Dynamic Bargaining and Scale Effects in the Broadband Industry[*]

Daniel Goetz [†]

This Version: January 15, 2019

Abstract

This paper measures the effect of horizontal mergers in the broadband internet industry when scale matters and bargaining takes time. Using data from 2010-2014, I document that longer negotiations over interconnection fees between U.S. internet service providers (ISPs) and Netflix are associated with large reductions in the transmission quality of Netflix streaming video. I write and estimate a structural model of dynamic bargaining between ISPs and Netflix where ISP scale and market power can both play a role in determining bargaining durations. I find that ISP scale is much more important than market power to explain bargaining outcomes, and show in counterfactuals that allowing mergers between ISPs serving geographically non-overlapping markets may still harm consumer welfare and reduce Netflix's incentive to invest in internet infrastructure.

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

---

[*] This paper is a revised version of Chapter 1 of my Ph.D. thesis. Thanks to Jan de Loecker and Bo Honoré for their support on this project, as well as Jakub Kastl, Myrto Kalouptsidi, Eduardo Morales, Nick Buchholz, Sharon Traiberman, the Princeton IO student workshop, and discussants at the CEA conference for valuable feedback. I am indebted to Octavian Carare and Paul LaFontaine at the FCC for their insights. 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 size of a bargaining party is important for determining bargaining outcomes in firm-to-firm negotiations. In settings ranging from negotiations between hospitals and insurers (Gowrinsankaran, Nevo, and Town, 2015), to negotiations between content providers and cable companies (Crawford and Yurukoglu, 2012), a recent literature has found that insurers or cable systems serving larger consumer markets downstream extract more of the upstream firm's surplus than intermediaries serving smaller downstream markets, conditional on relative outside options.[1] The role of size has practical implications for analyzing mergers, which typically increase both the size of the surplus the merged firm brings to the table, as well as its changing its outside option.

In this paper, I empirically quantify the effects of ISP size and ISP outside options on bargaining outcomes with Netflix, a large provider of internet streaming video. I estimate a structural model of a dynamic bargaining game played between Netflix and ISPs in 2013 over how to split the surplus from an investment made by Netflix in content transmission infrastructure. From 2007 through 2017 a wave of ISP mergers took place, almost exclusively between ISPs serving geographically non-overlapping markets, implying that changes in size—and not ISP outside options—is first order important in merger analysis here.[2] Over the same time period total data transmitted over the internet increased from 1.7 to 23 exabytes per day, a substantial portion of which was accounted for by Netflix in the U.S. (Cisco, 2017)[3]. This paper's contribution is to provide estimates of the role of scale in bargaining; to analyze how counterfactual mergers would have affected consumer welfare and investment; and to write and estimate a new structural model of dynamic bargaining.

The data is a cross-section of bargaining durations between Netflix and a set of U.S. ISPs. Starting in 2013, Netflix video stream quality decreased at a number of ISPs for prolonged periods; this degradation coincided with bargaining over the installation of newly developed Netflix equipment in ISPs' networks. Streaming quality was restored at an ISP when agreement was

---

[1] See also Grennan (2013) for hospital negotiations with stent suppliers, and Dafny, Ho, and Lee (2016) and Collard-Wexler, Gowrisankaran, and Lee (Forthcoming) for more hospital-insurer negotiations.

[2] Charter and TimeWarner, and CenturyLink and Qwest, are two recent large examples.

[3] 1 exabyte = 1 billion gigabytes.

reached, which happened at some ISPs only after a full year of degraded Netflix quality (see Figure 1).

Since I only observe bargaining durations and not the size or split of Netflix's surplus, I estimate a structural model of households, ISPs and Netflix that maps durations and ISP characteristics into bargaining outcomes. On the demand side, consumers choose internet access plans and ISPs based on whether Netflix quality is degraded at a particular ISP or not. On the supply side, ISPs bargain with Netflix by making offers to split Netflix's ISP-specific cost-savings surplus, which is different for each ISP and which ISPs do not observe. If Netflix rejects an ISP's offer, Netflix quality at that ISP remains low, ISPs realize their outside option as subscribers may switch ISPs, and ISPs update their beliefs about the surplus. Variation in bargaining delays across ISPs is explained by differences in outside options—how much profit an ISP loses during a period of low Netflix quality from switching subscribers—and by ISP characteristics that shift the surplus distribution. I parameterize Netflix's ISP-specific surplus distribution to allow for the possibility that Netflix achieves disproportionately more cost savings at an ISP serving many households than at a smaller ISP. If variation in outside options alone can explain the observed pattern of bargaining in the data, then the surplus scale parameter will be insignificantly estimated, and mergers between non-overlapping ISPs will have no effect in the counterfactual.

