Dynamic Bargaining and Size Effects in the Broadband Industry

Daniel Goetz *

December 12, 2019

Abstract

I estimate a model of dynamic bargaining between internet service providers (ISPs) and Netflix over interconnection fees and use it to evaluate counterfactual ISP mergers. I show that the size of the downstream consumer market an ISP serves is much more important than its disagreement point for determining bargaining outcomes with Netflix. I evaluate several mergers between ISPs serving non-overlapping markets and find that smaller ISP mergers would lengthen negotiations and reduce Netflix's share of the bargaining surplus, while larger mergers would have no significant effect.

Keywords: Mergers, Broadband Internet, Bargaining
JEL Classification: L41, L96, C73

*Thanks to Jan de Loecker, Bo Honoré, Jakub Kastl, Myrto Kalouptsidi, Eduardo Morales, Nick Buchholz, Sharon Traiberman, Octavian Carare, and Paul LaFontaine, as well as commenters from the Princeton IO student workshop, the Rotman School of Management, NYU, UCalgary, UBC Sauder, York U, Analysis Group, the CEA, CIREQ, the FCC, and UW-M. Financial support from the NET Institute (www.netinst.org) is gratefully acknowledged. Correspondence: Rotman School of Management, University of Toronto, Toronto, ON M5S3E6. Tel.: (647) 542 4505. Email: daniel.goetz@rotman.utoronto.ca. Web: http://www.dtgoetz.com

1 Introduction

The relative size of negotiating parties is important for determining outcomes in firm-to-firm bargaining. In negotiations between upstream and downstream firms, downstream firms that serve larger markets tend to extract a greater share of the bargaining surplus from the upstream firm.[1] However, downstream firms' outside options—how much profit they lose per customer if there is disagreement—has typically been the focal mechanism for explaining bargaining outcomes, leaving the effect of size in a residual. Since mergers between downstream firms can affect both the size of the market served as well as the downstream firm's outside option in bargaining, understanding how size affects bargaining in different contexts is important for both managers and competition policy.

In this paper, I build and estimate a structural model to quantify how the size and outside options of U.S. internet service providers (ISPs) affect bargaining outcomes with Netflix, a large provider of streaming video. The setting is well suited to the exercise: rich data on which ISPs can serve which residences allows for measurement of ISP size, household choice sets, and thus ISP outside options in bargaining. Moreover, ISP mergers in this industry are typically between firms serving non-overlapping markets, so that size is potentially the only channel through which consolidation would affect upstream bargaining.[2] With exponential growth in internet traffic, understanding how ISP size affects the division of upstream surplus is especially crucial for incentivizing investment.[3]

The main contribution of the paper is to show that ISP size substantially outweighs ISP outside options in determining bargaining outcomes, and to estimate the precise structural mechanism size acts through. In particular, I show that Netflix's bargaining surplus per household is convex in an ISP's size. Netflix thus has more to lose from disagreement with large ISPs, implying a weaker bargaining position and a smaller share of the bargaining surplus. The result is robust to

[1]See Gowrinsankaran, Nevo, and Town (2015) for hospital-insurer negotiation, Grennan (2013) for hospital negotiations with stent suppliers, and Crawford and Yurukoglu (2012) for negotiation between cable TV providers and channels. Dafny, Ho, and Lee (2016) and Collard-Wexler, Gowrisankaran, and Lee (2019) provide more hospital-insurer negotiation evidence.

[2]See Charter and Time Warner Cable (TWC), and CenturyLink and Qwest, for two recent large examples.

[3]Cisco (2017).

controls for ISP bargaining power a là the Nash bargaining residual. The form of the convexity is such that mergers between small ISPs lead to dramatic shifts in retained surplus while mergers between large ISPs do not, in contrast to approaches that average the Nash bargaining residuals of merging firms (Gowrinsankaran, Nevo, and Town, 2015).

The key data I observe is the duration of negotiations between Netflix and ISPs over interconnection fees. Starting in 2013, Netflix video stream quality decreased substantially at a number of ISPs while the parties negotiated. Successful bargaining marked a return to higher quality at that ISP, a cost savings for Netflix as it installed equipment in that ISP's network, and a lump sum transfer to the ISP.

Since I only observe bargaining durations in the data, I estimate a new structural model of households, ISPs and Netflix that allows for dynamic bargaining. On the demand side, households select internet plans and ISPs based on their observed choice sets, and on whether Netflix quality is degraded at a particular ISP or not. Exogenous variation in how ISPs upgrade plan speeds on their networks allows me to estimate price sensitivity. For each ISP and quarter, the model predicts how an ISP's subscriber profit depends on its negotiation status with Netflix; on average, ISPs are estimated to lose $2.2 million in profits per quarter of disagreement with Netflix.

The supply side uses a screening-style framework to map household demand and ISP characteristics into bargaining lengths and surplus splits. The base model explains variation in bargaining length through two mechanisms. First, if ISPs lose more profit during disagreement periods—i.e., their outside option is worse—they make more generous offers to Netflix and conclude agreement faster. Second, ISP characteristics—including network size—affect Netflix's unobserved surplus distribution, changing the ISP's sequence of optimal offers. A fixed cost parameter has the potential to shift the distribution left, implying that small ISPs may want to bargain more quickly since there is relatively less surplus to split. Size effects are not assumed, but estimated: if variation in ISP outside options alone can explain the observed pattern of bargaining in the data, then the recovered fixed cost will be small in magnitude and insignificant.

The data and structural estimation reveal that ISP size plays a significant role in determining bargaining outcomes. I estimate a fixed per-ISP cost of $18.3 million to Netflix of installing its cost-saving equipment, and a surplus increase of $8.17 per household. Removing the fixed cost and re-simulating bargaining durations leads the model to substantially overpredict (underpredict) bargaining length for small (large) ISPs. In a robustness check, I adapt the result from Collard-Wexler, Gowrisankaran, and Lee (2019) that relative patience can map to Nash bargaining power, and estimate a substantial fixed cost even after endowing the largest ISPs with the higher discount rates suggested by their financial balance sheets.

The fact that estimated bargaining surplus per household is convex in ISP size has important implications for counterfactual mergers. Since policy-relevant ISP mergers are between firms serving non-overlapping geographic markets, ISP outside options do not change after a merger. However, the size effect will generally imply that Netflix retains a lower share of a larger surplus when bargaining with the merged firm, since there is more surplus on the table but Netflix's bargaining position weakens. For Comcast and Time Warner Cable (TWC)—the largest cable internet providers at the time—the merger would lead to a negligible change in surplus and bargaining time, since the ISPs are so large that the non-linearities in surplus are minor. For a merger between smaller ISPs, a longer bargaining period will harm consumers, and the delayed agreement combined with a lower surplus share for Netflix will imply that Netflix may actually retain less surplus in absolute terms despite the cost efficiencies, weakening its incentive to invest.

Contribution to the Literature

My paper contributes to a growing literature on using bargaining models to assess counterfactual market structures, starting with Ho (2009) and Crawford and Yurukoglu (2012). Many of these papers—including Grennan (2013), Gowrisankaran, Nevo, and Town (2015), and Dafny, Ho, and Lee (2016)—rely on the Nash-in-Nash multilateral bargaining framework, which models bargaining outcomes as a function of agreement points, outside options generated from demand systems, and bargaining residuals, although Ho (2009) also models bargaining as take-it-or-leave-it offers as in this paper. Since Nash-in-Nash is scale invariant, large firms with better outcomes tend to have large bargaining residuals. This paper provides a more structural justification for the effect of size, which allows for consistent merger counterfactuals as opposed to making assumptions on how residuals combine when firms merge.

My paper also builds on a smaller literature devoted to empirically analyzing dynamic bargaining. Ambrus, Chaney, and Salitsky (2016) structurally estimate a dynamic bargaining game with one-sided offers and incomplete information, and I extend their framework to a setting that endogenizes whether an offer is made, and where the offering party receives outside payments.[4] A complementary complete information framework for multilateral dynamic bargaining is developed in Merlo and Wilson (1995) and Merlo and Tang (2012).

A few papers also examine how mergers affect surplus division through size effects. In particular, Chipty and Snyder (1999) analyze how cable system size affects the price paid to content providers. While they find that intermediaries have a stronger bargaining position if they do not merge, their analysis is reduced form and does not account for the size of intermediaries, only their number.[5] Dafny, Ho, and Lee (2016) show how non-overlapping mergers on the content (hospital) side can affect welfare if the content is provided through the same intermediary, which does not rely on size effects but which provides a complement to this paper's results on mergers between intermediaries serving non-overlapping markets.

Lastly, I contribute to the much broader literature on welfare analysis of mergers by highlighting new sources of welfare tradeoffs; in particular, that cost efficiencies on the content side can imply mergers harm consumers even without changes in market power, by increasing the chance of bargaining breakdowns and content degradation or foreclosure. This complements work such as Argentesi and Filistrucchi (2007); Chandra and Collard-Wexler (2009) who study the consequences of intermediary market power—but not size—on exit and entry of content providers.

[4] Multilateral bargaining also resembles joint optimal stopping problems, notable examples of which include Berry and Tamer (2006),Honoré and de Paula (2010) and Björkegren (2015). I draw on insights from this literature, especially with regards to how payoff interactions affect equilibrium existence and uniqueness.

[5] They find that the total sum of content providers' surplus with each intermediary is *convex* in the number of intermediaries, so that intermediaries do better to stay unmerged to keep surplus high and their bargaining position strong. In contrast, in my framework, due to the fixed cost the content provider's total surplus is *concave* in the number of intermediaries, as more intermediaries necessitates more fixed cost expenditures to reach the same market, and thus I conclude that mergers improve intermediaries' bargaining position.

2 Background

2.1 Residential Broadband and Netflix

The internet is a two-sided market. On one side are consumers, who purchase access to the internet in order to consume online services and content such as email and streaming video. On the other side are the providers of services and content, such as Google and Netflix, who charge consumers either indirectly, via advertisements, or directly, via subscription fees, for using services and viewing content. In the middle are layers of firms that intermediate the relationship between consumers and content providers. In what is to follow I refer to the service/content side of the market as content providers.

Consumers and small businesses interact with "last-mile" or "edge" internet service providers like Comcast and Verizon. A consumer's choice set for wired internet service depends on which ISPs have infrastructure connected to her house, since last-mile ISPs generally have the exclusive right to sell service on infrastructure they own. Service is differentiated by infrastructure technology (cable, fiber optic, etc.) across providers, and by tiered menus of plans varying by monthly price and download speed in megabits per second (MB/s) within providers.[6] By 2013, 70% of households had access to two or more wired providers offering maximum download speeds greater than 10MB/s. However, the industry is concentrated: for 91% of those consumers, at least one alternative was provided by the four largest last-mile ISPs: AT&T, Comcast, TWC, and Verizon.

Netflix and other large content providers seek to connect to last-mile ISPs, and have several options to do so. The largest, like Google or Microsoft, incur a large fixed cost to install infrastructure that allows them to connect directly with last-mile ISPs at low variable cost. Others buy access from "transit" ISPs like Level3 and Cogent, who connect with last-mile providers to transmit content to consumers. Using third parties to transmit content comes with a higher variable cost,

[6]Different plans do not provide access to different content: all plans have access to all content. Upload speed in MB/s, caps on how much content can be consumed in a month, and contract length are also plan features, but these are much less important: 92.5% of respondents in the 2013 Current Population Survey Internet Supplement list price, download speed, or reliability as the most important feature of service, from a list of choices that also includes upload speed, data usage caps, mobility, and bundling options.

and content providers must ensure they purchase sufficient access to meet consumer demand. To bypass purchasing enough transit access to meet demand at peak times, content companies can also pay to upload content to so-called "content delivery networks" (CDNs)—geographically distributed caches of servers in last-mile networks that ensure no consumer is far from a content source.

Netflix Bargaining Event

Starting in mid-2012, Netflix developed a strategy to transition from using mainly third parties to disseminate content, to using its own infrastructure. They developed a custom CDN, called Open Connect, and in so doing incurred a large fixed development and deployment cost.[7] Open Connect would save Netflix money in two ways. First, it would allow them to save on the variable cost of using third party CDNs. As the largest online video distributor, Netflix not only paid transit ISPs for connections and the CDNs for servers, but also pursued a policy of paying the fees that last-mile ISPs charged CDNs and transit ISPs carrying Netflix content.[8] 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.[9] Open Connect would allow Netflix to deliver service reliably and at lower cost.

[Figure 1 about here]

By mid-2013, Netflix had not installed Open Connect in the vast majority of last-mile ISP net-

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

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

[9] Netflix 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."

works, and had begun to report degraded quality of service to a number of U.S. ISPs.[10] I emphasize the quality degradation for the largest two U.S. ISPs by subscriber count, AT&T and Comcast, who in 2013 collectively accounted for 43% of all U.S. broadband subscribers, in Figure 1. Starting in mid-2013, the average transmission rate of Netflix data to subscribers at these ISPs dips far below trend, and is restored after varying amounts of time. ISPs including TWC (13% of subscribers) and Verizon (10.5%) also experience degradation, while Cox (5.5%) and Cablevision (3.3%) do not. I argue that these slowdowns and their resolutions correspond to periods of bargaining disagreement over the negotiated fees for installation of Open Connect. In Figure 1 Comcast service quality is fully restored during the first quarter of 2014, which corresponds to Netflix FCC filings indicating that by January, 2014, Netflix and Comcast had reached a deal on interconnection fees.[11] AT&T service quality is only restored later: in Netflix's April 2014 Q-10 filing, they state that AT&T still has not agreed to Open Connect interconnection,[12] but data from the Center for Applied Internet Data Analysis (CAIDA) indicates that AT&T began interconnecting with Netflix in August 2014—around the time AT&T service quality is restored. When describing the event in FCC filings at the end of 2014, Netflix notes that "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.[13]

3 Data

Data on bargaining delays, shares, prices and consumer ISP choice sets are described below.[14] Cross-sectional summary statistics are provided in Table 1 and patterns over time in Figure 2.

[Table 1 about here]

[10]Responsibility is difficult to determine. Cremer, Rey, and Tirole (2000) note that in theory, interconnection quality is determined by whoever values it the least.

[11]Petition to Deny, pg. 57, paragraph 2 — pg. 58 paragraph 2.

[12]Netflix 2014Q1 letter to investors, pg. 5 paragraph 3.

[13]Petition to deny, pg. 49, paragraph 2.

[14]More details about data availability and construction are provided in Appendix B.

[Figure 2 about here]

Bargaining Delays: Durations are derived from business filings and internet backbone monitoring reports for 49 ISPs. Bargaining begins simultaneously for all U.S. ISPs in 2013Q3, and ends either in that quarter ($t = 1$, no slowdown) or by 2014Q3 ($t = 5$, one year of slowdown).

Market shares: National subscriber data is gathered from ISPs' publicly available quarterly and yearly earnings reports (10-Q and 10-K) which provide numbers of internet plan subscribers, supplemented by data from Leichtman Research Group for privately held companies. I have quarterly data on 30 ISPs' subscribers from 2010Q1 to 2014Q4 inclusive. Some of these ISPs are purchased or merged before the beginning of bargaining, generating choice set variation for consumers. Total market size is recovered from yearly counts of the total number of occupied dwellings in the United States, from the American Community Survey.

Plan characteristics: I recover quarterly prices and download speed for all plans at the national level offered by 61 ISPs for 2010Q1 to 2014Q4 from a variety of sources, including the FCC Urban Rate Survey and Open Connectivity Database (see Appendix B.) The data includes the 30 ISPs for which I have share data, as well as all 49 ISPs in the bargaining delay data. ISPs add or drop plans from their menu across different regions depending on their infrastructure, but conditional on offering a plan it is advertised at the same price everywhere during the sample period.

Choice sets: Data on what ISP and plan choices are available to consumers comes from the U.S. National Broadband Map (NBBM). For each census block, typically covering several hundred housing units, I observe which ISPs provide service to that block and their maximum advertised speed at half-yearly intervals from 2010-2014. After cleaning, I have quarterly-level choice sets for 57 ISPs from 2010Q1 to 2014Q4, including the 49 bargaining ISPs, as well as ISPs for which I have market shares but which are merged before bargaining begins. I assume that all consumers have access to satellite internet as part of their choice set.

Plan microdata: The FCC's Measuring Broadband America (MBBA) program provides hourly internet quality testing data for an unbalanced panel of roughly 10 000 consumers, from which I recover quarterly, household-level ISP and plan choices for 17 ISPs. ISP participation is voluntary, but the FCC chooses the within-ISP sample of consumers to be representative of the national distribution of consumers across that ISP's plans.

3.1 Reduced Form Correlations

ISP bargaining patterns

To document how bargaining varies across ISPs, I run simple cross-sectional regressions of agreement times ($t \in (1 \ldots 5)$ on covariates, using the sample of 49 ISPs for which I have durations data and choice set data:

$$\text{Agreement Time}_f = \beta_1 \text{Market Size}_f + \beta_2 \text{Num. Competitors}_f + \boldsymbol{\gamma} X_f + \epsilon_f.$$

I emphasize two covariates: first, the number of homes passed by an ISP's infrastructure, which captures an ISP's market size; second, the average number of competitors an ISP faces in the markets it serves, which captures whether an ISP can expect to lose many subscribers during a slowdown and hence, its outside option in bargaining. I also include other features of ISPs' networks including their technology (DSL/Fiber, Cable).

