned} \text{Observed heterogeneity: } & \{\alpha_g^o, \alpha_p^o, \alpha_q^o\} \\ \text{Unobserved heterogeneity: } & \{\alpha_g^u, \alpha_p^u, \alpha_q^u\} \end{aligned}$$

To estimate these parameters, I match three "sets" of predicted moments to their sample analogues: (1) the covariance between unobserved firm-level heterogeneity and a set of instruments that shift firm markups; (2) the covariances of the observed technology type with observed consumer characteristics; and (3) the covariance of the set of instruments with the difference between the predicted conditional plan shares and observed conditional plan shares.

The first set of moments are useful in identifying the mean consumer valuation for each of the different types of technologies, prices and qualities. Given parameter guesses, ξ_{ft} will allow the model to exactly fit the observed market shares by shifting the mean utility for each firm at each date. Interacting these ξ_{ft} with a set of contemporaneous and lagged instruments will help identify the parameters even in the presence of latent switching costs (see Scherbakov (2015)).

The second set of moments match observed consumer attributes from the ACS to each consumer's chosen technology. These moments will be particularly useful in identifying α_g^o . Since bundling opportunities vary greatly by technology, allowing consumer valuation of different technologies to vary by a rich set of consumer characteristics will be crucial to recovering credible estimates the coefficients on download speed, and especially on price. If w denotes a set of technology dummies and consumer characteristics are denoted by x , then I fit the model's predictions for $E[x'w]$ and $E[w]$ to their sample analogues from the ACS.

The final set of moments will help identify unobserved heterogeneity in the valuation of price and download speed. I assume that the difference between the predicted conditional shares and the observed conditional shares is classical measurement error: $u_{j|ft} = s_{j|ft}(\theta, X) - \hat{s}_{j|ft}$, and is mean-independent of the instruments at the true parameter values.

6.1.1. Recovering $\xi_{ft}(\theta_d)$

Forming the first set of moments requires estimates of $\xi_{ft}(\theta_d)$. I recover the ξ_{ft} that rationalize the observed market shares according to the standard BLP algorithm.

Notice in order to choose the ξ_{ft} that matches the predicted to the observed shares, one must first calculate the predicted shares, which involves taking an integral taken over the joint distribution of (x, ν) . Distributional assumptions on ϵ ensure a closed form for the individual probability of purchase; however, for the remaining variables I follow Pakes (1986) and use simulation to approximate the integral's value.

6.1.2. The fitted moments

The first set of moments are, for instruments z_{ft} ,

$$G_{ns}^1(\theta_d) = \frac{1}{T} \sum_t \frac{1}{F_t} \sum_f z'_{ft} \xi_{ft}(\theta_d) \quad (1)$$

Where F_t gives the number of firms operating at time t , as some firms merge over the course of the sample.

ξ_{ft} is typically interpreted as a shock to demand for firm f 's products that ISPs observe but the econometrician does not. Here they are firm wide, and correspond intuitively to changes in the quality of bundled services, system-wide disruptions in reliability, or the quality of customer service. Since the shocks are observed by f , they are endogenous to "flexible" (contemporaneously set) inputs in the firm's profit maximization problem. I assume that only price is set contemporaneously.

The instruments for f at time t I use are variables that affect f 's price at time t but do not depend on ξ_{ft} . I assume that the menu of download speeds evolves exogenously, according to background technological progress that enables faster speeds over the same infrastructure. Functions of competitors' contemporaneous menu of download speeds, as well as f 's own speeds, are valid instruments under this assumption. I also use lagged functions of the instruments to separately identify preferences from switching costs, as in Scherbakov (2015): conditional on current plan characteristics, lagged characteristics can only have an effect on current shares in the presence of switching costs.