In the demand estimation, I combine plausibly exogenous geographic variation in the set of available ISPs and plans to identify household preferences for price, download speed and Netflix quality in a mixed-logit framework with random choice sets using observed purchase data. I estimate an average price elasticity of 5.9, substantially higher than previous estimates for internet service (Dutz, Orszag, and Willig, 2012; Carare, McGovern, Noriega, and Schwarz, 2015) but comparable to structural estimates for cable TV (Crawford and Yurukoglu, 2012), and show that dynamic bargaining led to a consumer welfare loss of 0.5% compared to if quality had remained on trend. On average, ISPs are estimated to lose \$2.2 million in profits per quarter of disagreement with Netflix.

Using the recovered demand curve and the structural model of bargaining, I estimate significant scale effects in Netflix's surplus distribution, with larger ISPs being associated with relatively more cost savings for Netflix. Each additional household in an ISP's market increases Netflix's cost savings by $8.17, but realizing the cost savings requires a fixed per-ISP cost of $18.3 million. For ISPs with small markets the surplus distribution is thus largely negative, so they conclude bargaining quickly, and often without the cost-savings investment actually being undertaken. Although small ISPs tend to have relatively lower disagreement payoffs, that alone is not enough to explain their shorter bargaining times; scale effects in the surplus distribution are necessary to rationalize the data. The significance of the fixed cost is robust to making smaller ISPs more impatient or allowing them to be worse at bargaining.

The significantly estimated fixed per-ISP cost implies mergers between non-overlapping ISPs have an effect on upstream bargaining. In such a merger, scale increases but market power does not, as no consumer experiences a reduction in their choice set.[4] A static bargaining merger analysis with scale effects would predict no welfare loss and overall surplus gains by eliminating a duplicated fixed cost. Moreover, Netflix's bargaining share of its investment surplus would fall, since the scale effect implies Netflix's outside option has weakened relative to its agreement payoff.[5] With dynamic bargaining and scale effects, the negotiation time can also increase and Netflix may therefore retain a lower share of a more heavily discounted future surplus.

I find that for Comcast and TimeWarner—the largest cable internet providers at the time—the merger would lead to a negligible change in bargaining time, and a negligible reduction in Netflix's retained surplus share. The ISPs are large enough that even after removing a duplicated fixed cost, Netflix's surplus at the merged ISP is essentially the sum of the constituent surpluses. For some mergers between smaller ISPs (e.g., Cablevision and Metrocast), I show that although the expected amount of total surplus may increase, Netflix actually retains less of it in present-value dollar terms, due to the combination of a weaker bargaining position and longer bargaining time. For smaller merging ISPs, a combination of scale effects and dynamic bargaining costs can be enough to reduce the incentive for Netflix to invest in its cost-saving technology after a merger even though post-merger the total amount of surplus has increased.

---

[4] Comcast and TimeWarner noted this when they petitioned for a merger in 2015.

[5] See Appendix A.

1.1 Related Literature

This paper relies on estimates of consumer substitution between downstream intermediaries to identify the parameters of an upstream bargaining model, which places it in a recent literature including Crawford and Yurukoglu (2012), Grennan (2013) as well as the aforementioned Gowrinsankaran, Nevo, and Town (2015) and Dafny, Ho, and Lee (2016). These papers assume a static Nash-in-Nash bargaining framework, and I extend their insights on how consumer substitution is a key source of upstream bargaining incentives to a dynamic setting. My model assumes take-it-or-leave-it offers in the bargaining game and so does not nest the Nash-in-Nash framework; since the offers are made by the intermediaries instead of the content providers it is complementary to Ho (2009), who also assumes take-it-or-leave-it offers but made by the content providers (hospitals) in her setting.

A clear antecdent to this paper's results on scale effects of mergers is Chipty and Snyder (1999), who analyze how cable system scale affects the price paid to content providers. 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. 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 scale effects but which provides a complement to this paper's results on mergers between non-overlapping intermediaries.