[Table 2 about here]

Results are presented in Table 2.[15] These regressions imply that both size and outside options affect bargaining outcomes, with size potentially playing the stronger role.

However, the number of competitors does not fully capture an ISP's outside option in bargaining. To correctly control for outside options' effect on bargaining durations requires estimating the demand response to a slowdown at an ISP, and having that response enter the ISP's optimal bargaining problem. In the next section I examine how households at affected ISPs reacted to the

[15] Results are robust to dropping AT&T. Using subscribers to proxy for market size limits the number of observations, but results on market size remain significant.

slowdown to motivate the structural demand model that will generate subscriber responses in Section 4.

Costs of bargaining for ISPs

The negative correlation between the number of competitors and agreement times in Table 2 implies that ISPs in more competitive markets may find longer bargaining more costly. Longer periods of disagreement will only be costly for ISPs if Netflix quality reductions cost an ISP subscribers, or induce an ISP to shrink its margins. Here I assess whether ISPs lost subscribers while negotiating with Netflix, and if they reduced their prices.

I run the following panel regression on the 30 ISPs for which I have subscriber data, using data from the beginning of 2010 through the end of 2014 at a quarterly frequency:

$$\Delta \log(\text{Subscribers}_{ft}) = \delta_1 \Delta \log p_{ft} + \delta_2 \Delta \log q_{ft} + \bar{\alpha}_f + \bar{\alpha}_t + \Delta\epsilon_{ft}. \tag{1}$$

Where p_{ft} and q_{ft} are the price and download speed of ISP f's entry level plan at time t.[16] I also drop ISP-quarters for which there is a merger or takeover, all of which occur entirely prior to 2013Q2. The estimated residual $\Delta\hat{\epsilon}_{ft}$ is ISP f's residual growth rate in period t, and is expected to be lower during the slowdown for ISPs that took time to bargain with Netflix.

Figure 3 presents the residual subscriber growth rates for the median ISP, with groups split by whether there was a bargaining delay with Netflix or not. The first vertical line indicates the beginning of bargaining, the dashed vertical line indicates the conclusion of the first wave of bargaining in early 2014, and the final line indicates when the last ISPs agreed. The median ISP that experiences a bargaining-induced quality degradation grows 0.5 percentage points slower during the slowdown than ISPs that concluded bargaining immediately, suggesting an economically significant role for subscriber substitution between ISPs in response to Netflix quality degradation.

Diff-in-diff regressions using shares as the dependent variable are presented as a baseline alongside structural demand in Section 4. Diff-in-diff regressions on the household-level MBBA subscription panel confirm that users are more likely to switch away from affected ISPs during

[16]The results are robust to alternative measures of ISP plan prices.

the slowdown, see Table C.1 in Appendix C.

[Figure 3 about here]

[Figure 4 about here]

[Figure 5 about here]

Two pieces of supplementary evidence support the claim that marginal households reacted to Netflix slowdowns by switching ISPs. First, Netflix has made ISP-specific streaming quality public and accessible since 2012; consumers at affected ISPs therefore had the information necessary to judge whether an ISP with better service existed in their choice set. Second, households in 2013 who owned a TV-based internet streaming device (e.g., AppleTV) were 36% more likely to switch ISPs than households without such a device.[17] Consumers who switched and owned a TV-based streaming device were also 19% more likely to report that reliability was their reason for switching relative to the average switcher.

Do ISPs also reduce their margins during the slowdown? To assess whether and how much ISPs reprice, I run the following regression:

$$\Delta \log p_{ft} = \delta_1 \Delta \log q_{ft} + \bar{\alpha}_f + \bar{\alpha}_t + \Delta\epsilon_{ft} \tag{2}$$

Figure 4 plots the median residual growth rates in price again split by whether the ISP took time to bargain with Netflix or not. I find that prices do not shrink or grow differentially for ISPs affected by the slowdown.

Not only does the marginal household value Netflix streaming quality, they value it enough to leave ISPs—or alternatively, not sign up with ISPs—whose streaming quality is degraded. However, ISPs do not appear to be changing their posted plan prices during this time. The structural

[17] Data from the Current Population Survey 2013 Internet Use Supplement, available at https://cps.ipums.org/cps/. 18.8% of households in the supplement reported having switched ISPs in the prior 3 years.

model will thus emphasize subscriber substitution and not changing markups as the cost of the slowdown for ISPs.

Netflix subscriber substitution

I argue that Netflix neither lost subscribers nor reduced its margin per subscriber during the slowdown, which will motivate modelling Netflix demand as inelastic in the structural model. The available data is the total number of Netflix quarterly subscribers, as well as the number of quarterly subscribers currently on a free trial (4.7% of total subscribers on average.) Netflix did not change its pricing during the slowdown, but there could have been a reduction in average margins if more free trials were offered.

Figure 5 shows the raw (non-residualized) growth rates in total and unpaid subscribers from 2010Q1 through 2015Q4. There is a strong seasonality in growth rates and a trend of slowing growth, but growth rates in total subscribers in the second half of 2013 largely resemble those in 2012. The growth rate in unpaid subscribers is higher at the end of 2013 compared to 2012, but does not spike substantially, suggesting that Netflix is not making more free trials available to compensate for reduced demand.

The weak evidence for substitution away from Netflix agrees with the ISP subscriber growth results, in that it points to a high consumer valuation for Netflix. Inelastic demand for streaming video may come from the lack of similarly priced alternatives for on-demand TV and movies,[18], from strong consumer sentiment that ISPs such as Comcast provide poor service[19], from the fact that Netflix users were high-valuation early adopters, or from widespread account sharing. I will assume that demand for Netflix is inelastic in the structural demand model, and that slowdowns affect consumer valuations for ISPs.

[18]By the beginning of 2013, Blockbuster's owner, DISH, had shut down 1100 of 1500 stores, and shuttered 1450 of 1500 by 2015. A monthly Netflix subscription granting unlimited streaming was $7.99 per month in 2013, while pay-per-view movies were anywhere from $2.99 to $5.99 for a one week rental.

[19]Consumer Reports National Research Center

4 Demand

4.1 Model

A household i in market m at time t chooses among internet service providers f in that market. Let $\mathcal{F}_{mt}$ index the set of firms in market m at time t.

Each firm f offers a menu of vertically differentiated plans $j \in \mathcal{J}_{fmt}$, which vary by market and over time. Conditional on choosing firm f, a household chooses between the available j offered by that f, enjoying the the following indirect utility:

$$U_{ifmt} = \delta_{ft} + \lambda_{ifmt} + \epsilon_{ifmt}, \tag{3}$$

where δ_{ft} is the mean utility each household derives from consuming f at time t and λ_{ifmt} is the household-specific value, which incorporates the household's plan choice. I assume that ϵ_{ifmt} is distributed type I extreme value and utility from choosing no internet (the outside option) is normalized to zero, which implies the following form for aggregate market shares:[20]

$$s_{ft} = \sum_{m|f \in \mathcal{F}_{mt}} w_{mt} \int \frac{\exp(\delta_{ft} + \lambda_{ifmt})}{1 + \sum_{f \in \mathcal{F}_{mt}} \exp(\delta_{ft} + \lambda_{ifmt})} \partial F(i)_{mt}, \tag{4}$$

where w_{mt} describes the weight of market m at time t, such that $\sum_m w_{mt} = 1$ for each t. I aggregate market-level shares to county-level since I only have national subscriber numbers. The model is thus similar to Goeree (2008) in that I do not observe product shares at the choice-set level; however, here the choice set distribution is observed in the data, while in that paper it is estimated.

[20] Although the model allows consumers to cancel internet completely in response to the Netflix slowdown, inelastic Netflix demand can be rationalized by account sharing (some sharer still has access to high quality) or by sourcing internet outside the home.

4.2 Specification

For mean utility, I assume the following functional form:

$$\delta_{ft} = \gamma \text{Disagree}_{ft} \times (t - t_0) + \bar{\alpha}_f + \bar{\alpha}_t + \bar{\alpha}_f^t ISP_f \times t + \xi_{ft}. \quad (5)$$

ISP and quarter fixed effects are included, as is an ISP-specific linear time trend to capture slow moving trends in ISP market shares.[21] Disagree_{ft} is a dummy indicating whether ISP f is still negotiating with Netflix at time t; γ should therefore be negative if the reduction in Netflix quality is affecting consumer utility. Interacting with time since 2013Q2 (t_0)captures the Figure 1 dynamics of an escalating slowdown in a reduced form way, but results are robust to removing the interaction. ξ_{ft} is an ISP-specific demand shifter that the ISP observes before setting prices, but which the econometrician does not observe, reflecting, for instance, shocks to bundled services.

Note that Disagree_{ft} does not have an m subscript, as I assume the bargaining between an ISP and Netflix affects that ISP's customers nationwide. Although there is some evidence that the slowdown had different effects across markets within the same ISP, there is no available data to identify this geographic heterogeneity.

Since firms are not uniquely identified with a price or download speed, these variables do not enter δ_{ft}. Households choose among f's menu of offered plans, and the effect on indirect utility is captured by λ_{ifmt}:[22]

$$\lambda_{ifmt} \equiv \max_{j \in \mathcal{J}_{fmt}} \left\{ -\exp(\alpha^p + \nu_i \sigma^p) \log(p_{jfmt}) + \exp(\alpha^q) \log(q_{jfmt}) \right\}, \quad (6)$$

where p_{jfmt} and q_{jfmt} are the price and download speed of firm f's plan j in market m at time t, respectively.[23] Heterogeneity in price disutility is captured by $\nu_i \sim \mathcal{N}(0, 1)$; the form of the price

[21] I allow an ISP's fixed effect to adjust after a merger.

[22] It is not possible for a consumer to reduce the disutility from the slowdown by upgrading their plan here, since Disagree_{ft} is not interacted with plan quality. With this assumption, ISPs will not want to prolong the slowdown to induce upgrades. I show using the MBBA microdata that there is no evidence of faster upgrading at affected ISPs in Table C.1, Appendix C.

[23] Logs are taken to remain consistent with the reduced form, and to account for wealth effects as in Berry, Levinsohn, and Pakes (1995); monthly internet prices in the U.S. can approach the cost of a car lease payment.

coefficient follows Berry, Levinsohn, and Pakes (2004). I also estimate a baseline specification where I do not account for plan choice, and households all earn the same average utility from f's plans if f is chosen; i.e. $\lambda_{ifmt} \equiv \tilde{\alpha}^p \overline{\log(p_{jfmt})} + \tilde{\alpha}^q \overline{\log(q_{jfmt})}$. In the baseline the price coefficient is unconstrained (i.e., it can be positive) to check the performance of the instruments in estimation.

4.3 Estimation and Identification

The parameters to estimate include the linear parameters θ^l in Equation 5 and the nonlinear parameters θ^{nl} in Equation 6, using observed data on aggregate shares, plan prices and download speeds, plan shares for a subset of ISPs, and the distribution of choice sets. I briefly outline the estimation procedure here.

For a guess of θ^{nl}, I recover estimates of mean utilities $\hat{\delta}_{ft}(\theta^{nl})$ via a modified Berry, Levinsohn, and Pakes (1995) inversion and use them to concentrate out the linear parameters in Equation 5 as in Nevo (2000). The residuals $\hat{\xi}_{ft}(\theta^{nl})$ from this equation are then interacted with a set of instruments Z_{ft} to form the sample analogue of moments $E[\xi_{ft}\boldsymbol{Z}_{ft}] = 0$ in GMM. This stage uses the bargaining data, subscriber data, choice set data, and price and plan feature data. To help identify the non-linear parameters, I follow Scherbakov (2015) in forming an additional set of moments based on conditional plan shares, $E[(s_{j|ft} - \hat{s}_{j|ft}(\theta^{nl}))\boldsymbol{Z}_{ft}] = 0$, which provide information on individuals' valuation of price and download speeds.[24] Data on $s_{j|ft}$ comes from the household choice panel data.

Since the unobserved component of mean utility ξ_{ft} is observed by a firm before setting prices it is endogenous to p_{jft}. The instruments Z_{ft} must therefore be uncorrelated with the transient demand shocks ξ_{ft}, but correlated with prices. While prices and menus of speeds can be flexibly chosen, an ISP's maximum offered download speed in an area is constrained by its installed technology, and upgrades to this maximum happen on a slow timeline and not based on quarterly shocks (see Figure 2). I therefore use contemporaneous and lagged values of functions of the maximum offered speeds of an ISP's competitors as instruments for price, as well as the number

[24]Information on the GMM estimation is provided in Appendix D.2.

of competitors an ISP faces in the markets it serves. Exogenous improvements in a competitor's maximum available download speed shift an ISP's markups in a way that is orthogonal to contemporaneous and transient demand shocks, and will therefore identify price disutility.

4.4 Results

The final column of Table 3 provides the main specification, which controls for household-specific choice sets, households' choice of plans within an ISP, and instruments for prices with exogenous network upgrading of competitors. I present results from intermediate models in Table 3 to demonstrate the value-added of these modelling choices.

[Table 3 about here]

Controlling for choice sets substantially increases the magnitude and precision of the estimated γ. Intuitively, some households that experience disutility from slower Netflix simply cannot substitute due to constrained choice sets; since the first model assumes they can always substitute, it does not discount the lack of substitution among constrained individuals and hence infers only minimal disutility from the slowdown. From the second to the third model, the price instruments increase the magnitude of the price coefficient—combined with a first stage F-statistic of 14.19 this gives confidence that the instruments are working as expected. Allowing consumers to actively substitute across plans in response to price changes in the fourth model, the average price elasticity increases to 5.9, in line with estimates for cable television (Crawford and Yurukoglu, 2012).

The demand results are the key input to the bargaining game, as firms will weigh the benefits of making a high fee offer to Netflix against the lost profit from consumers substituting away during the ensuing slowdown if the offer is rejected. Using the estimated parameters, I plot the percent subscriber loss from disagreement in 2013Q3 against the disagreement length in quarters, for the 49 ISPs in Netflix's data for which I also observe choice data and which could have the CDN installed (satellite providers are excluded.)[25]

[25]Details of the imputation of mean utility for ISPs not in the estimating sample are provided in Appendix D.

The positive relationship in Figure 6 implies that ISPs that stood to lose more subscribers if Netflix quality was reduced also bargained more quickly. However, many of the ISPs that had no period of slowdown are also fairly small; given the fixed cost to Netflix of setting up interconnection it is possible that small ISPs did not bargain to extract more of Netflix's surplus simply because they anticipated that there was not much surplus to extract. The bargaining model developed in Section 5 will disentangle the contribution of size and outside options, and allow for policy counterfactuals.

[Figure 6 about here]

4.5 Marginal Costs

Since firms care about lost profits and not lost revenue, the marginal costs of delivering an internet plan are necessary to accurately recover the firm's bargaining tradeoff. The largest cost in this industry is the fixed cost of actually laying the wired infrastructure; however, there are also variable costs that scale with the number of served households such as adding customer support staff, network maintenance, and—for smaller ISPs—paying for interconnection with backbone and Tier 1 carriers. These costs vary by ISP, but not by plan.

I first recover a weighted average ISP-specific demand elasticity within a quarter. Using an ISP's lowest priced plan, I invert this ISP-level demand elasticity to find the marginal cost of that plan, and assume that is the marginal cost to the ISP of delivering service for all of its plans in that quarter.[26]

With the marginal costs in hand, I can recover the profit of each ISP in each quarter as a function of time and the vector of disagreements. Define $a_{ft} = 1 - \text{Disagree}_{ft}$ and $\boldsymbol{a}_t \equiv (a_{1t}, \ldots, a_{Ft})$, so $\boldsymbol{a}_t$ gives the vector of agreement at time t. Then quarterly profits for f as a function of $\boldsymbol{a}_t$ are

$$\pi_{ft}(\boldsymbol{a}_t) = 3 \cdot \sum_{j \in \mathcal{J}_{ft}} s_{jft}(\boldsymbol{a}_t)(p_{jft} - mc_{ft}). \tag{7}$$

[26] A list of ISP demand elasticities (averaged across quarters), marginal costs and average margins are provided in Appendix D.3.

where we multiply by 3 to reflect that prices are monthly, but time periods in the model are quarterly.

5 Bargaining Model

In this section I develop a model of dynamic bargaining that can rationalize the observed delays in the data. The model is necessary to separate the effects of size and outside options, and for assessing delays under counterfactual market structures.

5.1 Model

Upstream bargaining is a dynamic game played between all ISPs indexed by $f = 1, \dots, F$, and the upstream content provider, Netflix. Time is discrete and runs from t_0 to a terminal period T.