The second set of moments use data from the ACS. The ACS provides micro-level data on which technology g each i in the ACS chooses to connect to the internet, as well as demographics x . The moments are:

$$G_{ns}^2(\theta_d) = \sum_m \frac{n_m}{n} \sum_g \frac{n_{gm}}{n_m} w_{gm} \left[\frac{1}{n_{gm}} \sum_{i_{gm}=1}^{n_{gm}} x_{i_{gm}} - E[x|c_{imt} = g, \theta_d] \right] \quad (2)$$

Where n is the total number of ACS respondents, n_m is the number in market m , and n_{gm} is the number reporting they use technology g in market m . Recall that w_g is a set of technology dummies.

The second term in the square brackets is the model's prediction of the mean characteristics of an individual choosing technology g . I estimate it using the simulated individuals:

$$E[\mathbf{x}|c_{imt} = g, \theta_d] \approx \frac{(ns)^{-1} \sum_r \mathbf{x}_r Pr(c_{rmt} = g|\mathbf{x}_r, \nu_r, \theta_d)}{s_g}$$

The final set of moments use data from the MBBA testing program. These moments are similar in spirit to the second set of moments, interacting the difference between actual and predicted conditional shares with an instrument. The moments are:

$$G_{ns}^3(\theta_d) = \frac{1}{T} \sum_t \frac{1}{\tilde{F}_t} \sum_f z'_{ft} \left[\frac{n_{j|ft}}{n_{ft}} - s_{j|fmt}(\theta_d) \right] \quad (3)$$

Where $\tilde{F}_t$ is the number of firms present in the MBBA data at time t (with $\tilde{F}_t \leq F_t$), $n_{j|fmt}$ is the number of individuals in the data who choose j as their plan—given that they choose f in market m at time t —and n_{fmt} is the number of individuals who choose firm f in market m at time t . The model's predicted conditional share requires integration, so once again we simulate to recover:

$$s_{j|ft}(\theta_d) = \sum_{m|f \in m} \tilde{w}_m \frac{\frac{1}{ns} \sum_r \mathbf{1} \left[j = \underset{j'}{\operatorname{argmax}} \{ \alpha_{rp} p_{jft} + \alpha_{rq} q_{jft} \} \right] s_{rfmt}(\theta_d)}{s_{fmt}}$$

I stack $G^1(\cdot)$, $G^2(\cdot)$ and $G^3(\cdot)$ and use the two-step GMM method of θ_d from the stacked moments. As Hansen (1982) shows, provided that $ns \rightarrow \infty$, and MBBA and ACS sample sizes go to infinity, the estimator will be consistent. A discussion of the assumptions on the simulated sample size and how the simulation error affects the variances is postponed until later.

6.2. Supply

Since ISP profits are estimated separately using demand data, the only parameters to estimate on the supply side are associated with μ_f . I assume that μ_f is the product of an observable function of ISP attributes w_f , and a shock that is observed by Netflix but not by the ISP or the econometrician, ζ :

$$\mu_f = \exp(\mu' w_f) \exp(\sigma_\zeta \zeta_f)$$

I assume that ζ is distributed truncated lognormal, with each draw ζ_f independent and identically distributed.

To estimate $\{\mu, \sigma_\zeta\}$, I will rely on maximum likelihood. From the data section, I only have

observed agreement/disagreement timings $\{T_f\}$ so the likelihood function will have the form:

$$\mathcal{L}(\{T\}_f; \mu, \sigma_\zeta) = P(\hat{T}_1(\mu, \sigma_\zeta) = T_1, \dots, \hat{T}_F(\mu, \sigma_\zeta) = T_F)$$

In theory, for the likelihood to be well-identified, one of two situations must be true: the model may give unique predictions for agreement times for a given set of parameter values. Alternatively, the model may have multiple equilibria for given $\{\mu, \sigma_\zeta\}$, but as long as the data is only generated by one equilibrium and I select the correct predicted equilibrium when forming the likelihood I will recover the correct parameter values.

To construct the likelihood I must find the predicted agreement times as a function of parameters, which will follow from recovering the optimal policies. Below I detail how I recover optimal policies and in doing so, select an equilibrium.²³

6.2.1. Policy functions

I start by recovering Netflix's policy function. Recall that in each period t , after fees have been proposed by all ISPs that have not made an acceptable offer in a previous period, Netflix has the chance to accept or reject each offer. In a given period, Netflix's action space is enormous: with $J > 30$ ISPs, at the beginning of the game Netflix has 2^J possible actions.