There is a small but growing set of tools to empirically analyze dynamic bargaining. A complete information framework for multilateral dynamic bargaining is developed in Merlo and Wilson (1995) and Merlo and Tang (2012), where delays arise if the value of surplus to be split is stochastic and may rise over time. My model complements their analysis by allowing for delay with non-stochastic surplus and introducing downstream competition among bargaining parties. Ambrus, Chaney, and Salitsky (2016) have structurally estimated 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.[6]

There is a substantial literature on structural empirical analysis of mergers (Nevo, 2000a), and much of the focus has been on cost-savings versus market concentration. In my setting upstream interactions are crucial for understanding how this tradeoff occurs; Evans (2003), Evans (2010) and Evans and Schmalensee (2013) argue that two-sided markets, like the market for internet service, face a unique set of antitrust issues, since market power may also be exerted against the content side of the market. Empirical work including Argentesi and Filistrucchi (2007) and Chandra and Collard-Wexler (2009) focuses on the consequences of intermediary market power on content-side prices, and on exit and entry of content providers. I contribute to this literature by showing how intermediary scale—independently of market power—and an upstream dynamic bargaining model can imply new sources of welfare tradeoffs in merger analysis. In the empirical results, I provide a counterpoint to the suggestion in Becker, Carlton, and Sider (2010) that competition between ISPs can be enough to ensure positive outcomes for households, since competition plays a secondary role in mitigating the costly slowdown here.

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.

---

[6] 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.

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.[7] 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, TimeWarner, 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, 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

---

[7] 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.

Connect, and in so doing incurred a large fixed development and deployment cost.[8] 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.[9] 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.[10] Open Connect would allow Netflix to deliver service reliably and at lower cost.

By mid-2013, Netflix had not installed Open Connect in the vast majority of last-mile ISP networks, and had begun to report degraded quality of service to a number of U.S. ISPs.[11] 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 Time-Warner (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.[12] AT&T service quality is only restored later: in Netflix's April

---

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

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

[10] 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."

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

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

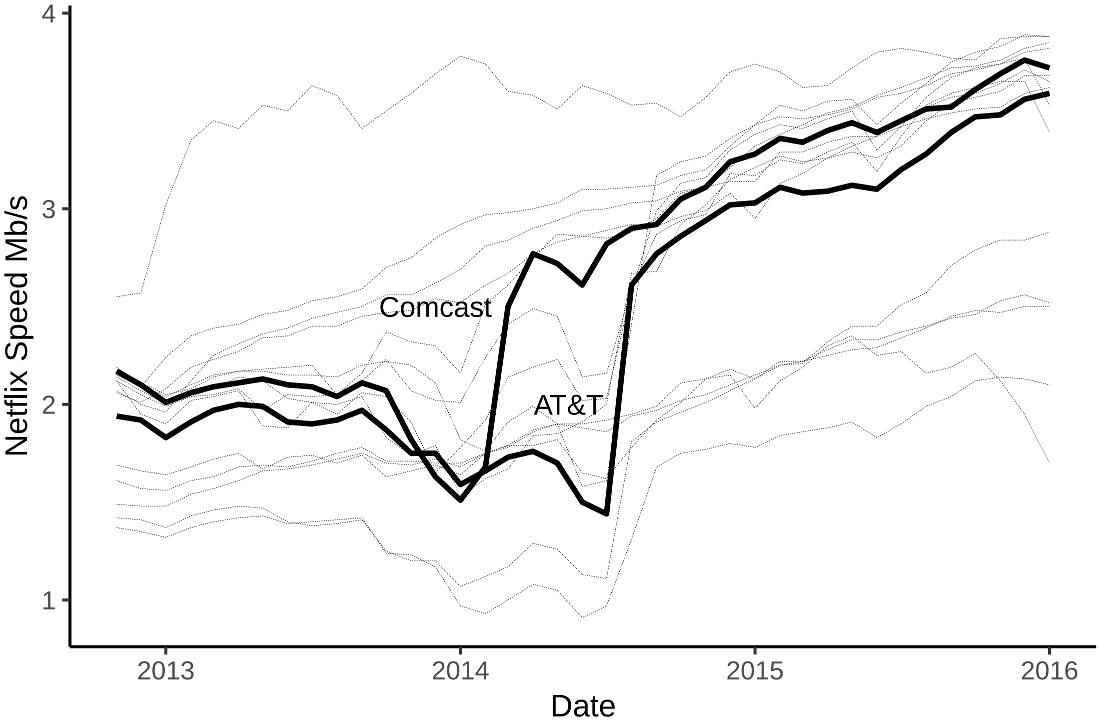

Figure 1: Average Netflix throughput to 16 U.S. ISPs

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.