At time t_0, Netflix may pay an upfront R&D cost to draw a random vector of ISP-specific surpluses $\boldsymbol{\mu} \equiv (\mu_1, \dots, \mu_F)$, which only Netflix observes. The vector $\boldsymbol{\mu}$ corresponds to Netflix's ISP-specific cost-savings from installing their CDN servers in that ISP's network, relative to the status quo of continuing to use third-party transmission; randomness reflects that this is a new (to Netflix), custom technology whose payoff may not be observable until after R&D.

μ_f is drawn from distribution $G(\mu | w_f, \theta^s)$, where w_f are ISP observables and θ^s a vector of parameters. The dependence on observables implies, for example, that larger ISPs may be associated with higher Netflix surplus on average. I assume that conditional on observables, the μ_f are independent.

The bargaining protocol consists of repeated offers of fees made each period by the uninformed party, the ISPs, which Netflix may accept or reject. Until an ISP's offer is accepted, Netflix quality at that ISP will remain low.

Actions and Timing

Starting from t_0, within each period t:

1. Any ISP f whose prior offers have not been accepted observes its own vector of mean-zero

idiosyncratic bargaining costs $\boldsymbol{\epsilon}_{ft}$. f chooses a lump sum interconnection fee $\tau_{ft} \geq 0$, or else to not make an offer that period, $\tau_{ft} = \varnothing$.

2. Netflix accepts or rejects each offer. If Netflix accepts f's offer, it pays f τ_{ft} and realizes surplus μ_f, retaining $\mu_f - \tau_{ft}$. Quality is immediately restored at that ISP.
3. Each f observes whether offers were made and which offers Netflix accepted, and competes for consumers. f receives flow payoffs $\pi_{ft}(\boldsymbol{a}_t)$ that depend on the acceptance vector, earns τ_{ft} if their offer is accepted, and update their beliefs about $\boldsymbol{\mu}$ depending on $\boldsymbol{a}_t$.

Substantively, the model is a screening model with one-sided asymmetric information where the uninformed party is making offers. Delay will arise if the idiosyncratic bargaining cost implies that it is not optimal for the ISP to make an offer that period, or if Netflix rejects an offer, anticipating that the likelihood of a more generous offer next period outweighs the cost of waiting. Note that because the μ_f are conditionally independent, f does not learn about its own surplus distribution from the rejected offers of other ISPs. ISP interaction therefore comes through the dependence of the flow payoffs on $\boldsymbol{a}_t$.

ISP's Problem

The ISP maximizes expected discounted profits conditional on its beliefs about $\boldsymbol{\mu}$, its bargaining costs $\boldsymbol{\epsilon}_{ft}$, other ISPs' strategies, and Netflix's optimal dynamic acceptance strategy. Define $\mathcal{B}_{ft}$ to be f's beliefs about the distribution of $\boldsymbol{\mu}$ at time t. Then f's value function can be written as:

$$V_{ft}(\mathcal{B}_{ft}, \boldsymbol{\epsilon}_{ft}) = \max \Big\{ \max_{\tau_{ft}} \mathbf{E}_{\boldsymbol{\tau}_{-ft}, \boldsymbol{a}_t} \left[a_{ft}\tau + \pi_{ft}(\boldsymbol{a}_t) + \beta V_{ft+1}(\mathcal{B}_{ft+1}, \boldsymbol{\epsilon}_{ft+1}) \right] + \epsilon_{1ft}, \quad \mathbf{E}_{\tau_{-ft}, \boldsymbol{a}_{-ft}} \left[\pi_{ft}(0, \boldsymbol{a}_{-ft}) + \beta V_{ft+1}(\mathcal{B}_{ft+1}, \boldsymbol{\epsilon}_{ft+1}) \right] + \epsilon_{0ft} \Big\} \tag{8}$$

I index value functions by f to account for heterogeneity in π_{ft} across ISPs, and by t since, with a terminal period T, the problem is non-stationary.

The value function does not depend on opponents' beliefs $\mathcal{B}_{-ft}$, because all beliefs about $\boldsymbol{\mu}$ are symmetric if we assume pure strategies. This follows because each ISP observes whether an offer was made by each opponent; conditional on strategies, an ISP can therefore infer what offers were made, and therefore what its opponents know about their own μ_f. Thus while Netflix has asymmetric information with respect to ISPs, ISPs are symmetrically uninformed. I therefore write $\mathcal{B}_t = \mathcal{B}_{ft}$.

f does not observe other ISPs' idiosyncratic draws $\boldsymbol{\epsilon}_{-ft}$ when it makes an offer, so it must take an expectation over $\boldsymbol{\tau}_{-ft}$. This expectation is only over the probability of making an offer since, given symmetric information, conditional on a strategy f knows what offers $-f$ will make if it does make an offer. f does not know the exact μ_f, requiring it to take an expectation over the probability of a realization of $\boldsymbol{a}_t$ given its offer, which will depend on f's beliefs, other ISPs' offers, and Netflix's strategy.

Netflix's Problem

After initiating bargaining at t_0, Netflix chooses whether to accept or reject f's offers in each period. Consistent with the reduced form evidence, I assume that Netflix subscriptions are inelastic, so that Netflix is only bargaining over its share of the surplus. Netflix period profits are therefore:

$$\pi_{xt}(\boldsymbol{a}_t, \boldsymbol{\tau}_t) = \bar{\pi}_{xt} + \sum_f (\mu_f - \tau_{ft}) a_{ft}, \tag{9}$$

where $\mu_{ft} - \tau_{ft}$ is only earned in the period of agreement where τ_{ft} is paid. Given the ISPs' beliefs and strategies, after observing $\boldsymbol{\tau}$ Netflix's problem can be written,

$$V_{xt}(\boldsymbol{\tau}_t, \mathcal{B}_t) = \max_{\boldsymbol{a}_t} \pi_{xt}(\boldsymbol{a}_t, \boldsymbol{\tau}_t) + \beta \mathbf{E}_{\boldsymbol{\tau}_{t+1}} \left[V_{xt+1}(\boldsymbol{\tau}_{t+1}, \mathcal{B}_{t+1}) | \boldsymbol{\tau}_t, \mathcal{B}_t \right] \tag{10}$$

where next period's expectation is taken over the probability of receiving an offer; i.e., the probability of $\tau_{ft} = \varnothing$ versus $\tau_{ft} \geq 0$. Conditional on receiving an offer, given ISP beliefs and strategies that offer is known.

5.2 Equilibrium, Estimation and Identification

The estimation strategy will be full maximum likelihood over the probability of observing the agreement timings in the data. To do so will require computing the model's equilibria for any given parameterization.

In the standard bilateral screening problem, the uninformed party (the ISP) makes a unique sequence of decreasing offers in a Perfect Bayesian Equilibrium (PBE), by trading off extracting surplus against pushing expected agreement to the discounted future. In the current framework, an ISP will trade off capturing a greater share of the surplus against both delay and the lost subscribers from disagreement as Netflix quality remains low. All else equal, greater subscriber loss will necessarily imply faster bargaining. While the introduction of these side-payments—and the idiosyncratic bargaining cost—does not affect the unique, decreasing sequence, the dependence of those side-payments on other ISPs' actions complicates equilibrium analysis.

The key requirement for PBE existence in this setting is that $\pi_{ft}(\boldsymbol{a}_t)$ is supermodular in the agreement vector. This condition is similar to one in the mobile phone adoption game of Björkegren (2015), where each agreement (adoption) increases the incentive for future agreement.[27]

Instead of solving the full model, I compute equilibria in the model under two alternative assumptions on ISPs' beliefs about the actions of their competitors: the first is that f believes all $-f$ will never agree, and the second is that f believes all $-f$ agree immediately. Defining subscriber profits in the first case as $\underline{\pi}_{ft}(a_{ft})$ and in the second case as $\bar{\pi}_{ft}(a_{ft})$, we have:

$$\underline{\pi}_{ft}(a_{ft}) \equiv \pi_{ft}(a_{ft}, \mathbf{0}) \tag{11}$$

$$\bar{\pi}_{ft}(a_{ft}) \equiv \pi_{ft}(a_{ft}, \mathbf{1}) \tag{12}$$

These assumptions imply a bilateral bargaining game with a unique equilibrium between any given ISP and Netflix. If the profit function from Section 4 is supermodular, then the parameters estimated under the two alternative assumptions will bound the true parameter values; if it is

[27]For other uses in structural estimation see also Jia (2008).

not supermodular, then the computation may have difficulty converging in initial guesses away from the true parameter values.[28] Note that embedded in these alternatives is the additional assumption that Netflix's strategy is separable across ISPs. In the next section I write $\pi_{ft}(a_{ft})$ as a placeholder for either $\bar{\pi}_{ft}$ or $\underline{\pi}_{ft}$.

Equilibrium

I consider pure strategy Markov Perfect Bayesian Equilibria. The unique equilibrium will consist of a set of strategies $\boldsymbol{\tau}_f$ for the ISP, a set of beliefs for the ISP B_f, and a set of cutoffs for Netflix such that strategies are optimal in expectation and beliefs are updated rationally according to strategies.

Define $\mu_{ft}(\mathcal{B}_{ft}, \boldsymbol{\tau}_f)$ to be Netflix's time t cutoff conditional on facing ISP beliefs $\mathcal{B}_{ft}$ and strategy $\boldsymbol{\tau}_f$; that is, Netflix types with $\mu_f \geq \mu_{ft}(\cdot)$ will accept τ_{ft}. Then the value function of the ISP can be written:

$$V_{ft}(\mathcal{B}_{ft}, \boldsymbol{\epsilon}_{ft}) = \max \Bigg\{ \max_{\tau_{ft}} G(\mu_f \geq \mu_{ft}(\mathcal{B}_{ft}, \tau_{ft})|\mathcal{B}_{ft}) \left(\tau_{ft} + \sum_{t' \geq t} \pi_{ft'}(1) \right) + \\ G(\mu_f < \mu_{ft}(\mathcal{B}_{ft}, \tau_{ft})|\mathcal{B}_{ft}) \left(\pi_{ft}(0) + \beta \mathbf{E}_{\boldsymbol{\epsilon}_{ft}} \left[V(\mathcal{B}_{ft+1}, \boldsymbol{\epsilon}_{ft+1}) | \mathcal{B}_{ft}, \tau_{ft}, a_{ft} = 0 \right] \right) + \epsilon_{1ft}, \\ \beta \mathbf{E}_{\boldsymbol{\epsilon}_{ft}} \left[V(\mathcal{B}_{ft}, \boldsymbol{\epsilon}_{ft+1}\right] + \epsilon_{0ft} \Bigg\}, \qquad (13)$$

which emphasizes that it is the $G(\cdot)$ distribution which is to be estimated.

Given an ISP strategy, at each time t that Netflix faces an offer, they will choose to accept the offer if it is more valuable than waiting for a possible offer at $t+1$:

$$\mu_f - \tau_{ft} \geq \beta \mathbf{E}_{\boldsymbol{\epsilon}_{ft+1}} \left[V_{xt+1}(\tau_{ft+1}, \mathcal{B}_{ft+1}) | \tau_{ft}, \mathcal{B}_{ft} \right]. \qquad (14)$$

Setting Equation 14 to equality will imply a cutoff value of μ_{ft} that depends on beliefs and the future sequence of offers.[29] I assume that a zero fee offer by an ISP is always accepted by Netflix.

[28] In unreported Monte Carlo simulations with supermodular profit and two ISPs, it is straightforward to solve the model numerically and achieve converge.

[29] Note that if $\beta = 0$ for Netflix, then Netflix optimally accepts whenever $\mu_f - \tau_{ft} \geq 0$. if $\boldsymbol{\epsilon}_{ft}$ has zero variance,

This was Netflix's explicit policy for bargaining ISPs, but also accommodates non-bargaining ISPs by providing a way to end bargaining in the first period with no surplus transfer.

Given the cutoff rule for Netflix acceptance, beliefs about the distribution of $G(\cdot)$ evolve according to a series of truncations. If τ_{ft} is rejected and Netflix's cutoff that period was μ_{ft}, then a rational belief for the ISP next period is that $\mu_f < \mu_{ft}$. Beliefs after any given history can therefore be summarized as an upper bound on the support of the surplus distribution.

Parametric Assumption on Surplus Distributions

The supply model should allow for the identification of size effects in the surplus distribution, to disentangle outside options versus size in bargaining outcomes. Size effects will imply that ISPs of different sizes experience different bargaining outcomes, conditional on outside options; however, in order to identify if there are actually size effects in bargaining, the distribution family picked for the surplus distribution needs to be able to accommodate size-invariance as a null hypothesis.

I assume that μ_f is drawn from a normal distribution

$$\mu_f \sim \boldsymbol{\sigma}_\zeta' \boldsymbol{w}_{2f} \cdot \mathcal{N}(0,1) + \boldsymbol{\lambda}' \boldsymbol{w}_{1f}, \tag{15}$$

where $\boldsymbol{w}_{1f}$ is a vector of observables that shift the mean of the surplus draw and $\boldsymbol{w}_{2f}$ is a vector of observables that shift the variance of the surplus draw. The ISPs and Netflix know $(\boldsymbol{\lambda}, \boldsymbol{\sigma}_\zeta)$ as well as the values of the covariates $\boldsymbol{w}_{1f}$ and $\boldsymbol{w}_{2f}$.

The next proposition implies that this distribution can accommodate size-invariance:

Proposition 1 *If each element of $\boldsymbol{w}_{1f}$ and $\boldsymbol{w}_{2f}$ is linear in market size, and if $\sigma_{f\epsilon}$ is also linear in market size, then scaling up market size will scale up the optimal offers $\boldsymbol{\tau}_f^*$ without changing the surplus split or expected disagreement length.*

Under Proposition 1, increasing an ISP's market size will linearly scale up the average surplus, the surplus variance, and the idiosyncratic bargaining cost variance. Since $\pi_{ft}(\cdot)$ also increases linearly in market size, the relative tradeoff the ISP faces in choosing its screening strategy will

then Netflix accepts if $\mu_f - \tau_{ft} \geq \beta(\mu_f - \tau_{ft+1})$ and the offer sequence is deterministic.

be unaffected, so the probabilities of agreement and offers at each t will not change. With proposition 1, testing for market size effects will involve including either a constant in $\boldsymbol{w}_{1f}$ or a term that is non-linear (e.g., quadratic) in market size, and then testing whether the coefficient on that non-linear term is statistically different from zero. If the coefficient is estimated as zero, then increasing the size of the ISP without changing its profit per household during agreement and disagreement periods—e.g., through a non-overlapping merger—will have no effect on disagreement duration, and size will not be important to bargaining outcomes.[30]

Estimation

The supply side parameters to estimate are θ^s, a vector that governs how ISP observables affect the distribution of surplus through $G(\mu_f|w_f, \theta^s)$, and the variance of the distribution of $\boldsymbol{\epsilon}_{ft}$, $\sigma_{f\epsilon}$.[31] The data used is the cross section of disagreement durations, as well as ISPs' estimated profit as a function of agreement from the demand section.[32]

Estimation will proceed by maximum likelihood. For a guess of θ^s, I solve the optimal cutoffs for Netflix and the optimal belief-contingent offers for an ISP, as well as its probabilities of making offers. At time $T+1$ (2015Q1) I impose that the ISP and Netflix both receive zero, so that there is a return to the status quo, and solve the value function and policy function for ISPs and Netflix by backwards induction.[33] Since each ISP has different period profits, I solve optimal strategies and value functions in the bilateral game separately for each ISP-Netflix pair.[34]

By adding the probabilities of all the possible paths by which agreement can occur at time t, the contribution to the likelihood of observing agreement at time t can be recovered. For instance, agreement in the second period may result from a rejected offer in the first period or from making no offer in the first period, followed by an accepted offer. Given the set of agreement times $\{y_f\}$ and the distribution of predicted times $\hat{y}(\theta^s)$, the negative log-likelihood to be minimized can be

[30] This proposition implies that the model includes as a base case the market-size invariance properties of the static Nash-in-Nash bargaining used in the literature.

[31] This variance cannot be normalized since π_{ft} appears in the ISP's value function without a coefficient, which uses up the one available normalization. $\sigma_{f\epsilon}$ may also vary with f.

[32] I drop the ISP Mediacom when performing the estimation as an outlier, see Appendix B.

[33] Ambrus, Chaney, and Salitsky (2016) assume a 50-50 split at $T+1$.

[34] The full estimation algorithm, which is also used for counterfactual simulations, is provided in Appendix E.

written as

$$\ell(\boldsymbol{y};\theta^s,\boldsymbol{w}) = -\sum_f \log \mathbf{P}_f(\hat{y}_f(\theta^s) = y_f) \tag{16}$$

The idiosyncratic error $\boldsymbol{\epsilon}$ plays a key role in ensuring the likelihood exists for a wide range of parameter values, which aids in estimation. Without $\boldsymbol{\epsilon}$, some values of θ^s might imply that bargaining for f ends with certainty before y_f, implying $\ell(\theta^s,\cdot) = -\infty$ for those values.