To simplify this problem I assume for now that $\beta_N = 0$ so that Netflix is myopic. Combined with the separability of the marginal gains to agreement with each ISP μ_f , assuming that $\beta_N = 0$ implies that Netflix considers each offer independently. In particular, although Netflix might like to commit to a more sophisticated strategy *ex ante*, once offers have been made it can do no better than accepting all τ_f such that $\tau_f \leq \mu_f$.²⁴

Given Netflix's strategy, I now characterize ISPs' optimal state contingent policies. ISP f has the following information structure: Netflix's optimal responses are separable, so that Netflix's choice to accept or reject any other ISP $-f$'s offer does not give f any information about μ_f . Only Netflix's response to f 's offer gives information about the distribution of μ_f . Based on Netflix's optimal rule to accept if $\tau_f \leq \mu_f$, if f 's offer τ_{ft} in period t is rejected, they know that the conditional distribution of μ_f is now $F(\mu_f | \mu_f < \tau_f, w_f)$.

Therefore, if there were no profit interactions, ISP f 's optimal strategy would be a strictly decreasing sequence of offers. With profit interactions, f still optimally chooses a strictly decreasing sequence of offers. First, the screening incentive remains. Second, since ISPs are substitutes from the consumer's perspective, consumers will naturally leave ISPs with relatively lower qual-

²³This may be the unique equilibrium—I do not have a proof for uniqueness, and have not yet experimented with different initial conditions to see if they converge to the same policy functions.

²⁴One can also set $\beta_N = \beta$, and maintain independence by requiring the analogue of the static Nash-in-Nash bargaining assumption. That is, when considering how to best respond to f , Netflix takes its decision rules with respect to all other ISPs as given. I assume myopia for now as it reduces the computational burden.

ity of service; therefore ISPs have an incentive to agree more quickly to steal business from rivals, and prevent their own business from being stolen. These incentives increase the marginal gains to agreement for f compared to the setting without profit interactions. Given Netflix's myopic strategy, these incentive remains regardless of f 's competitors' strategies. In other words, f has a dominant strategy to make a strictly decreasing sequence of offers, so in equilibrium all firms make a strictly decreasing sequence of offers.

While the above argument holds regardless of the size of f 's state space, the model quickly becomes computationally intractable if firms care about the states and strategies of all other firms. I adopt a cutoff rule to mitigate this curse of dimensionality. First, for f consider all $-f$ such that each $-f$ overlaps at least 20% of f 's footprint. Starting with Comcast, determine all such $-f$ —Comcast's primary competitors—and then determine the $-f$ of those primary competitors. Following this procedure until no new firms are added will form a minimal group that cares about the actions of all other members of that group, even if some members' footprints do not overlap.

The minimal group constructed in such a way is Comcast, AT&T, TimeWarner and Verizon. I assume when solving these firms' optimal policies that they pay attention to the markets that are exclusively served by members of the group. Other ISPs react to the policies of the biggest four, but do not affect the four in turn. Going by the 20% cutoff, a few smaller regional ISPs also optimize in response to each other—it is the smallest ISPs who have the largest state space, as they must pay attention to all the large firms that can affect their flow profits, either directly or indirectly.

6.2.2. Numerical example for two ISPs

The model with profit interactions can induce bunching in agreement times, even for small differences in profit. I construct a numerical example with two ISPs and map out the joint probability distribution for agreement timings; results are shown in Figure 8.²⁵

The probability of simultaneous agreement in the model with profit interactions increases to 19.9% from 14.3% in the model without interactions. Indeed, the joint distribution of probabilities exhibits a much higher "ridge" along the diagonal of agreement times, reflecting the fact that it is quite costly to be in a situation where an ISP's opponent has concluded agreement while that ISP has not.