2014 Q-10 filing, they state that AT&T still has not agreed to Open Connect interconnection,[13] 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.[14]

---

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

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

3 Data

Data on bargaining delays, shares, prices and consumer ISP choice sets are described below. A summary of which data is available for which ISP is provided in Appendix B.

Bargaining Delays: I gather data on the Netflix quality degradation event from several sources, including the Netflix data in Figure 1, data from MLab—an independent measurement company—and CAIDA. In addition I draw extensively from business filings. The full data construction description is provided in Appendix B, as is the argument for why the slowdowns are not simply ISP-specific technical issues but represent a slowdown of Netflix content in particular, and a discussion of which party is responsible for the slowdown. 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). In total I have durations for 49 ISPs, henceforth referred to as the core ISPs, of which 40 either agreed in 2013Q3 or did not bargain.

Market shares: Data on market shares are gathered from ISPs' quarterly and yearly earnings reports (10-Q and 10-K) which are available for all publicly traded companies in the U.S. Total internet subscriber numbers are given every quarter; combined with data on consumer choice sets, these numbers imply ISPs' market shares. The reports also contain ancillary data on mergers, which provide a source of variation in available plans. Some ISPs are privately held—e.g. Cox and RCN—in which case I use estimates of the subscription base from Leichtman Research Group. Market share movements are dominated by trends and mergers. I have quarterly data on 30 ISPs from 2010Q1 to 2014Q4 inclusive, some of which are purchased or merged before the beginning of bargaining, with 524 observations in total.

Plan characteristics: The menu of monthly prices and download speeds each ISP offers are gathered primarily 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. ISPs add or drop plans from their menu across different regions, but conditional on offering a plan it is advertised at the same price everywhere during the sample period. I have quarterly observations for 61 ISPs, including the 30 for which I have share data, as well as all 49 ISPs in the bargaining delay data.

Choice sets: Most consumers have access to only one or two wired internet options. Data on what choices are available to consumers comes from the National Broadband Map (NBBM), a government initiative with data available from 2010 through 2014 which collects information at half-yearly intervals on ISP connections at the census block level. For each census block ($<$ 1000 housing units typically), ISPs report whether they provide service to that block and their maximum advertised speed. The maximum advertised speed truncates ISP plan offerings in that block, generating geographic variation in menus within an ISP. Combined with census data on exact counts of households within each block, this data gives the weight of households across choice sets for any level of geographic aggregation. I assume that all consumers have access to satellite internet as part of their choice set. The share of consumers with access to two or more high speed (≥ 25 Mb/s) providers increases from below 20% to almost 70% during the sample. I impute 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.

Plan microdata: I construct a time series of within-ISP plan shares using data from the FCC's Measuring Broadband America (MBBA) program. The program consists of hourly wired ISP testing data for an unbalanced panel of roughly 10 000 households from 2012 through 2014, including download speeds and latency. ISP participation is voluntary, but the within-ISP sample of consumers is chosen to be representative of the distribution of consumers across that ISP's plans. Using a household's ISP and their tested download speeds, I back out which plan within an ISP's menu each household subscribes to in each quarter. There are 17 ISPs represented total, of which 16 overlap with the group of ISPs for which I observe share data.

Table 1: Correlations with Agreement Times

	(1)	(2)	(3)	(4)
Mkt. Size	1.921^{**}		1.874^{**}	1.876^{**}
	(0.314)		(0.319)	(0.323)
Num. Comp.		-0.246	-0.124	-0.121
		(0.172)	(0.133)	(0.141)
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. $^{*}p<0.1$; $^{**}p<0.05$

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:

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

I emphasize two covariates: first, the number of households an ISP can serve given its infrastructure, which captures an ISP's 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. Outside options are first order important in the empirical static bargaining literature (Collard-Wexler, Gowrisankaran, and Lee, Forthcoming), and larger firms have also been found to retain a larger share of surplus in bargaining by Crawford and Yurukoglu (2012). I also include other features of ISPs' networks including their technology (DSL/Fiber, Cable).

Results are presented in Table 1. Later agreement times are significantly positively correlated with size, and negatively correlated with the number of competitors. These regressions suggest that both size and ISP outside options may be playing a role in determining bargaining outcomes, with a potentially stronger role for size.

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 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 1 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. I first assess whether ISPs lost subscribers, then examine whether 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.[15] 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 2 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 bar-

---

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

gaining 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.[16] Regression results that control for choice sets and that use shares as the dependent variable are developed in the structural demand section.