Identification

In the base specification, I assume $\boldsymbol{w}_{1f} = [1, M_f]$, $w_{2f} = M_f$, and $\sigma_{f\epsilon} = w_{3f} \times \sigma_\epsilon$, with $w_{3f} = M_f$, where M_f is the number of households that ISP f could serve in 2013Q3. There are therefore four parameters to identify: $\theta^s \equiv (\lambda_1, \lambda_2, \sigma_\zeta, \sigma_\epsilon)$. I assume a quarterly discount rate of $\beta = 0.975$.

The key feature to identify is whether $\lambda_1 < 0$; i.e., if there are size effects pushing the mean of the surplus distribution down by the same amount for all ISPs. With $\lambda_1 < 0$ and $\lambda_2 > 0$, the mass of the normally distributed surplus for an ISP will lie mostly below zero for smaller ISPs and mostly above zero for large ISPs, which will induce a shorter bargaining time for small ISPs and a longer one for large ISPs, as in the data. However, since small ISPs also tend to have worse outside options, they will bargain more quickly as they have more to lose from periods of disagreement even with positive surplus on the table. The correlation between bargaining time and market size—conditional on outside options—will identify (λ_1, λ_2).

σ_ζ and σ_ϵ must be weakly positive, and at least one must be strictly positive to generate non-zero delays since otherwise the surplus distribution is degenerate. Their joint identification relies on cross-sectional variation in delay times. Neither can be too large, since if either is increased substantially it would imply large gains to bargaining for all ISPs—either because there is more surplus mass above zero, or because there are idiosyncratic benefits to waiting—which would counterfactually imply that most small ISPs would have a positive expected bargaining time. A more right-skewed distribution of agreement times implies σ_ϵ is relatively large, since in that case it benefits ISPs to wait to make offers, while a left-skewed distribution implies that σ_ζ is relatively large.

Finally, the level differences in $\pi_{ft}(1)$ and $\pi_{ft}(0)$ pin down the scale of the estimated parameters. For instance, doubling all θ^s will imply that lost subscriber profits play less of a role in speeding up bargaining, because the importance of lost profits relative to available surplus has decreased. If variation in these profits helps explain cross-ISP variation in agreement times, then such a doubling would decrease the likelihood of predicting the observed agreement timings.

As in Merlo and Tang (2012), additional data could help identify the model and loosen parametric restrictions. Observing when offers were made would help identify the variance of the idiosyncratic bargaining cost σ_ϵ, while observing the value of the final accepted offer would help identify λ and the mean of the surplus distribution.[35]

5.3 Results

I scale the profit π_{ft} to be in hundreds of millions of dollars, and the market size to be in hundreds of millions of households. Since profit is the normalizing factor, all parameters are therefore estimated in hundreds of millions of dollars. Estimated parameters and standard errors under alternative assumptions are presented in Table 4.

[Table 4 about here]

The baseline results in specification (1) imply a fixed cost λ_1 of realizing the surplus of 18.3 million USD, and an increase in surplus value λ_2 to Netflix of 817 million USD per 100 million households passed, or $8.17 per household.

The specification (1) results are computed under the assumption that all other ISPs immediately reach agreement with Netflix. Specification (2) finds a slightly higher fixed cost under the assumption that other ISPs disagree with Netflix indefinitely: this follows because the gain from agreement for an ISP is lower when all other ISPs have not agreed, leading to less costly delays and slower bargaining, and therefore a higher fixed cost to speed negotiations. Because

[35]In Monte Carlo exercises reported in Appendix E.1 the maximum likelihood function is mostly well-behaved around the true parameters even with a limited number of observations and only durations observed.

the estimated parameters are so similar in specifications (1) and (2), I will use (1) going forward.[36]

I provide two checks on whether the estimates of the surplus shifter λ_1 is reasonable. First, I plot predicted disagreement times from specification (1) against actual disagreement times in Figure 7 and find that the fit is fairly close to the 45 degree line. Compared to specification (3) which assumes $\lambda_1 = 0$ and relies only on differences in outside options to explain bargaining time variation, including the size effect λ_1 substantially improves the ability of the model to replicate the patterns in the data.

[Figure 7 about here]

Second, I show that the model's predicted payments from Netflix to ISPs accord with industry estimates. Although the exact details of ISP negotiated agreements are confidential, an expert with domain knowledge approximated Netflix's cost savings from avoiding third-party transit fees to Comcast's network at 12 million USD per year.[37] Using the estimation's assumed yearly discount rate of $\beta^4 = 0.904$ implies a lump sum payment from Netflix to Comcast of 124 million USD, if the savings were in perpetuity. The model's estimate is in the correct range: I estimate that Comcast's expected fee was 155 million USD, where the expectation is of the optimal fee with respect to the probability of each history of offers/non-offers and rejections that would lead the bargaining game to end in the third quarter, as it did in the data.

The role of size in delays and surplus splits

Using the estimated model, I explore the role of size effects and outside options in explaining variation in bargaining delays. If size effects are unimportant in prolonging disagreement—and the associated period of lower consumer welfare—then there is no reason not to allow mergers between ISPs that serve separate markets. Moreover, the model will provide estimates of the

[36]The similarity of (1) and (2) also suggests that there is not much precision to be gained from a full solution method that makes ISP strategies contingent on other ISPs' states.

[37]A copy of the page https://blog.streamingmedia.com/2014/02/heres-comcast-netflix-deal-structured-numbers.html is available through the Internet Archive. Thanks to Paul LaFontaine at the FCC for providing this reference.

expected surplus splits between each bargaining pair, which will have implications for the degree to which the fixed cost improves Netflix's bargaining position.

To understand the role of the fixed cost on bargaining outcomes, I simulate the counterfactual expected disagreement durations and surplus splits for each ISP when $\lambda_1 = 0$. The results are shown in Figure 8, where they are plotted against the baseline predictions from Specification (1) for comparison.

[Figure 8 about here]

The main finding is that the fixed cost explains most of why large ISPs take longer to bargain, and most of why small ISPs bargain quickly. Without the fixed cost, the model predicts counterfactually long bargaining times for the small ISPs, and counterfactually short bargaining times for large ISPs, although smaller ISPs do still take slightly less time to bargain on average owing to their worse outside options.

Intuitively, removing the fixed cost does two things. First, for large ISPs that were already bargaining over a positive surplus, it weakens Netflix's bargaining position by widening the gap between their agreement and disagreement payoffs. This induces Netflix to more quickly accept a less generous offer, which explains why large ISPs tend to bargain faster and receive a larger share of the surplus after removing the fixed cost. Second, for small ISPs that had no positive surplus to bargain over, there now is one. Even though their bargaining position is weaker than the larger ISPs', they now find it optimal to make offers.

Importantly, the result that Netflix retains less surplus after removing the fixed cost also appears in the static Nash bargaining model (see Appendix A.) That the dynamic model makes the same prediction validates that the underlying marginal incentives governing ISP behaviour in the model are operating as expected.

Robustness

In this section, I evaluate whether variation in bargaining power can explain why large ISPs take longer to bargain than smaller ISPs. I draw on the insight from Collard-Wexler, Gowrisankaran, and Lee (2019) that relative patience maps directly into the Nash bargaining residual when the Nash-in-Nash bargaining framework is microfounded as a sequence of bilateral alternating offer games. Patience operates also maps to bargaining power in my setting, as a more patient ISP will set a more slowly decreasing sequence of optimal offers compared to a less patient ISP, screening Netflix more finely and extracting a greater share of its surplus in expectation.

I re-estimate Specification (1) using a firm's weighted average cost of capital (WACC) as a proxy for its patience, since the WACC is used by firms in practice to discount future revenue and costs. WACC data for 2013Q3 is taken from Bloomberg for all publicly traded firms, and imputed for private firms based on ISP size, technology, and subsidiary status.[38] In the data, the 10 largest firms have a WACC that is 11% lower than the remaining firms, implying lower costs of waiting.

Results are reported in Table 4, Specification (4), where I find that the magnitude of the estimated fixed cost is not much changed from the baseline, although the standard error does increase.

6 Counterfactuals

6.1 Mergers

With the model estimated, it is straightforward to do merger analysis between non-overlapping ISPs. Since mergers typically only occur between ISPs using the same last-mile transmission technology, the non-overlapping case is the empirically relevant one. To evaluate a merger between f and f', I resimulate the bargaining game with the estimated parameters in the first column of Table 4, new market size $M_f + M_{f'}$, and new profits $\pi_{ft} + \pi_{f't}$.

In this framework, a merger will impact consumer welfare by changing how long bargaining–and hence, the period of lower quality Netflix—will last at the merging firms; as per the reduced form, I do not focus on pricing. A merger will also affect the size of Netflix's surplus by removing duplicated fixed costs, and what share of its surplus Netflix retains by altering Netflix's bargaining

[38] See Appendix E.3 for details.

positions and ISP's screening incentives.

[Table 5 about here]

In Table 5, I simulate several bilateral mergers between Cable ISPs. In the first row, the Comcast-TWC merger was proposed towards the end of 2014 and would have combined the two largest cable internet providers, with connections to over half of U.S. households. I find that their merger would have actually slightly decreased disagreement times, slightly reduced ISPs' share of the Netflix surplus, and that Netflix would therefore have increased the amount of surplus it retained by 4.38 million USD. Intuitively, without a fixed cost, the bargaining model has the property that linearly scaling up a firm's market size and profits leads to the exact same relative tradeoffs for ISPs in making offers—it is the fixed cost, and the extra bargaining power it gives Netflix, that leads to different bargaining outcomes across different sized firms. Comcast and TWC are already so large that although their combined surplus is larger than the sum of the constituent firms, the non-linearity from the fixed cost is not important given the total size of the surplus, and so the effect of their merger is essentially that of a linear scaling up.

In the second row, the Charter-TWC merger actually went through in 2015. While in this case there is a slight reduction in bargaining times, there is again essentially no change in the surplus split. Charter and TWC were both large at the time of the merger, and so the size effect is still mostly inoperative.

In contrast, having smaller ISPs merge can have much more substantial effects on disagreement times and ISP shares and hence, Netflix's total retained surplus. Ex ante, both Cablevision and Metrocast—two smaller ISPs— face the incentive to make quick, generous offers to avoid losing subscriber revenue since there is very little expected surplus. However, when they merge, the removal of the duplicated fixed cost substantially increases Netflix's surplus, and the ISPs thereby increase their bargaining power and share.

Although mergers between non-overlapping firms do reliably generate efficiencies in this context, the analysis shows that those efficiencies might be completely captured by the merging firms. Because Netflix is the party that generates the surplus through its investment in infrastructure, capturing these efficiencies is not a mere transfer of surplus, but might actually disincentivize

Netflix from making its investment. In other words, because of how size effects degrade the upstream supplier's bargaining position, a merger between two non-overlapping firms that does not improve their relative outside option can still worsen the hold-up problem faced by an upstream supplier. While past papers have estimated that larger firms have a larger bargaining residual (Grennan, 2013), by explicitly estimating the size effect my paper allows for internally consistent counterfactuals to quantify the consequences of firm size separate from outside options.

6.2 Investment incentives and welfare effects

If the fixed cost λ_1 corresponds to an outlay by Netflix for its cost-saving interconnection infrastructure, then its removal roughly maps to a policy of a tax credit for investment in the internet backbone. Since policies supporting fixed broadband investment have been popular in the U.S. in the past, it is worth asking how the removal of fixed costs would change Netflix's incentive to invest.

As with the merger analysis, removing λ_1 has a theoretically ambiguous effect on Netflix's retained surplus. First, although there is now more surplus on the table, Netflix has a worse bargaining position over the inframarginal surplus, and so its total take may decrease. Second, the surplus Netflix retains will be discounted more heavily if ISPs take more time to bargain. Subsidizing Neflix's investment may therefore *decrease* its incentive to deploy the technology, by weakening its bargaining position enough that it is unable to reap the benefits of the greater surplus.

I find that on balance, Netflix's surplus increases under the tax credit regime, but much of the value of the credit is captured by ISPs. From the 10 ISPs with which Netflix had positive surplus ex ante, Netflix increases its surplus by 66 million USD (from 800 to 866), which is less than 40% of the tax credit value of 10×18.3 million. From the small ISPs Netflix captures 142 million USD in surplus where before surplus was zero.[39]

For a given consumer, the welfare effect is ambiguous and depends on which ISPs are in her choice set, as suggested by Figure 8. Households with bargaining ISPs in their choice set benefit

[39] All values are expected surplus discounted to 2013Q3.

from faster agreement, while those with small ISPs in their choice set are harmed by a now-positive disagreement duration. The welfare changes in practice are quite small: for consumers with a large ISP in their choice set (90.1% of households) welfare improves by 0.03% on average, while for consumers with only small ISPs in their choice sets (9.9% of households) welfare falls by 0.2%.

7 Conclusion

In this paper I analyze the consequences of ISP sizes and outside options on bargaining in the broadband internet industry. Using bargaining durations between Netflix and U.S. ISPs and rich demand data, I build and estimate a structural model of dynamic bargaining and show that size is far more important than outside options in determining the length of bargaining and the surplus split. With small ISPs, counterfactual mergers between non-overlapping firms harm both consumer welfare and Netflix's incentive to invest in the internet backbone.

The development of empirical modelling tools to accommodate upstream bargaining has substantially enriched our understanding of the consequences of mergers, with the caveat that merged firms' residual bargaining power is still poorly understood. While this paper explores how structurally modelling size effects can lead to novel insights compared to bargaining residuals, incorporating size effects into workhouse multilateral bargaining frameworks is still a necessary further step.

References

Ambrus, A., E. Chaney, and I. Salitsky (2016): "Pirates of the Mediterranean: An Empirical Investigation of Bargaining with Transaction Costs," Discussion paper, Economic Research Initiatves at Duke.

Argentesi, E., and L. Filistrucchi (2007): "Estimating Market Power in a Two-Sided Market: The Case of Newspapers," *Journal of Applied Econometrics*, 22(7), 1247–1266.

Berry, S., J. Levinsohn, and A. Pakes (1995): "Automobile prices in market equilibrium," *Econometrica*, pp. 841–890.

——— (2004): "Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Car Market," *Journal of Political Economy*, 112(1), 68–105.

Berry, S., and E. Tamer (2006): "Identification in Models of Oligopoly Entry," *Advances in Economics and Econometrics*, 2, 46–85.

Björkegren, D. (2015): "The Adoption of Network Goods: Evidence from the Spread of Mobile Phones in Rwanda," Discussion paper.

Chandra, A., and A. Collard-Wexler (2009): "Mergers in Two-Sided Markets: An Application to the Canadian Newspaper Industry," *Journal of Economics and Management Strategy*, 18, 1045–1070.

Chipty, T., and C. M. Snyder (1999): "The Role of Firm Size in Bilateral Bargaining: A Study of the Cable Television Industry," *The Review of Economics and Statistics*, 81(2), 326–340.

Cisco (2017): "The zettabyte era–trends and analysis," *Cisco white paper*.

Collard-Wexler, A., G. Gowrisankaran, and R. S. Lee (2019): "âĂIJNash-in-NashâĂİ bargaining: a microfoundation for applied work," *Journal of Political Economy*, 127(1), 000–000.

Crawford, G., and A. Yurukoglu (2012): "The Welfare Effects of Bundling in Multichannel Television Markets," *American Economic Review*, 102(2), 643–685.

Cremer, J., P. Rey, and J. Tirole (2000): "Connectivity in the commercial Internet," *The Journal of Industrial Economics*, 48(4), 433–472.

Dafny, L., K. Ho, and R. Lee (2016): "The Price Effects of Cross-Market Hospital Mergers," Discussion paper, National Bureau of Economic Research.

Goeree, M. S. (2008): "Limited Information and Advertising in the U.S. Personal Computer Industry," *Econometrica*, 76(5), 1017–1074.

Gowrinsankaran, G., A. Nevo, and R. Town (2015): "Mergers When Prices Are Negotiated: Evidence from the Hospital Industry," *American Economic Review*, 105(1), 172–203.

Grennan, M. (2013): "Price Discrimination and Bargaining: Empirical Evidence from Medical Devices," *American Economic Review*, 103(1), 145–177.

Ho, K. (2009): "Insurer-Provider Networks in the Medical Care Market," *American Economic Review*, 99(1), 393–430.

Honoré, B., and A. de Paula (2010): "Interdependent Durations," *Review of Economic Studies*, 77(3).

Jia, P. (2008): "What Happens When Wal-Mart Comes to Town: An Empirical Analysis of the Discount Retailing Industry," *Econometrica*, 76(6), 1263–1316.

Merlo, A., and X. Tang (2012): "Identification and Estimation of Stochastic Bargaining Models," *Econometrica*, 80(4), 1563–1604.

Merlo, A., and C. Wilson (1995): "A Stochastic Model of Sequential Barganing with Complete Information," *Econometrica*, 63, 317–399.

Nevo, A. (2000): "A Practitioner's Guide to Estimation of Random-Coefficients Models of Logit Demand," *Journal of Economics and Management Strategy*, 9(4), 513–548.