7. Estimation results

---

²⁵ $\sigma = 1$, $\kappa = 0$, $\pi = (1.1, 1, 0.95, 0.9)$ for the ordering in ???. ISPs are symmetric.

7.1. Demand estimates

I first present results from the estimation of the price and speed sensitivities in Table 7. Moving from the logit, to microBLP, and finally to the full model with variation in choice sets, the sign and magnitude of the estimates are fairly stable. The slowdown dummy is negative, but somewhat marginal in all specifications.

As expected from the conditional mean tables, income and household size are both correlated with a reduction in the price elasticity of demand (recall that the coefficients multiplied by the variables enter W in $-e^{-W}$, the price coefficient.) Valuation of download speed decreases with age as expected, but also decreases with household size, which is the opposite of expectations and bears closer inspection.

I look at two simple measures of model fit. For the full model, I check whether the shares of technologies predicted by the model match reasonably well those in the data for the beginning of 2013. Results are reported in Table 9—the predictions are quite close.

I also examine the short run price and speed elasticities of substitution. For speed, turning on the agreement dummy on increases the share by roughly 0.8 percent. Table 8 reports the price elasticities by technology. Despite the model flexibility, the elasticities (reported by technology) are still below one in general.

7.2. Marginal value of agreement estimates

The demand estimates from the previous section are combined with the dynamic multilateral bargaining assumptions and disagreement durations to estimate the distribution of Netflix's marginal value of agreement vis-à-vis each ISP. Estimates of the coefficients of the mean shifting variables w_f , as well as the standard error of the dispersion of information ζ , are presented in Table 10.

From the table, it is clear that interconnecting with providers whose network is largely DSL to the end consumer is not as valuable to Netflix as interconnecting with other technologies. This result comes from the empirical regularity that DSL providers with as much to lose from delayed agreement as alternative technologies will, on average, agree more quickly. That is, there must be less to be gained from a fine screening strategy. Moreover, the larger a provider's network, the greater the value in interconnection. The fact that this value is logarithmic in the ISP's pass-by rate is a crucial parametric choice, as it will imply that a large increase in pass-by—due to a merger for instance—does not lead to a substantial shift in the mean of the distribution.

The model fits the empirical regularities of the data: for the four largest ISPs, *ex ante* expected agreement times are 3.6 quarters for AT&T, 2.8 quarters for Comcast and TimeWarner, and 2.9 quarters for Verizon, matching the fact that AT&T experienced a much longer delay in agreement than any of the other largest ISPs.

8. Counterfactual

8.1. Comcast-TimeWarner merger

With both the supply and demand estimates in hand, I turn to the paper's counterfactual—the effect of a merger between Comcast and TimeWarner.

There are three dimensions on which to evaluate the merger. Do the *ex ante* predicted negotiated fees change? Do bargaining times change? And does consumer welfare change? The first criterion speaks to the increased degree of extractive power held by the larger merged firm—which affects the split of surplus in upstream bargaining, but not total surplus. The latter two criteria speak directly to dead weight loss and consumer welfare, which are the typical targets for judging the value of a merger. In this case the delay in agreement generates a deadweight loss: since it is only to split surplus, and not create it, and since delay reduces firm profits and consumer welfare, it is inefficient.

The table of mean fees and disagreement times, pre (Table 11) and post (Table 12) merger, reveal the heterogeneity in responses to the merger across ISPs. While the average overall delay increases for the largest firms post merger, this is mainly driven by AT&T — the disagreement length for the merged firm actually goes down. This follows largely from the fact that while an increase in footprint increases the mean of the distribution of surplus, an ISP will only be able to take advantage and screen more finely if there is relatively less to lose in flow profits from disagreement. Essentially, losses in profits from subscribers increase linearly in footprint while the benefit of agreement does not, implying that for the estimated parameters, making firms larger will decrease bargaining length even as it increases per-ISP fees.²⁶

Finally, to speak to the inefficiency created by dynamic bargaining, I estimate consumer welfare under the factual and counterfactual scenarios. I find that the median U.S. consumer values the internet options they have available at 113 USD per month, dropping to 109 USD if all available options have degraded Netflix QoS. I compute the *ex ante* expected welfare for the median consumer over the 6 periods of bargaining in both scenarios, and find that their welfare decreases by 1.9 percent if Comcast and TimeWarner were allowed to merge. However, this number is somewhat misleading, as the customers of Comcast and TimeWarner actually benefit from a 3.1 percent increase in welfare under the new scenario as bargaining times decrease. AT&T's customers suffer a 4.2 percent drop in welfare.