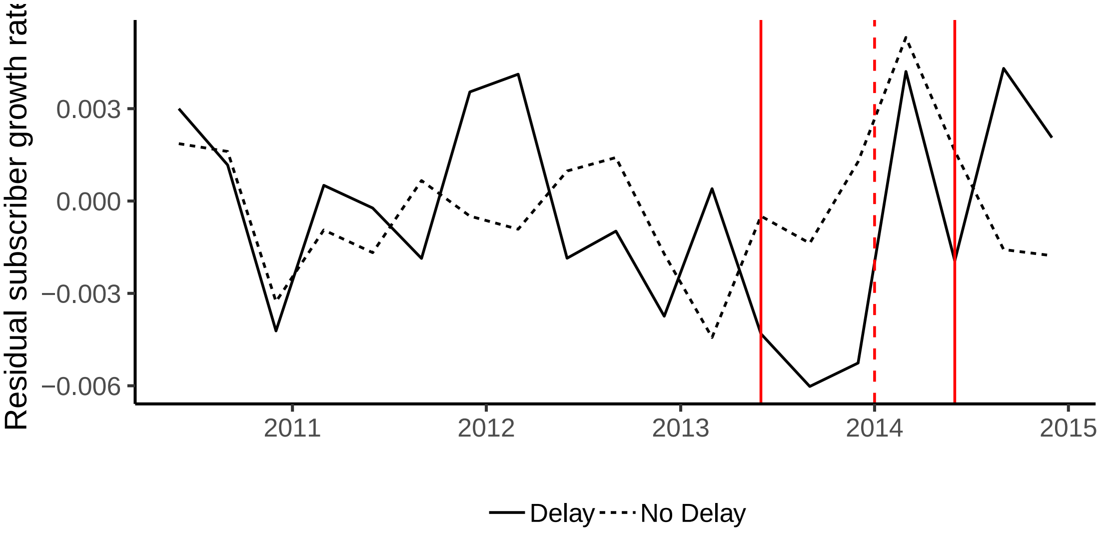

Figure 2: Residual ISP subscriber growth rates

Note: Lines are median residuals from Equation 1 by delay/no delay..

Two pieces of supplementary evidence support the claim that marginal households reacted to 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 participating in the 2013 Current Population Survey (CPS) Internet Use Supplement who owned a TV-based internet streaming device (e.g., AppleTV) were 36% more likely to switch ISPs than households without such a device, where 18.8% of CPS households reported having switched ISPs in the prior

---

[16] Regressions using a similar diff-in-diff design on microdata from the MBBA program confirm that users are more likely to switch away from affected ISPs during a slowdown, see the first two columns of Table C.1 in Appendix C.

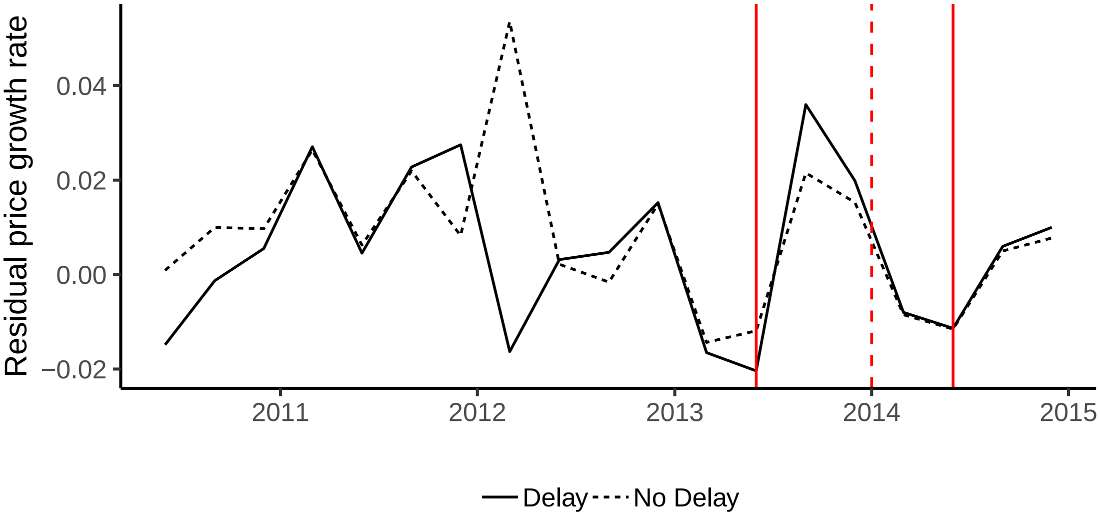

Figure 3: Residual ISP price growth rates

Note: Lines are median residuals from Equation 2 by delay/no delay.