Scherbakov, O. (2015): "Measuring Consumer Switching Costs in the Television Industry," Discussion paper.

Main Text Figures and Tables

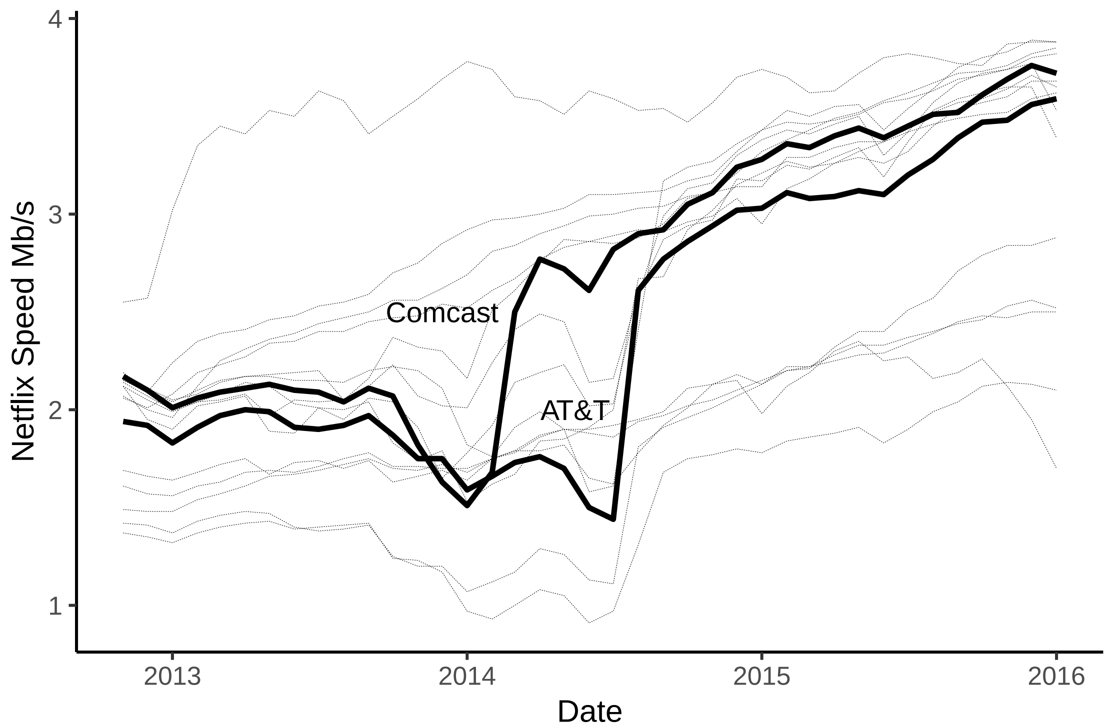

Figure 1: Average Netflix throughput to 16 U.S. ISPs

Note: 16 of 49 ISPs are shown. This subset of ISPs has data from the beginning of Netflix's monitoring program, while other ISPs are added in periodically.

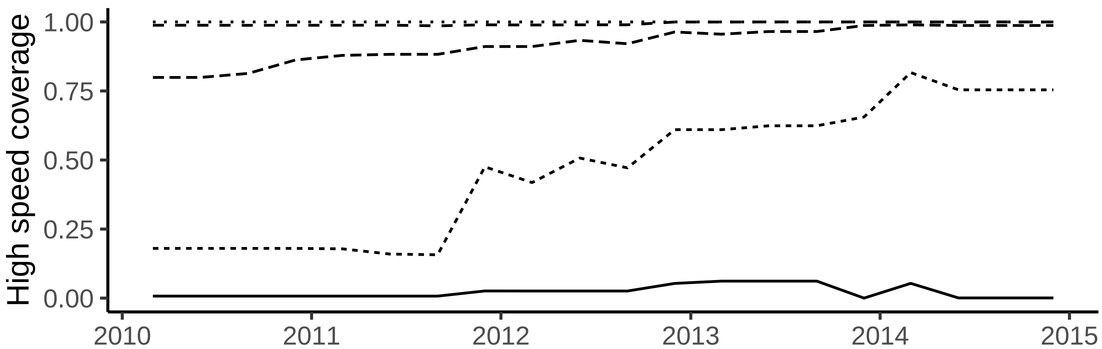

(a)

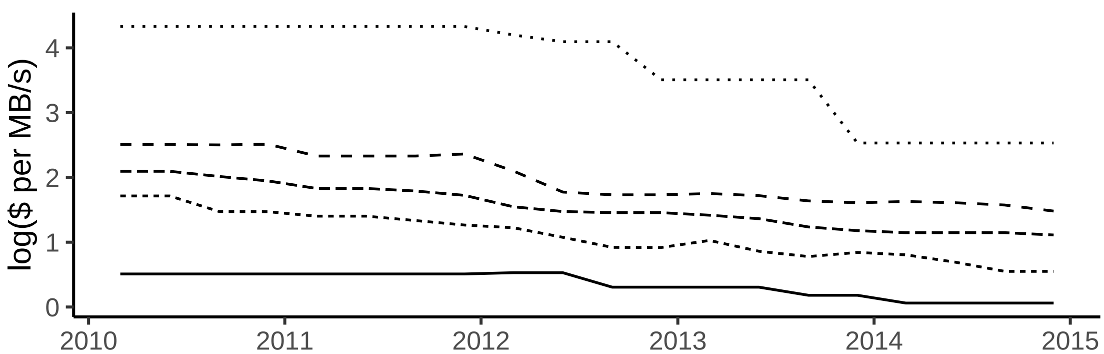

(b)

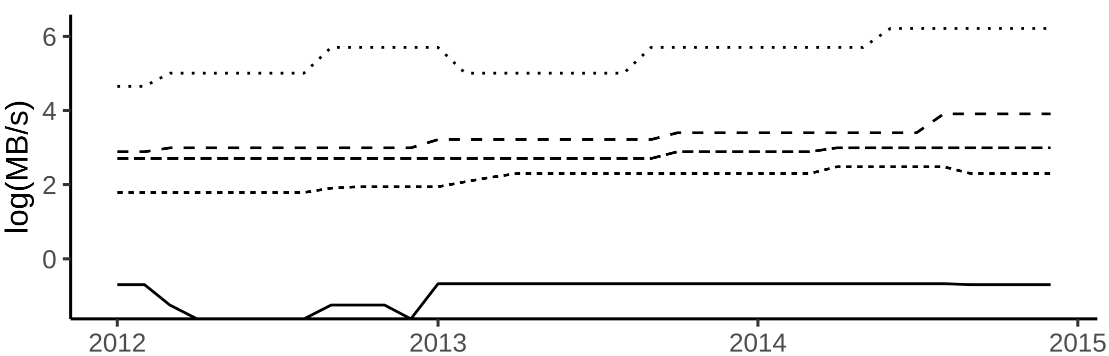

(c)

Figure 2: Variation over time

(a) **ISP coverage** (Share of ISPs' market with 50Mb/s or higher maximum advertised speed, quantiles.) (b) **ISP prices** (log price per Mb/s, ISP median plans, quantiles.) (c) **Household choices** (log Mb/s, quantiles.)

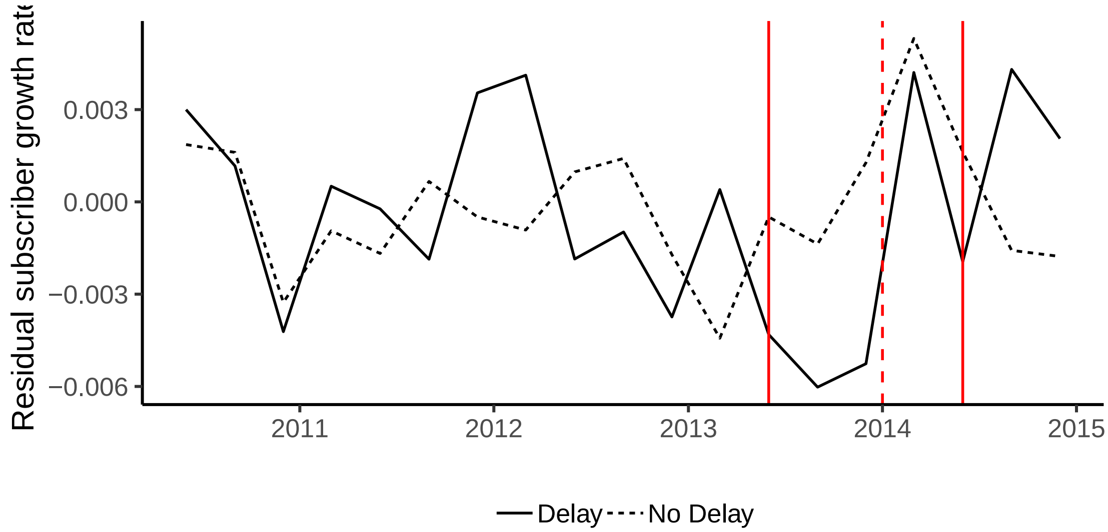

Figure 3: Residual ISP subscriber growth rates

Note: Lines are median residuals from the regression of subscriber growth on covariates in Equation 1, split by whether an ISP experienced a bargaining delay or not.

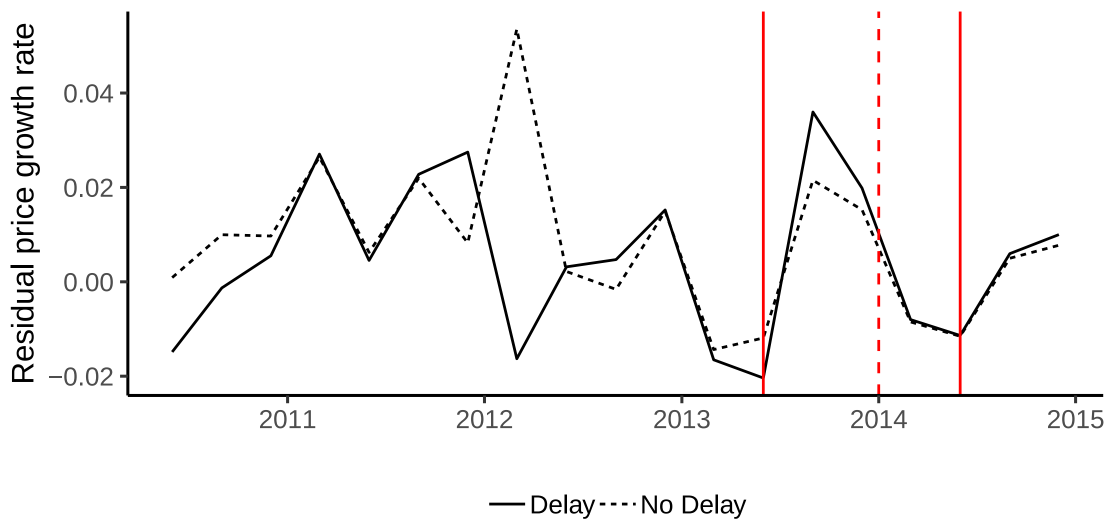

Figure 4: Residual ISP price growth rates

Note: Note: Lines are median residuals from the regression of price growth on covariates in Equation 2, split by whether an ISP experienced a bargaining delay or not.

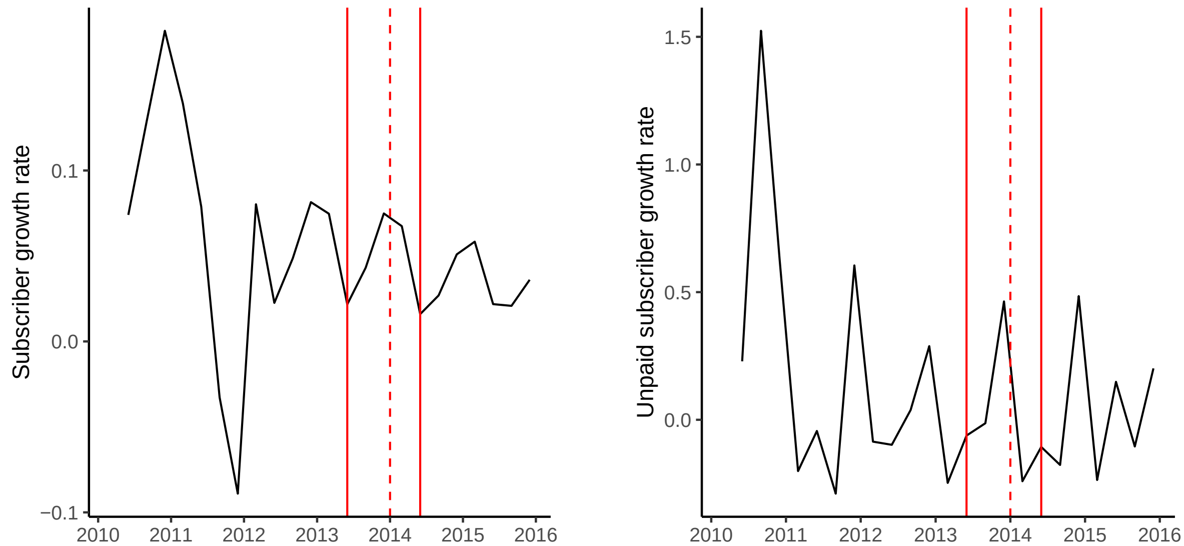

Figure 5: Raw Netflix growth rates

Note: Post-2012Q1 subscriber growth uses only streaming subscribers, pre-2012Q1 includes streaming and DVD.

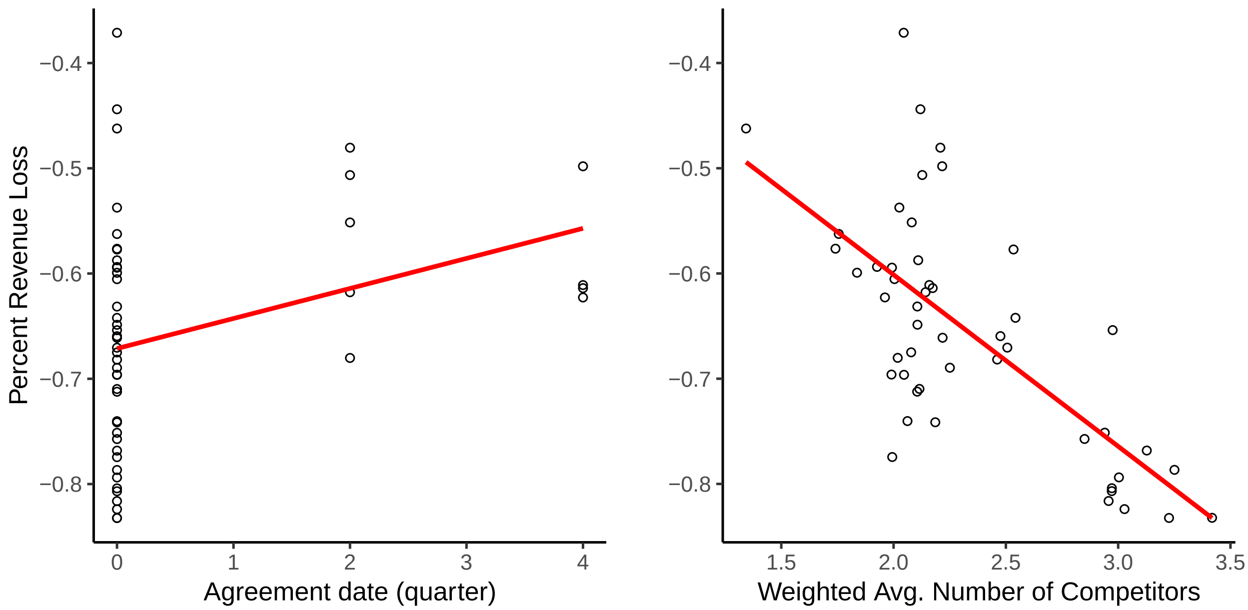

Figure 6: Revenue loss variation

Note: The percent revenue loss is computed in the first quarter in which disgreement is possible, 2013Q3, and is conditional on all other ISPs facing no slowdown. All ISPs face a a satellite ISP competitor.