---

²⁶Testing alternate parametrizations for the mean shifter is crucial to making sure this result is not simply a artifact of the parametrization.

9. Conclusion

This paper combines a model of dynamic multilateral bargaining with demand for broadband internet to evaluate the effect of a proposed merger. I estimate an industry model of demand, plan choice, pricing and interconnection bargaining using data on plan prices, consumer choice sets and bargaining delays between major U.S ISPs and the leading purveyor of streaming video content, Netflix. To map disagreement times and ISP competition into interconnection fees, I develop a dynamic bargaining model with asymmetric information. ISPs make take-it-or-leave-it offers to learn about Netflix's benefit from interconnection, while simultaneously competing for subscribers who value Netflix quality of service. Intuitively, if delaying agreement over interconnection degrades quality of service to subscribers, then the opportunity cost of lost subscriptions identifies the fee.

I structurally estimate the model and recover fixed interconnection fees ranging from 44 to 69 million USD. I find that a proposed merger between TimeWarner and Comcast that was challenged by the Federal Communications Commission would have slightly raised interconnection fees and bargaining length, reducing consumer welfare by 1.9 percent. This analysis could be improved by estimating a richer model of demand, incorporating individuals' preferences for different types of streaming content in two stage utility function as in Crawford and Yurukoglu (2012). By incorporating substitution between types of streaming content into consumers' decisions, the information encoded in the bargaining event could be used as a laboratory to test out the effects of vertical integration between ISPs and video streaming services. This work is ongoing.

References

- Ambrus, Attila, Eric Chaney, and Igor Salitsky, "Pirates of the Mediterranean: An Empirical Investigation of Bargaining with Transaction Costs," Technical Report 2016.
- Berry, Steven, James Levinsohn, and Ariel Pakes, "Automobile prices in market equilibrium," *Econometrica*, 1995, pp. 841–890.
- , —, and —, "Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Car Market," *Journal of Political Economy*, 2004, 112 (1), 68–105.
- Björkegren, Daniel, "The Adoption of Network Goods: Evidence from the Spread of Mobile Phones in Rwanda," Technical Report 2015.
- Crawford, Gregory and Ali Yurukoglu, "The Welfare Effects of Bundling in Multichannel Television Markets," *American Economic Review*, 2012, 102 (2), 643–685.
- Dafny, Leemore, Kate Ho, and Robin Lee, "The Price Effects of Cross-Market Hospital Mergers," Technical Report, National Bureau of Economic Research 2016.
- Economides, Nicholas and Benjamin E. Hermalin, "The Economics of Network Neutrality," *RAND Journal of Economics*, 2012, 43 (4), 602–629.
- Fershtman, Chaim and Ariel Pakes, "Dynamic Games with Asymmetric Information: A Framework for Empirical Work," *Quarterly Journal of Economics*, 2012, 127 (4), 1–51.
- Gans, Joshua, "Weak Versus Strong Network Neutrality," *Journal of Regulatory Economics*, 2015, 47 (2), 183–200.
- Goeree, Michelle Sovinsky, "Limited Information and Advertising in the U.S. Personal Computer Industry," *Econometrica*, 2008, 76 (5), 1017–1074.
- Gowrinsankaran, Gautam, Aviv Nevo, and Robert Town, "Mergers When Prices Are Negotiated: Evidence from the Hospital Industry," *American Economic Review*, 2015, 105 (1), 172–203.
- Grennan, Matthew, "Price Discrimination and Bargaining: Empirical Evidence from Medical Devices," *American Economic Review*, 2013, 103 (1), 145–177.
- Honoré, Bo and Aureo de Paula, "Interdependent Durations," *Review of Economic Studies*, 2010, 77 (3).
- Jia, Panle, "What Happens When Wal-Mart Comes to Town: An Empirical Analysis of the Discount Retailing Industry," *Econometrica*, 2008, 76 (6), 1263–1316.
- Lee, Robin S. and Tim Wu, "Subsidizing Creativity through Network Design: Zero-Pricing and Net Neutrality," *Journal of Economic Perspectives*, 2009, 23 (3), 61–76.
- Peitz, Martin and Florian Schuett, "Net Neutrality and Inflation of Traffic," *International Journal of Industrial Organization*, 2016, 46, 16–62.