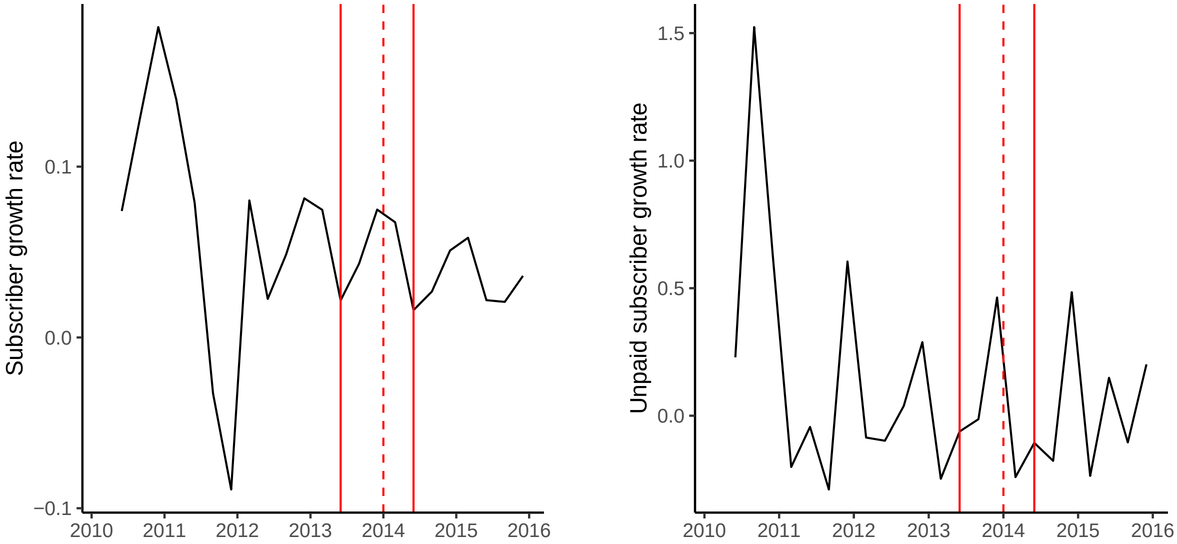

Figure 4: Raw Netflix growth rates

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

3 years.[17] Consumers who switched and own a TV-based streaming device were also 19% more likely to report that reliability was their reason for switching relative to the average switcher.

---

[17] Data available from https://cps.ipums.org/cps/.

Do ISPs also reduce their prices 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 3 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 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 4 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], or from the fact that Netflix users at this stage might be categorized as early adopters with high valuations. In my structural demand model, I will assume that demand for Netflix is completely inelastic, and that slowdowns affect consumer valuations for ISPs.

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 the outside option is normalized to zero, which implies the following form for aggregate market shares:

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

where w_{mt} describes the weight of market m at time t, such that $\sum_m w_{mt} = 1$ for each t. I

---

[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

aggregate beyond market-level shares since I only have national subscriber numbers. The model is thus similar to Goeree (2008) in that I do not observe shares at the choice-set level; however, here the choice set distribution is observed in the data, while in that paper they are estimated.

4.2 Specification

For mean utility, I assume the following functional form:

$$\delta_{ft} = \gamma \text{Disagree}_{ft} \times t + \bar{\alpha}_f + \bar{\alpha}_t + \bar{\alpha}_f^t ISP_f \times t + \xi_{ft}. \tag{5}$$

ISP and quarter fixed effects are included, as is an ISP-specific linear time trend to capture slow moving trends in ISP market shares.[20] 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.[21] 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}:

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

where p_{jfmt} and q_{jfmt} are the price and download speed of firm f's plan j in market m at time t,

---

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

[21] Including the interaction with time will improve precision in most estimates and captures the dynamics of an escalating slowdown in a reduced form way, but results are robust to removing the interaction.

respectively.[22] Heterogeneity in price disutility is captured by $\nu_i \sim \mathcal{N}(0,1)$; the form of the price coefficient follows Berry, Levinsohn, and Pakes (2004). I also provide a baseline specification for comparison where households randomly choose among plans conditional on choosing an f, so that $\lambda_{fmt} \equiv \tilde{\alpha}^p \overline{\log(p_{jfmt})} + \tilde{\alpha}^q \overline{\log(q_{jfmt})}$. In the baseline the price coefficient is unconstrained (i.e., it can be positive) to check the performance of 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 the standard Berry, Levinsohn, and Pakes (1995) inversion and use them to concentrate out the linear parameters in Equation 5 as in Nevo (2000b). 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 $E[\xi_{ft}\boldsymbol{Z}_{ft}] = 0$ for a GMM procedure. 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.