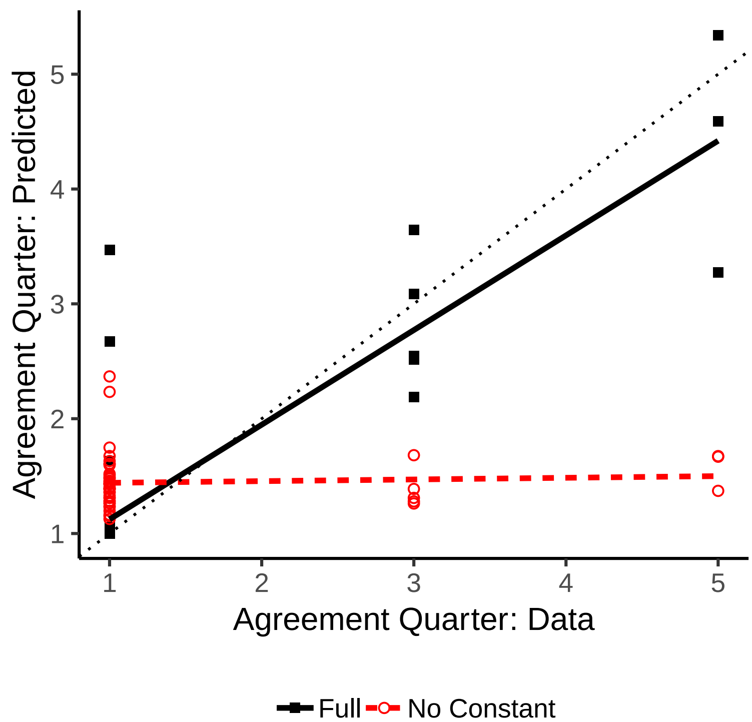

Figure 7: Supply Model Fit

Note: The full model predictions use specification (1) with $\bar{\pi}_{ft}$. The ISP Mediacom is dropped in the estimation.

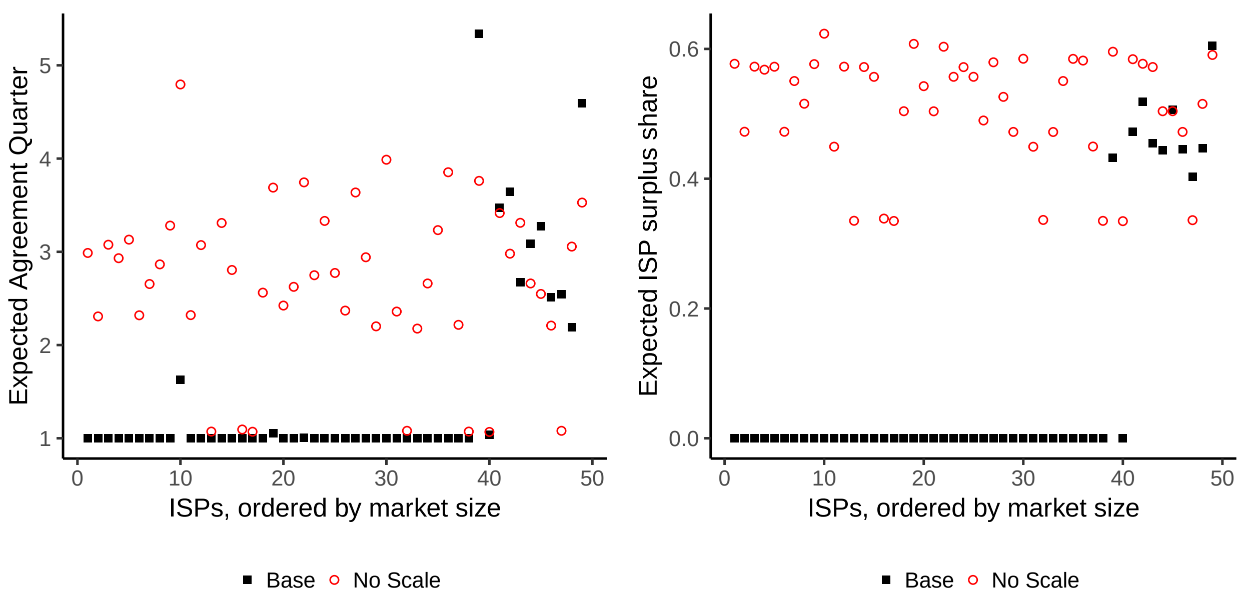

Figure 8: Size effects, expected delays and surplus splits

Note: ISP market size for ordering purposes is taken in 2013Q3. The no-size-effect simulations are performed with $\lambda_1 = 0$ for the estimates from specification (1) in Table 4. The base simulations are the model predictions from specification (1).

Table 1: Cross-sectional summary statistics

Variable	Num. ISPs	Mean	Std.Dev.	Min.	Max.
Bargaining data					
Agreement Time	49	1.53	1.21	1	5
Subscriber data					
Subscribers (000)	31	3,037.58	4,896.70	13	21,962
Choice set data					
Market (000)	55	4,998.71	11,145.50	45.13	52,678.75
Competitors	55	1.33	0.47	0.34	2.42
Competitors, incl. Wireless	55	3.33	0.47	2.34	4.42
Price and plan feature data					
Basic Plan Speeds (MB/s)	60	8.88	16.28	0.50	100
Median Plan Speeds (MB/s)	60	42.69	95.72	1	550
Basic Plan Price ($)	60	43.26	15.58	1	109
Median Plan Price ($)	60	68.10	25.09	29.95	199.95
Household choice panel data					
Chosen Plan Speed (MB/s)	17	21.12	21.51	0.29	500

Note: Num. ISPs indicates the number of ISPs present in that dataset. Agreement time is in quarters, with $t = 1$ in 2013Q3. Competitors are a population-weighted average based on market overlap. Median plan price and speed are within-ISP-menu in each quarter.

Table 2: Correlations with Agreement Times

	(1)	(2)	(3)	(4)
Mkt. Size	0.805^{**}		0.785^{**}	0.786^{**}
	(0.139)		(0.136)	(0.137)
Num. Comp.		-0.247^{**}	-0.125^{**}	-0.121^{*}
		(0.103)	(0.062)	(0.070)
Controls	No	No	No	Yes
Observations	49	49	49	49
R^2	0.443	0.042	0.453	0.453

Note: The dependent variable is the agreement time ($t \in (1, \ldots, 5)$. Market size and number of competitors are standardized beta coefficients. Standard errors are heteroskedasticity robust, *p<0.1; **p<0.05.

Table 3: Demand Model Estimates

		Controlling for Choice Sets		
	(1)	(2)	(3)	(4)
Disagree$\times t$	-0.006*	-0.010**	-0.009**	-0.010**
	(0.003)	(0.005)	(0.004)	(0.005)
$\overline{\log(p)}$	-0.027	-0.064*	-0.192	
	(0.021)	(0.035)	(0.207)	
$\overline{\log(q)}$	-0.002	0.008	0.034	
	(0.010)	(0.019)	(0.046)	
$\bar{\alpha}^p$				-13.32**
				3.244
σ^p				-0.778**
				0.092
$\bar{\alpha}^q$				-15.033**
				3.244
Implied Price Elast.	0.027	0.064	0.192	5.901
ISP FE	✓	✓	✓	✓
Date FE	✓	✓	✓	✓
ISP-specific trend	✓	✓	✓	✓
IV			✓	✓
First stage F-stat			14.19	
Observations	523	523	493	—

Note: Specification (1) is multinomial logit, (2) adds in sub-national choice sets, (3) instruments for price, which includes lagged instruments. (4) adds in non-random plan choice and is estimated with two-step GMM. First stage F-stat jointly tests excluded instruments, and is not present for (4) since price parameters enter non-linearly. Standard errors are cluster robust at the ISP level. *p<0.1; **p<0.05.

Table 4: Supply Model Estimates

	Base			Robustness
	(1)	(2)	(3)	(4)
Surplus mean shifters (w_{1f})				
Constant (λ_1)	-0.183^{*}	-0.190^{*}		-0.171^{*}
	(0.096)	(0.116)		(0.102)
M_f (λ_2)	8.173^{**}	8.222^{**}	35.504	7.13^{**}
	(2.231)	(1.758)	(1638)	(1.552)
Surplus variance shifters (w_{2f})				
M_f (σ_ζ)	1.951^{**}	2.338^{**}	0.771	2.661^{**}
	(0.338)	(0.389)	(1.083)	(0.489)
Bargaining shock shifters (w_{3f})				
M_f (σ_ϵ)	0.011^{**}	0.011^{**}	0.678	0.012^{**}
	(0.005)	(0.004)	(25.43)	(0.005)
N	48	48	48	48

Note: Specification (1) uses $\bar{\pi}_{ft}$ (all competitors agree instantly), specification (2) uses $\underline{\pi}_{ft}$ (all competitors never agree). In (4), the four largest ISPs have a higher discount rate than other ISPs. *p<0.1; **p<0.05, standard errors computed using 50 bootstrap replications. The ISP Mediacom is dropped as an outlier, see Appendix B.

Table 5: Cable Mergers

	Pre-Merger		Post-Merger		Netflix ΔSurplus	
ISPs	Length	ISP Share	Length	ISP Share	mil. USD	Percent
Comcast, TWC	2.31	0.45	2.22	0.46	4.38	1.34
Charter, TWC	2.69	0.44	2.12	0.46	7.22	4.32
Cablevision, TWC	2.69	0.45	2.95	0.49	-1.63	-1.26
Cablevision, Metrocast	3.31	0.44	4.90	0.51	-1.10	-7.33

Note: Length refers to the expected length of disagreement in quarters, ISP share is the expected fraction of surplus the ISP extracts. In the pre-merger case these are weighted averages across the merging ISPs, weighting by market size, and in the post-case case it is for the single merged firm. The last column divides the USD MM Netflix gain/loss in surplus by Netflix's expected retained surplus pre-merger.

Online Appendix

A Size Effects in Static Bargaining

In this section I show that in the bilateral Nash bargaining solution, introducing a size effect into Netflix's surplus implies that as market size increases but the ISP's outside option remains the same, Netflix's share of the marginal surplus it brings to bargaining will fall.

Let r^a be the per-household ISP surplus if agreement is reached with Netflix, r^{na} be the corresponding surplus if there is no agreement with Netflix (the disagreement payoff), and $\Delta r \equiv r^a - r^{na} > 0$. Let s be the value of Netflix's per-household cost savings surplus, let $f > 0$ be the fixed cost associated with that surplus, and assume a disagreement payoff of zero. Let M be the total market size, and τ be the negotiation transfer from Netflix to the ISP. Lastly, let β be the bargaining power of the ISP, and $1 - \beta$ the corresponding bargaining power of Netflix, where $\beta \in (0, 1)$. Then τ solves:

$$\max_\tau \left(M \cdot r^a + \tau - M \cdot r^{na}\right)^\beta \times \left(M \cdot s - f - \tau\right)^{1-\beta}.$$

The solution is $\tau = M \cdot (\beta s - (1 - \beta)\Delta r) - \beta f$. Notice that if not for f, the per-household transfer τ/M would be invariant to market size. That is, the standard Nash bargaining model predicts no effect of size on the surplus split. Papers such as Grennan (2013) and Crawford and Yurukoglu (2012) find that $\beta = f(M)$ with $f'(M) > 0$, but the relationship is not explicit, so counterfactual choices of β must necessarily be ad hoc.

In the paper, I am concerned with how much of Netflix's marginal surplus $S = Ms - f$ the ISP extracts. The ISP's share of this is τ/S. Taking the derivative with respect to M, the expression is

$$\frac{\partial(\tau/S)}{\partial M} = \frac{f \cdot (1 - \beta)\Delta r}{(Ms - f)^2} > 0 \qquad \Leftrightarrow \qquad \beta < 1$$

which is always positive if $\beta < 1$; i.e., the ISP increases its share of Netflix's marginal surplus as market size increases, and Netflix decreases its share. Intuitively, because Netflix gains relatively more from bargaining as the market size increases, its outside option is becoming relatively worse, which weakens its bargaining position. If $\beta = 1$, it implies that the ISP has all the bargaining power, and so extracts all of Netflix's surplus regardless of market size, so that the derivative is zero.

B Data

B.1 Supply Data

I draw on Netflix's throughput data, as well as MLab measurement data, Center for Applied Internet Data Analysis (CAIDA) data on interconnection (described below), and Netflix qualitative business filings data to argue that (1) quality reductions were for business reasons, not technical ones and (2) the duration of quality reductions corresponds to the duration of bargaining over interconnection.

A key document in my analysis is the public (partially redacted) version of Netflix's "Petition to Deny". The Petition to Deny is a legal document filed by Netlix to the Federal Communications Commission 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. While I take statements in the document about agreement timings at face value, statements that the ISPs are to blame for the slowdown will need to be evaluated.

CAIDA provides information on whether an Autonomous System (AS)—a collection of IP addresses under the control of one network provider, such as Comcast, Netflix OpenConnect, or a third-party operator such as Limelight or Cogent—is directly interconnected with any other AS. Two ASs that are not directly interconnected can still exchange information, but must rely on a third party AS (or chain of third parties) that is connected to both original Autonomous Systems to do so. The interconnection data is publicly available, and provided at monthly intervals.

Ruling Out ISP Technical Difficulties

From the Netflix throughput data it is clear that some ISPs experience a slowdown and others do not. However, it is possible that affected ISPs were suffering from general network problems and not Netflix-specific problems. I falsify this hypothesis with data from MLab, an independent research group that measures the rate of transmission of data between last-mile ISPs and transit ISPs. Since Netflix primarily uses the service of the transit ISP Cogent during this time, if slowdowns occur between Cogent and the affected ISPs but not between other transit providers and the affected ISPs, then it will be evidence of a Netflix-ISP specific problem. Figure B.1 illustrates that starting in July 2013, there was a sharp drop in Netflix streaming quality to Comcast, Verizon and TWC. Throughput of other content to these ISPs was not affected, and throughput of Netflix to ISPs such as Cox—a large provider with roughly 5.5% share of subscribers nationwide—was also not affected. Although the MLab data is not comprehensive in its coverage of ISPs, it strongly suggests that slowdowns occurred pairwise. Combined with the qualitative data cited in Section 2.1, it is clear that the slowdowns were business driven.

Assigning Responsibility for the Slowdown

Although Netflix's standing offer to install Open Connect servers was not taken up by the largest ISPs until 2013, transmission of Netflix content to the end users remained reliable until mid-2013. At that point, Netflix, the last-mile ISPs, or both Netflix and last-mile ISPs either actively precipitated a collapse in reliable transmission, or failed to proceed with status quo network upgrades.[40] The end result was that quality of service plummeted for the largest ISPs.

From Netflix's point of view, it is the ISPs that were at fault: Netflix argues that 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

[40] Harvard Business School case N9-616-007.

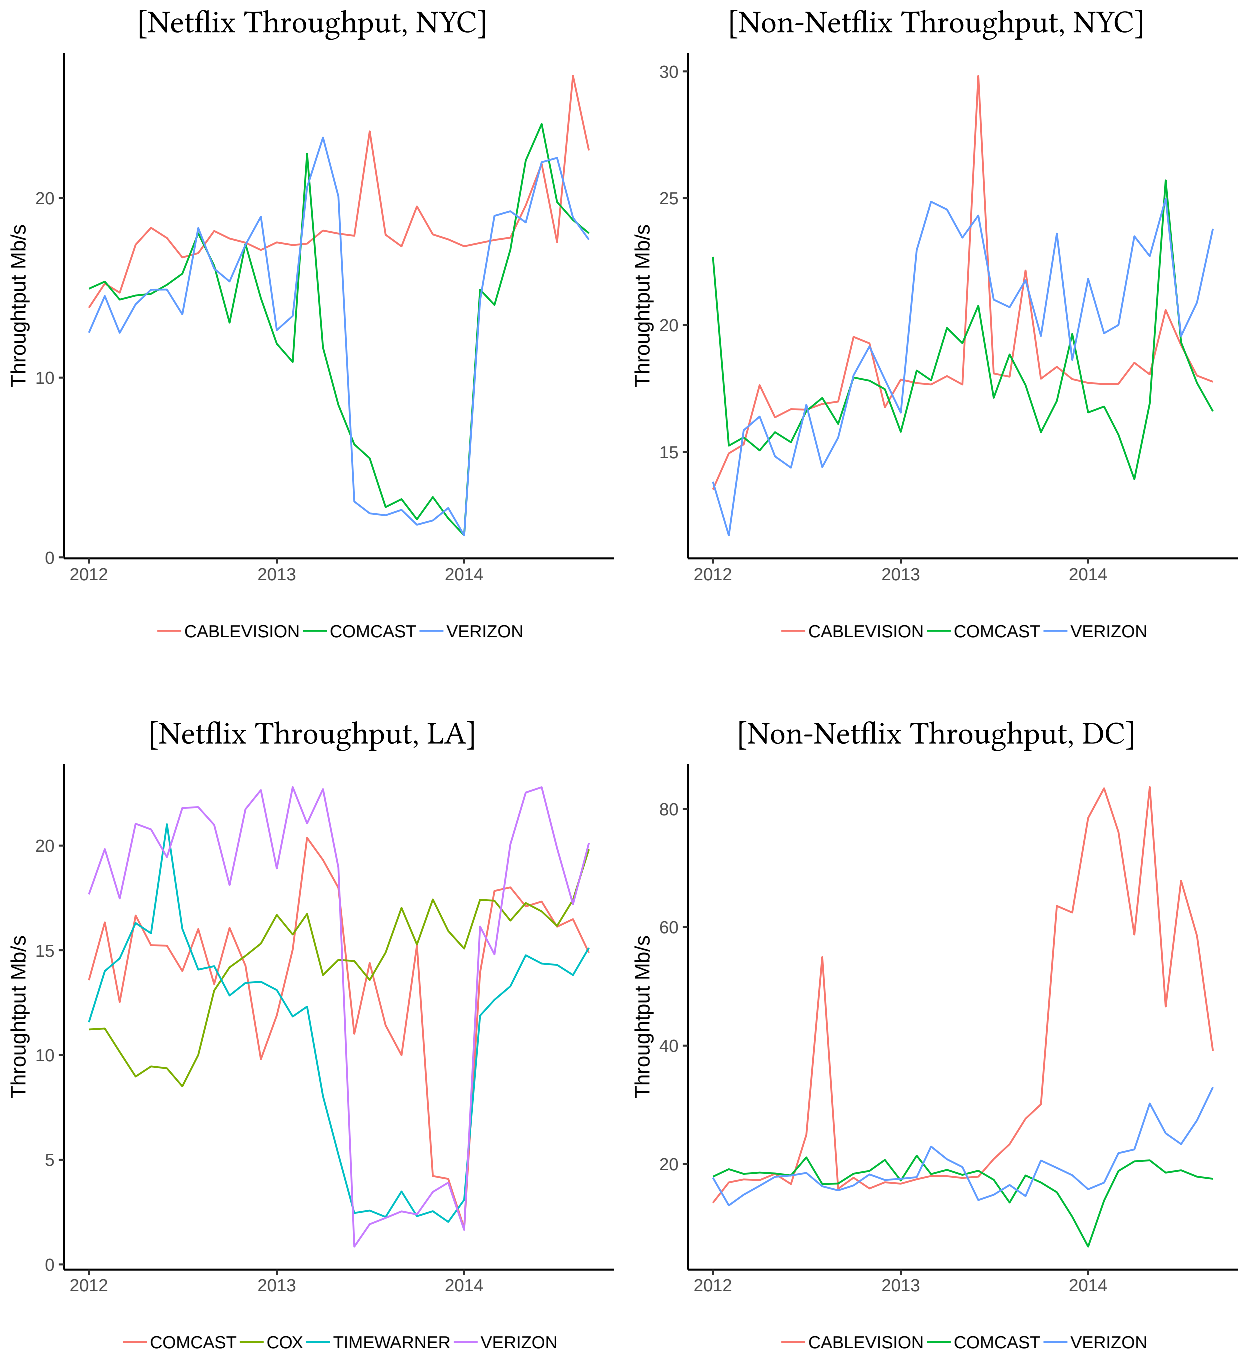

Figure B.1: ISP-Netflix Pairwise Throughput in NYC, LA and DC greater than the present discounted value of the costs." [41]