Rubinstein, Ariel, "Perfect Equilibrium in a Bargaining Model," *Econometrica*, 1982, 50 (1), 97–110.

Scherbakov, Oleksander, "Measuring Consumer Switching Costs in the Television Industry," Technical Report 2015.

Yildiz, Muhamet, "Waiting to Persuade," *Quarterly Journal of Economics*, 2004, 119 (1), 223–248.

A Figures and Tables

Figure 1: Average Netflix throughput to 27 U.S. ISPs

Table 1: Choice set stats

Market	Fraction	Cumulative Fraction
AT&T-COMCAST	0.154	0.154
AT&T-TIMEWARNER	0.098	0.252
COMCAST-VERIZON	0.091	0.343
CENTURYLINK-COMCAST	0.067	0.410
TIMEWARNER-VERIZON	0.053	0.463
AT&T-CHARTER	0.051	0.514
CABLEVISION-VERIZON	0.044	0.558
AT&T-COX	0.030	0.588
HUGHES	0.026	0.614
CENTURYLINK-COX	0.026	0.639
RCN-VERIZON	0.019	0.658
COMCAST-FRONTIER	0.017	0.675
CENTURYLINK-CHARTER	0.016	0.691
FRONTIER-TIMEWARNER	0.016	0.707
AT&T-SUDDENLINK	0.013	0.721
AT&T-BRIGHTHOUSE	0.013	0.734
CENTURYLINK-TIMEWARNER	0.013	0.746
COX-VERIZON	0.012	0.759
AT&T-WIDEOPENWEST	0.012	0.771
BRIGHTHOUSE-VERIZON	0.011	0.782

Figure 2: Average Internap Throughput

Figure 3: Average Cogent Throughput

Table 2: Conditional means of demographics from the ACS

	Cat.	FamInc	HHnum	Eng	Male	Age
Uncond.	—	89.220	2.193	0.970	0.591	51.860
Internet	No	43.774	1.789	0.939	0.512	61.543
	Yes	97.801	2.489	0.979	0.614	51.098
Tech	Cable	100.704	2.460	0.979	0.606	50.446
	Dialup	63.112	2.067	0.975	0.590	61.992
	DSL	90.060	2.531	0.979	0.625	52.297
	Fiber	121.480	2.607	0.984	0.636	51.684
	Mobile	79.529	2.495	0.970	0.602	48.401
	Satellite	77.424	2.440	0.981	0.631	54.976

Figure 4: Aggregate subscribers

Figure 5: Sample attrition

Figure 6: Sample upgrade path from MBBA

Figure 7: ISP pricing over time

Table 3: ISP disagreement durations

AT&T	12
AT&T U-Verse	12
Brighthouse	5
Charter	6
Centurylink	12
Comcast	6
Frontier	6
Mediacom	12
TimeWarner	6
Verizon	6
Verizon FiOS	6

Table 4: Consumer substitution

	<i>OLS</i>				<i>FE</i>			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Disagree	−0.022* (0.012)	−0.015*** (0.005)			−0.013** (0.006)		−0.013** (0.006)	
Disagree2			−0.018 (0.011)	−0.011** (0.004)		−0.009 (0.006)		−0.009 (0.006)
$\Delta \log(\text{price})$							−0.002 (0.004)	−0.002 (0.004)
Constant	0.013*** (0.003)	0.006*** (0.001)	0.013*** (0.003)	0.006*** (0.001)				
Observations	541	526	541	526	526	526	526	526
R ²	0.006	0.019	0.005	0.012	0.641	0.638	0.641	0.638
Adjusted R ²	0.004	0.017	0.003	0.010	0.604	0.600	0.604	0.600