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. 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 of competitors an ISP faces in the markets it serves. Exogenous improvements in a competitor's technology, com-

---

[22] Note that in this specification, it is not possible for a consumer to reduce the disutility from the slowdown by upgrading their plan. 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.

petitor entry, and mergers 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

I present results from four models in Table 2: the first model does not instrument for price, assumes households choose plans randomly, and further assumes all households have access to all ISPs—i.e., $\mathcal{F}_{mt} \equiv \mathcal{F}_t$. The second model adds in sub-market-share level variation in choice sets; the third adds in instruments for price; and the fourth model adds in non-random choice of ISP plans.

Moving from the first to the second column, controlling for choice sets substantially increases the magnitude and precision of the estimated γ. 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 1.8, in line with the upper end of estimates (1.5) in Dutz, Orszag, and Willig (2012).[23]

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.)[24]

---

[23] The authors find price elasticity is decreasing over time, shrinking to 0.7 in 2008. My data covers 2010-2014, and may reflect the increasing use of the internet as an entertainment alternative to cable, which has much higher elasticities, see Crawford and Yurukoglu (2012).

[24] Details of the imputation of mean utility for ISPs not in the estimating sample are provided in Appendix D.1.

Table 2: 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			✓	✓
Observations	523	523	493	—

Note: Specification (1) is multinomial logit, (2) adds in sub-national choice sets, (3) instruments for price including lagged values, which explains the drop in observations. (4) adds in non-random plan choice and is estimated with two-step GMM. Standard errors are cluster robust at the ISP level. *p<0.1; **p<0.05.

The positive relationship in Figure 5 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 scale effects and market power, and allow for policy counterfactuals.

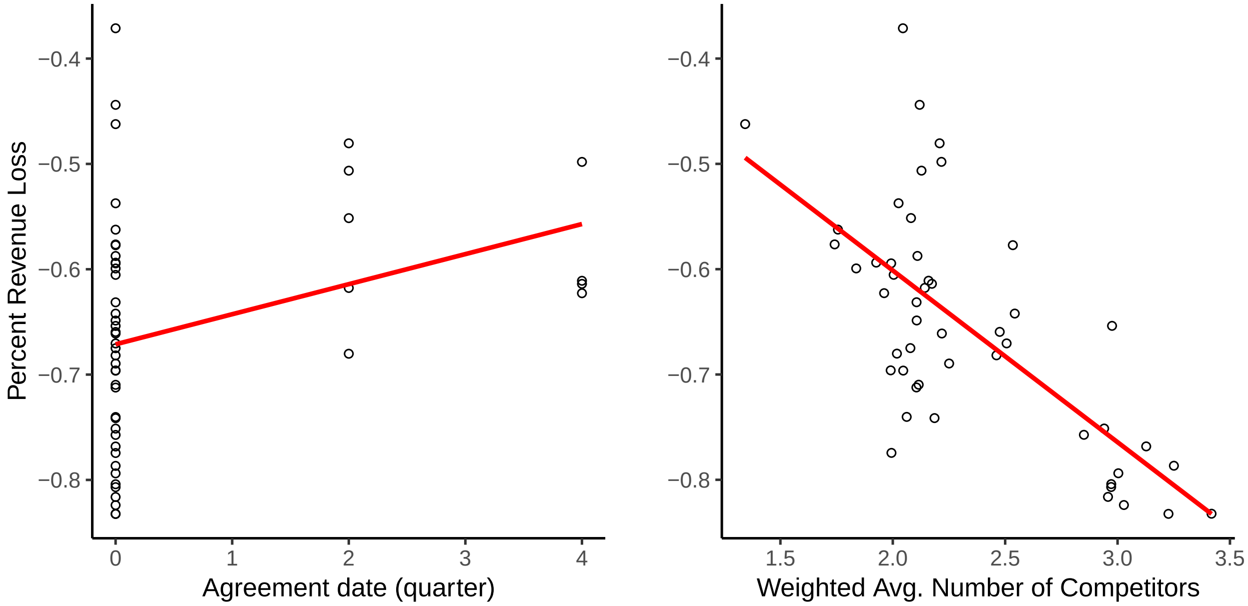

Figure 5: Revenue loss variation

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.