However, from CAIDA data, Netflix cancelled service with a crucial third party in late 2013. Limelight was a large CDN that Netflix had relied on to smooth delivery of its services, but after November 2013—the low point of quality degradation from Netflix's throughput data—they no longer interconnected. Meanwhile, Comcast claimed that Netflix's throughput problems could be solved if they purchased more bandwidth from transit providers—a statement reported and dismissed by Netflix.[42]

In the reduced form analysis and demand model estimates, I show that Netflix had the incentive to degrade quality of service to induce faster agreement times from ISPs. Since the marginal consumer appears to be more elastic with respect to switching ISPs than canceling Netflix, there is no incentive for ISPs to slow down traffic.

Constructing Durations

I assume that bargaining begins for all ISPs in 2013Q3 based on the sharp drop in the MLab data. The duration of disagreement is constructed in the following steps:

1. Using the Netflix throughput data through end of 2015, for each ISP for which there is data, I regress throughput on a linear time trend and code a "slowdown" dummy as 1 if throughput falls below 80% of its predicted value. This is a necessary condition for ISPs to be considered as having a lengthy slowdown, and provides candidate disagreement durations.

2. If an ISP is explicitly mentioned in Netflix's Q-10 or Petition to Deny Filings as having reached agreement or not by a certain time, I adjust the disagreement durations to reflect this (Comcast, TWC, and Verizon reach agreement in 2014Q1, while AT&T does not. No other ISPs are mentioned.)[43]

[41] Petition to deny, pg. 52, paragraph 2.

[42] Statement of Ken Florance, Vice President of Content Delivery at Netflix since 2012.

[43] In its first quarter letter to investors issued on April 21, 2014, pg. 5 paragraph 3, Netflix notes "now nearly all cable Internet households receive great quality Internet video", implying that TWC also concluded negotiations by the first quarter of 2014. From the same document, Netflix mentions the extremely poor streaming quality 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

3. If any remaining ISP appears in the CAIDA data as interconnecting with OpenConnect at a certain time, I adjust the disagreement data to reflect this timing

4. All remaining ISPs are those that did not experience a quality degradation. Netflix pursued a policy of installing Open Connect infrastructure in the networks of medium size ISPs, so I assume that these ISPs either reached agreement with Netflix immediately, or did not negotiate at all, which will have the same information content in the bargaining model.[44] CAIDA data indicates that some small ISPs reached agreement earlier—for instance, RCN interconnects with Open Connect in late 2012. I assume that ISPs that agree immediately do so in the first period of bargaining in 2013Q3.

Note that although I use the ISP Mediacom in the demand estimation, I drop it in the supply estimation. This is based on discussions at the FCC that Mediacom shares a substantial portion of its backbone infrastructure with AT&T, and therefore did not have the ability to bargain independently.

B.2 Plan characteristic data

ISP plan menus are sourced from the FCC Urban Rate Survey and Open Connectivity Database. 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, which often include plan features.

B.3 Data Availability

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.

[44] "if an ISP has an individual market area serving a population of at least 100,000 subscribers, Netflix will install Open Connect appliances at that location at no charge to the ISP.", pg.49, paragraph 2, Netflix Petition to Deny.

Table B.1: ISP Data Availability Matrix

ISP	Pricing	Market	Shares	Microdata	Bargaining
Acd.Net	✓				
Alaska Communications	✓	✓			✓
Armstrong	✓	✓			✓
AT&T	✓	✓	✓	✓	✓
Atlantic Broadband	✓	✓			✓
Bendbroadband	✓	✓			✓
Bresnan	✓	✓	✓		
Brighthouse	✓	✓	✓	✓	✓
Broadstripe	✓	✓	✓		
Buckeye	✓	✓			✓
Cable One	✓	✓	✓		✓
Cablelynx	✓				
Cablevision	✓	✓	✓	✓	✓
Cde Lightband	✓	✓			✓
Centurylink	✓	✓	✓	✓	✓
Charter	✓	✓	✓	✓	✓
Cincinnati Bell	✓	✓	✓		✓
Clearwire	✓			✓	
Comcast	✓	✓	✓	✓	✓
Comporium	✓	✓			✓
Consolidated	✓	✓	✓		✓
Cox	✓	✓	✓	✓	✓
Earthlink	✓	✓	✓		✓
EPB	✓	✓			✓
Fairpoint	✓	✓	✓		✓
Fidelity	✓	✓			✓
Frontier	✓	✓	✓	✓	✓
GCI	✓	✓			✓
Google Fiber	✓	✓			✓
Grande	✓	✓			✓
GVTC	✓	✓			✓
GWI	✓	✓			✓
Hargray	✓	✓			✓
Hawaiian Telcom	✓	✓	✓		✓
HTC	✓	✓			✓
Hughes	✓	✓	✓		
Insight	✓	✓	✓		
Knology	✓	✓	✓		
Lumos	✓	✓			✓
MCTV	✓	✓			✓
Mediacom	✓	✓	✓	✓	✓
Metrocast	✓	✓			✓
Metronet	✓	✓			✓
Midcontinent	✓	✓			✓
Newwave	✓	✓			✓
North State	✓	✓			✓
Qwest	✓	✓	✓		
RCN	✓	✓	✓	✓	✓
Shentel	✓	✓			✓
Sonic.net	✓	✓			✓
Suddenlink	✓	✓	✓		✓
Surewest	✓	✓	✓		
TDS Telecom	✓	✓	✓	✓	✓
TWC	✓	✓	✓	✓	✓
Veracity	✓	✓			✓
Verizon	✓	✓	✓	✓	✓
Wave Broadband	✓	✓		✓	✓
Wideopenwest	✓	✓	✓		✓
Wildblue	✓	✓	✓	✓	
Windstream	✓	✓	✓	✓	✓

C ISP switching and upgrading in microdata

Table C.1: Consumer switching behaviour

	Dependent variable:					
	Switch ISPs		Upgrade Plan		Downgrade Plan	
	(1)	(2)	(3)	(4)	(5)	(6)
Disagree	0.002^{*}	0.002^{*}	-0.008^{*}	-0.005	-0.001	-0.001
	(0.001)	(0.001)	(0.005)	(0.005)	(0.002)	(0.002)
$\Delta\log(p_{min})$	-0.004^{*}		-0.004		-0.001	
	(0.002)		(0.011)		(0.005)	
$\Delta\log(q_{min})$	0.0002		-0.0001		-0.006^{**}	
	(0.001)		(0.004)		(0.002)	
$\Delta\log(\bar{p})$		-0.003		-0.088^{**}		0.001
		(0.004)		(0.018)		(0.006)
$\Delta\log(\bar{q})$		-0.001		0.020^{**}		-0.001
		(0.001)		(0.005)		(0.002)
Household FE	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	33,955	33,955	33,955	33,955	33,955	33,955
R^2	0.001	0.001	0.008	0.010	0.002	0.002

Note: p_{min} is the price of an ISP's entry-level plan in that quarter, while $\bar{p}$ is its mean plan price. Similar definitions hold for download speed. Standard errors are clustered at the household level, 5971 clusters. p<0.1;**p<0.05

Greater ISP switching during the slowdown

In the reduced form section I find that households are switching away from affected ISPs during the slowdown by analyzing aggregate subscriber data. I show this relationship holds in the

Measuring Broadband America household microdata by running the following specification:

$$\Delta \text{ISP}_{ift} = \beta_1 \text{Disagree}_{ft} + \beta_2 \Delta \log(p_{ft}) + \beta_3 \Delta \log(q_{ft}) al\bar{p}ha_t + \Delta \epsilon_{ift} \tag{17}$$

The specification is a first-differenced fixed effect regression at the level of a household-quarter. Disagree_{ft} does not enter as a difference, reflecting the escalating effect of the slowdown on consumers' utility. The dependent variable ΔISP_{ft} will take a value of 1 if a consumer switches ISPs between quarters $t-1$ and t. Results are reported in columns (1) and (2) of Table C.1, for different measures of ISP-level prices p_{ft} and download speeds q_{ft}. The coefficient β_1 is estimated positive and significant (at the 10% level), suggesting that in the microdata households at ISPs affected by the Netflix slowdown did switch ISPs at a higher rate. The predicted 0.2% increase in a household's probability of leaving an affected ISP compared to a non-affected ISP is slightly smaller than the reduced 0.5% growth rate at estimated ISPs in the subscriber level data. This discrepancy is expected given that some of the cost to ISPs comes not from households leaving, but from households not signing up in the first place.

No plan upgrading during the slowdown

The demand specification in the structural model implies that households do not benefit from trying to upgrade their plan speed to mitigate the effect of the slowdown, as there is no interaction between the slowdown dummy and plan speed. I verify that households at affected ISPs do not upgrade at a faster rate than households at non-affected ISPs with the following specification:

$$\Delta \text{plan}_{ift} = \beta_1 \text{Disagree}_{ft} + \beta_2 \Delta \log(p_{ft}) + \beta_3 \Delta \log(q_{ft}) al\bar{p}ha_t + \Delta \epsilon_{ift} \tag{18}$$

where Δplan_{ift} takes a value of 1 if a household *increases* its plan speed in specifications (3) and (4) of Table C.1, and a value of 1 if a household *decreases* its plan speed in (5) and (6). The expectation is that none of β_1 will be significant. The remainder of Equation 18 is as in Equation 17. Results are reported in Table **??** and while the coefficient is negative and significant in column (3), there is no evidence that households differentially *upgraded* their plan speeds during the slowdown at affected ISPs.

D Demand Model

D.1 Demand Estimation

The demand estimation algorithm closely follows Berry, Levinsohn, and Pakes (1995), see Algorithm 1.

D.2 Imputating Mean Utility for Omitted ISPs

The model parameters $\boldsymbol{\theta}^d$ are fitted with the 30 ISPs for which I have data on shares. However, Netflix bargained with the 49 ISPs for which I have choice set data and which they have reported in their throughput data. To compute the revenue elasticities due to disagreement for the remaining 19 ISPs using the estimated parameters, I must make an assumption on how to recover mean utilities and ISP-specific time trends for these omitted ISPs.

I assume that each of the estimated fixed effects and time trends for the observed 30 ISPs is a linear function of observables; in particular, technology type and the number of residential units in network in the second quarter of 2013. This restricts the number of ISPs to 24 (since 7 ISPs are absorbed into larger competitors by then.) Output is provided below.

D.3 Marginal Cost Recovery

I compute the full set of plan-ISP-quarter level elasticities numerically, then take a weighted average (using plan shares as weights) within an ISP to recover column (1) in Table D.2.

Denote by ε_{jft} the price elasticity of demand for plan j offered by ISP f at time t, $\underline{\varepsilon}_{ft}$ as the elasticity of the entry-level plan, and $\underline{\mathrm{p}}_{ft}$ as the price of the entry-level plan. I then recover the ISP-quarter marginal cost as

$$mc_{ft} = \underline{\mathrm{p}}_{ft}\left(1 - \frac{1}{\underline{\varepsilon}_{ft}}\right).$$

This number is reported as column (2) in Table D.2. When the elasticity would imply a negative marginal cost, I instead set the marginal cost to zero.

Algorithm 1 Demand Model Estimation Algorithm

1: Randomly sample from normal distribution $(\nu_r)_{r=1}^{R}$ in each market m, and pick inner tolerance $\epsilon^{inner} > 0$
2: Pick $\hat{\theta}^d_{nl}$
3: Pick a vector of mean utilities $\boldsymbol{\delta}^{(0)} \equiv (\delta^{(a)}_{ft})_{f,t}$

- For any $\boldsymbol{\delta}^{(a)}$ for $a \in \mathbf{N} \cup \{0\}$, compute the model's predicted shares $\hat{\boldsymbol{s}} = (\hat{s}_{ft}(\hat{\theta}^d_{nl})_{f,t}$ as

$$\sum_{m|f \in \mathcal{F}_{mt}} w_{mt} \frac{1}{R} \sum_r \frac{\delta^a_{ft} + \hat{\lambda}_{rfmt}(\hat{\theta}^d_{nl}, \nu_r)}{1 + \sum_{f \in \mathcal{F}_{mt}} \exp(\delta^a_{ft} + \hat{\lambda}_{rfmt}(\hat{\theta}^d_{nl}, \nu_r))}$$

where

$$\hat{\lambda}_{rfmt}(\hat{\theta}^d_{nl}, \nu_r) = \max_{j \in \mathcal{J}_{fmt}} \{-\exp(\hat{\alpha}_p + \nu^p_r \hat{\sigma}^p) \log(p_{jfmt}) + \exp(\hat{\alpha}^q) \log(q_{jfmt})\}$$

- Update mean utilities using the contraction mapping

$$\boldsymbol{\delta}^{(a+1)} = \log \boldsymbol{s} - \log \hat{\boldsymbol{s}}(\hat{\theta}^d_{nl}, \boldsymbol{\delta}(\hat{\theta}^d_{nl})) + \boldsymbol{\delta}^{(a)}(\hat{\theta}^d_{nl})$$

- Iterate until $\max \left| \boldsymbol{\delta}^{(a+1)}(\hat{\theta}^d_{nl}) - \boldsymbol{\delta}^{(a)}(\hat{\theta}^d_{nl}) \right| < \epsilon^{inner}$, recover converged $\boldsymbol{\delta}(\hat{\theta}^d_{nl})$

4: Concentrate out the linear parameters as in Nevo (2000) to recover $\boldsymbol{\xi}(\hat{\theta}^d_{nl})$

$$\delta(\hat{\theta}^d_{nl}) = \boldsymbol{X}\boldsymbol{\beta} + \boldsymbol{\xi}$$

5: Using the instruments $\boldsymbol{Z}_{ft}$, form two sets of sample moments as in Berry, Levinsohn, and Pakes (1995) and Scherbakov (2015) respectively. Call the stacked vector of these moments $\boldsymbol{g}(\theta^d_{nl}, \boldsymbol{X}, \boldsymbol{Z})$

$$\frac{1}{N \cdot T} \sum_{f,t} \boldsymbol{Z}_{ft} \xi_{ft}(\hat{\theta}^d_{nl})$$

$$\frac{1}{\cdot} \sum_{j,f,t} \boldsymbol{Z}_{ft}(s_{jft} - \hat{s}_{jft}(\hat{\theta}^d_{nl}))$$

6: For a weighting matrix W^0, form the GMM objective function $g'(\hat{\theta}^d_{nl}) W g(\hat{\theta}^d_{nl})$
7: Iterate 2-6 until convergence, $\theta^{d,*}_{nl}(W^0)$
8: Form W^1 as $\boldsymbol{g}(\hat{\theta}^{d,*}_{nl})\boldsymbol{g}'(\hat{\theta}^{d,*}_{nl})$, and repeat 2-6 until convergence, recover $\theta^{d,*}_{nl}(W^1)$

Table D.1: Imputing ISP fixed effect and time trend parameters

	Dependent variable:	
	feval	ttrends
	(1)	(2)
log(msize)	0.116	0.0002
	(0.102)	(0.003)
techDSL	-0.174	-0.010
	(0.277)	(0.009)
techSATELLITE	-5.291***	0.018
	(0.580)	(0.020)
Observations	24	24
F Statistic (df = 3; 20)	26.988***	1.060
Note:	*p$<$0.05; **p$<$0.01; ***p$<$0.005	

Finally, for average margins, I divide $p_{jft} - mc_{ft}$ by p_{jft} for each ISP plan in each quarter, then take a weighted average (using plan shares as weights) to recover the average ISP margin.