Note:

Standard errors are clustered at the ISP level. *p<0.1; **p<0.05; ***p<0.01

Table 5: Change in switching probabilities

	Upgrade speed	
	(1)	(2)
$slow_{it}$	-0.010*** (0.001)	-0.008* (0.003)
$\log(d_{it-1})$		-0.071*** (0.002)
$slow_{it} \times \log(d_{it-1})$		-0.001 (0.001)
t		0.001*** (0.000)
FE	i	i
Clusters	9816	9816
R^2	0.03	0.054

***, ** and *: 0.1%, 1% and 5% significance. Standard errors clustered at the unit level.

Table 6: Synchronicity in agreement timings among primary competitors

	(1)	(2)	(3)
<i>prim_comp</i>	0.015*	0.008	0.010*
	(0.008)	(0.007)	(0.006)
<i>speed_{it}</i>	0.16		
	(0.14)		
<i>streaming_video_{it}</i>		-0.01	
		(0.03)	
<i>bundle_{it}</i>			0.04*
			(0.01)
<i>N</i>	11	11	11
<i>T</i>	24	24	24
<i>R</i> ²	0.01	0.015	0.018

***, ** and *: 0.1%, 1% and 5% significance. Standard errors clustered at the unit level.

Figure 8: No interactions

Figure 9: Profit interactions

Table 7: Non-linear parameter estimates

Plan Characteristic and Household Characteristic	Logit	microBLP	Full Model
Slowdown			
	-0.01 (0.01)	-0.01 (0.008)	-0.01 (0.006)
Price:			
Constant	-1.29 (0.53)	-1.35 (0.47)	-1.37 (0.52)
Income	0.03 (0.01)	0.02 (0.01)	0.01 (0.01)
Household size	1.00 (0.38)	1.08 (0.47)	0.77 (0.3)
Unobserved		0.62 (0.25)	0.67 (0.22)
Download speed:			
Age	-0.01 (0.00)	-0.01 (0.00)	-0.01 (0.00)
Speak English	0.61 (0.23)	1.23 (0.44)	1.23 (0.4)
Household size	-0.81 (0.28)	-0.88 (0.34)	-0.89 (0.32)
Male	0.01 (0.01)	0.02 (0.01)	0.02 (0.01)
Unobserved		0.00 (0.01)	0.00 (0.00)

Table 8: Own-price elasticity by technology

Technology	Logit	microBLP	Full Model
DSL	-1.223	-1.172	-1.012
Fiber	-3.199	-3.100	-1.264
Cable	-0.564	-0.501	-0.705
Mobile	-0.302	-0.311	-0.805
Dialup	-0.298	-0.500	-0.835
Satellite	-0.783	-0.862	-0.599

Table 9: Model fit: predicted versus actual tech choice

	Category	Data	Model
Internet	Yes	0.783	0.749
	No	0.217	0.256
Technology	Cable	0.558	0.541
	DSL	0.205	0.218
	Fiber	0.18	0.16
	Satellite	0.022	0.025

Table 10: Supply model estimates

Variable	Estimate
Constant	-0.540
DSL dummy	-0.130
log(Pass-by)	0.080
σ_{ζ}	0.980

Table 11: *Ex ante* mean delays and fees, no merger

ISP	Delay (Quarters)	Fee (100m USD)
AT&T	3.600	1.297
COMCAST	2.780	1.222
TIMEWARNER	2.780	1.222
VERIZON	2.965	1.085

Table 12: *Ex ante* mean delays and fees, merger

ISP	Delay (Quarters)	Fee (100m USD)
AT&T	4.093	1.351
MERGED	2.460	1.330
VERIZON	2.873	1.070