Table D.2: ISP-level elasticities, marginal costs, and margins

ISP	Elasticity	Marginal Cost (USD)	Margin $\sum_{j\in f} s_{jf}\left(\frac{p_{jf}-c_f}{p_{jf}}\right)$
Alaska Communications	12.338	67.311	0.400
Armstrong	4.279	14.877	0.523
AT&T	18.892	34.355	0.171
Atlantic Broadband	1.520	48.584	0.089
Bendbroadband	8.611	17.618	0.688
Brighthouse	9.649	38.806	0.229
Buckeye	0.101	0	1
Cable One	0.684	0.537	0.990
Cablevision	8.310	41.126	0.222
Cde Lightband	0.060	0	1
Centurylink	3.577	19.678	0.487
Charter	5.533	29.276	0.264
Cincinnati Bell	12.516	15.323	0.640
Comcast	1.702	32.639	0.397
Comporium	2.211	0	1
Consolidated	5.899	24.537	0.349
Cox	9.476	34.675	0.305
Earthlink	2.634	3.768	0.816
EPB	2.070	54.942	0.206
Fairpoint	3.305	10.190	0.719
Fidelity	9.673	44.809	0.231
Frontier	9.336	33.136	0.181
GCI	0.046	0	1
Google Fiber	0.002	0	1
Grande	4.079	19.738	0.439
GVTC	11.756	33.972	0.179
GWI	2.584	58.744	0.075
Hargray	7.389	26.269	0.303
Hawaiian Telcom	7.115	20.105	0.518
HTC	10.680	24.963	0.387
Lumos	0.00000	0	1
MCTV	9.948	41.315	0.213
Mediacom	6.847	38.707	0.337
Metrocast	7.978	31.481	0.301
Metronet	0.024	0	1
Midcontinent	3.492	17.688	0.483
Newwave	8.731	32.559	0.260
North State	7.993	31.912	0.218
RCN	10.637	44.367	0.228
Shentel	0.001	0	1
Sonic.net	0.00000	0	1
Suddenlink	9.047	12.528	0.587
TWC	3.012	16.336	0.660
Veracity	4.955	32.006	0.381
Verizon	0.475	0	1
Wave Broadband	6.016	20.154	0.543
Wideopenwest	9.738	36.694	0.342
Windstream	0.424	15.532	0.756
Alaska Communications	4.335	32.789	0.470

Note: for computing the margin, s_{jf} is the share of plan j within ISP f, which add to 1 for each f.

E Bargaining Model

E.1 Monte Carlo

To verify whether the objective function is well behaved around the optimal parameters, I pick a θ^s and covariates, generate data from the model, and plot the objective function around the true parameter values.

I choose $\theta^s = (-1, 3, 0.8, 0.01)$, market size uniformly distributed between 0.1 and 2, agreement profits uniformly distributed between 0.4 and 0.5, and disagreement profits distributed between 0.3 and 0.4. I generate $J = 50$ bargaining durations.

I plot the value of the objective function in a range of $\theta^s_k \pm \text{abs}(\theta^s_k)$ for $k = 1, 2$ and $\exp(\log(\theta^s_k) \pm \text{abs}(\log(\theta^s_k)))$ for $k = 3, 4$. Since the last two elements of θ^s are variances, the objective function is not defined for negative values, so I search over $\log(\theta^s_k)$ for these two values and exponentiate before inputting into the model. Other parameters are held fixed at their optimal values in each plot. Output is presented in Figure E.1. The objective function is largely well behaved, but does exhibit local minima and θ^s_3 may be poorly identified. θ^s_1 and θ^s_s are likely to be well estimated however.

E.2 Estimation Algorithm

See Algorithm 2 for details.

E.3 Robustness

I observe the weighted average cost of capital (WACC) for 24 ISPs: Alaska Communications, ATT, Atlantic Broadband, BendBroadband, Cable One, Cablevision, Centurylink, Charter, Cincinnati Bell, Comcast, Consolidated, Earthlink, Fairpoint, Frontier, GCI, Google Fiber, Hawaiian Telcom, Shentel, Suddenlink, TDS Telecom, TimeWarner, Verizon, WideOpenWest, and Windstream. The remaining firms are privately held and do not report balance sheet information. For these firms, I predict their WACC based on their households passed, technology (DSL or Cable), and whether

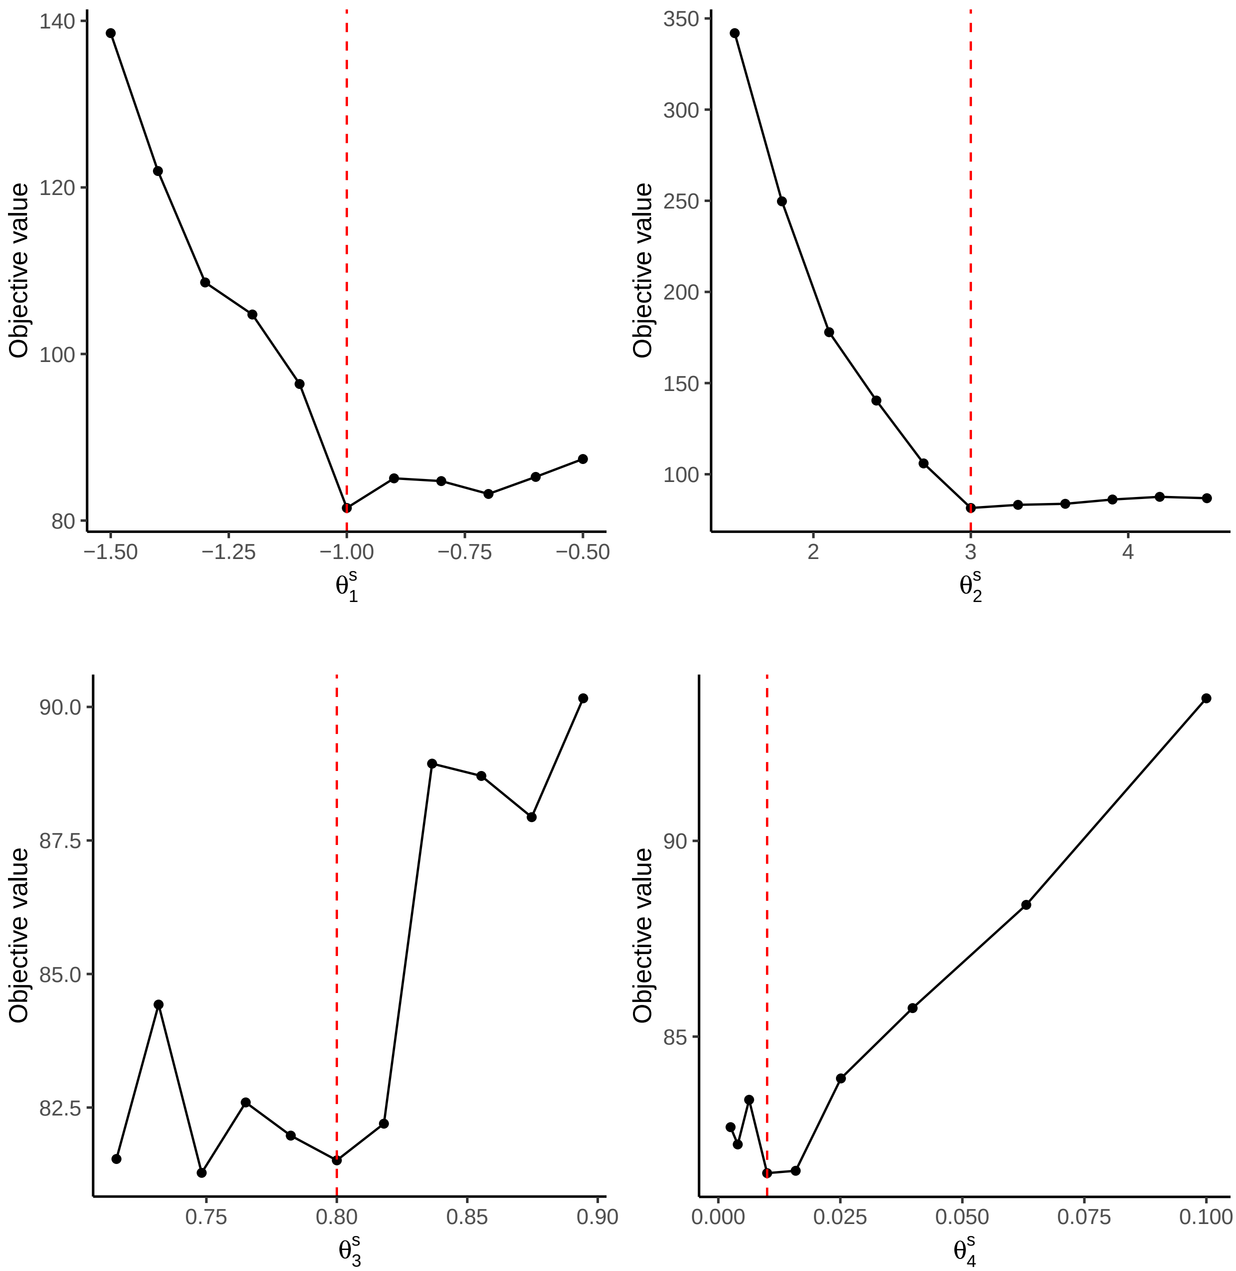

Figure E.1: Objective Function Behaviour, 50 simulated observations

Algorithm 2 Bargaining Model Estimation Algorithm

1: Choose $K \in \mathbf{N}$ gridpoints and discount rate δ
2: Pick $\hat{\theta}^s$
3: Compute $\hat{\theta}^s_f = \hat{\theta}^s \times X_f$
4: Pick state space gridpoints $\boldsymbol{b} = (b_{(k)})_{k=1}^K$
5: Compute cumulative profits if agree for f: $\Pi_{ft} \equiv \sum_{t' \geq t} \delta^{t'-t} \pi_{ft'}(a_{ft'} = 1)$
6: Initialize value functions at time $N+1$: ISP f expected value function $EV_{f,N+1}(b_{(k)}) = \Pi_{f,N+1}$, Netflix value function $V_{x,N+1}(b_{(k)}, b_{(k')}) = 0$

- $EV_{f,\cdot}$ argument is the rational belief about the maximum possible surplus of Netflix
- $V_{x,\cdot}$ first argument is ISP's belief about the maximum possible surplus of Netflix, second argument is Netflix's (unobserved) actual surplus

7: Given $EV_{f,t+1}$ and $V_{x,t+1}$, for each state $b_{(k)}$

$$b_t^*(b_{(k)}) - p_t(b_t^*(b_{(k)})) = \delta V_{x,t+1}(b_t^*(b_{(k)}), b_t^*(b_{(k)}))$$

$$V_{f,t}(b_{(k)}) = \max_{b_t^*} P(b' > b^* | b' \leq b) \cdot (p_t(b_t^*(b_{(k)}) + \Pi_t) + P(b' < b^*(b_{(k)}) | b' \leq b) \cdot \delta EV(b_t^*(b_{(k)})$$

$$EV_t(b_{(k)}) = E_{\boldsymbol{\epsilon}}\left[\max\left\{V(b_{(k)}) + \epsilon_1, \delta EV(b_{(k)}) + \epsilon_0\right\}\right] = \log\left(\exp\left(\frac{V(b_{(k)})}{\sigma_{f,\epsilon}}\right) + \exp\left(\frac{\delta EV(b_{(k)})}{\sigma_{f,\epsilon}}\right)\right)$$

$$\rho_t(b_{(k)}) = \frac{\exp\left(\frac{V_t(b_{(k)})}{\sigma_{f,\epsilon}}\right)}{\exp\left(\frac{V_t(b_{(k)})}{\sigma_{f,\epsilon}}\right) + \exp\left(\frac{EV_t(b_{(k)})}{\sigma_{f,\epsilon}}\right)}$$

$$V_{x,t}(b_{(k)}, b_{(k')}) = \rho_t(b_{(k)}) \cdot \big(\mathbf{1}[b_{(k')} \geq b_t^*(b_{(k)})] \cdot (b_{(k')} - p_t^*(b_{(k)})) + \mathbf{1}[b_{(k')} < b_t^*(b_{(k)})] \cdot \delta V_{x,t+1}(b_t^*, b_{(k')})\big) + (1 - \rho_t(b_{(k)})) \cdot \delta V_{x,t+1}(b_{(k)}, b_{(k')})$$

8: Iterate 7 backwards through time until have $p_t(b_{(k)})$, $\rho_t(b_{(k)})$, $V_{f,t}$, $EV_{f,t}$, $V_{x,t}$ $\forall t = 1, \ldots, N$
9: Iterate 3-7 for each ISP, until have complete set of optimal policies and value functions $\forall f$
10: For each f, compute the distribution of predicted agreement times $\hat{y}(\theta^s)$ based on the optimal policies and form the log likelihood

$$\ell(\boldsymbol{y}; \theta^s, \boldsymbol{w}) = \sum_f \log \mathbf{P}_f(\hat{y}_f(\theta^s) = y_f)$$

11: Iterate 2-10 until log likelihood is minimized for θ^s

they are owned by a larger firm or are part of a conglomerate (subsidiary). Estimated coefficients are reported in Table E.1. In Figure E.2, I plot the size ranking of ISPs against their actual WACC for ISPs where WACC is observed, and against the predicted WACC for all ISPs.

Table E.1: Predicting ISP-level WACC

	Dependent variable: log(WACC)
log(Mkt. Size)	−0.030
	(0.028)
DSL	−0.068
	(0.128)
Subsidiary	0.099
	(0.145)
Observations	24
R^2	0.088

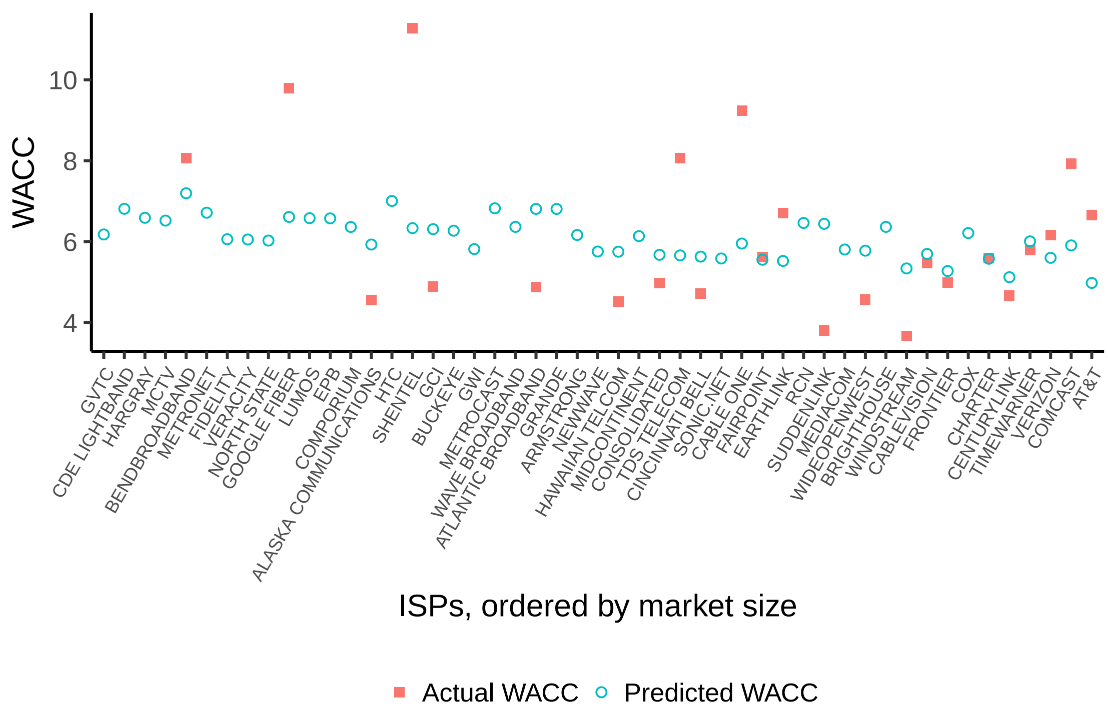

Figure E.2: Actual and Predicted WACCs