4.5 Marginal Costs

Since firms care about lost profits and not lost subscribers, 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 the ISP 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.[25] From the reduced form pricing results in Figure 3, I assume that the ISPs do not reprice

---

[25] A list of ISP demand elasticities (averaged across quarters), marginal costs and average margins are provided in Appendix D.2.

during the disagreement so using margins instead of prices does not change the relationships in Figure 5. 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 $\boldsymbol{a}_t \equiv (a_{1t}, \dots, a_{Ft})$, 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

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

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 scale and market power, 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:

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

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.[26] Since the profit function I estimate in Section 4 is not supermodular (see Appendix D.3), I am not guaranteed to find an equilibrium for parameter guesses away from the true parameter values.[27]

---

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

[27] In unreported Monte Carlo simulations with supermodular profit and two ISPs, it is easy to solve the 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, and the parameters estimated under the two alternative assumptions will bound the true parameter values. 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:

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

---

numerically.

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 $t+1$:

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

Setting Equation 14 to equality will imply a cutoff value of μ_{ft} that depends on beliefs and the future sequence of offers.[28] I assume that a zero fee offer by an ISP is always accepted by Netflix. 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 scale effects in the surplus distribution, to disentangle market power versus scale in bargaining outcomes. Scale effects will imply that ISPs of different sizes experience different bargaining outcomes, conditional on market power; however, in order to identify if there are actually scale effects in bargaining, the distribution family picked for the surplus distribution needs to be able to accommodate scale-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

---

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

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 scale-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 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 scale of the ISP without changing its market power—e.g., through a non-overlapping merger—will have no effect on disagreement duration, and scale will not be important to bargaining outcomes.[29]

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 $\boldsymbol{\epsilon}_{ft}$, $\sigma_{f\epsilon}$. 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.

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

---

[29] 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.

by backwards induction.[30] Since each ISP has different period profits, I solve optimal strategies and value functions in the bilateral game separately for each ISP-Netflix pair.[31]

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 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 $\ell(\theta^s,\cdot) = -\inf$ 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 scale effects pushing the mean of the surplus distribution down regardless of ISP size. 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 less market power, they will bargain more quickly as they have more to lose from periods of disagreement even with positive surplus on the table, so it is joint variation in bargaining times and market size—conditional on market power—that identifies $\boldsymbol{\lambda}$.

---

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

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

The above discussion involves identifying the mean shifters conditional on the standard deviations, σ_ζ and σ_ϵ. 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 identification relies on cross-sectional variation in delay times. 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 is important in explaining cross-ISP variation in agreement times, 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.[32]

5.3 Results

I normalize the profit π_{ft} to be in hundreds of millions of dollars, and the market size to be in hundreds of millions of households. Estimated parameters and standard errors under the alternative assumptions on subscriber profit $\underline{\pi}_{ft}$ and $\bar{\pi}_{ft}$ are presented in Table 3.

Specification (1) is the baseline result with $\bar{\pi}_{ft}$, where subscriber substitution during disagreement is computed under the assumption that all other ISPs immediately reach agreement with

---

[32] 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, although θ^s_3 might be poorly identified.

Table 3: Supply Model Estimates

	Base			Robustness	
	(1)	(2)	(3)	(4)	(5)
Surplus mean shifters (w_{1f})					
Constant (λ_1)	-0.183*	-0.190			-0.254
	0.096				
M_f (λ_2)	8.173**	8.222	35.504		7.976
	2.231				
Surplus variance shifters (w_{2f})					
M_f (σ_ζ)	1.951**	2.338	0.771		1.632
	0.338				
Bargaining shock shifters (w_{3f})					
M_f (σ_ϵ)	0.011**	0.011	0.678		0.013
	0.005				
M_f^2					1.337
N	48	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), and specification (1) uses $\bar{\pi}_{ft}$. In (4), the four largest ISPs have a higher discount rate than other ISPs. (5) allows scale effects in the idiosyncratic bargaining shock. *p$<$0.1; **p$<$0.05, standard errors computed using 50 bootstrap replications. The ISP Mediacom is dropped as an outlier.

Netflix. Since π_{ft} determines units, the estimated parameters are in hundreds of millions of dollars. The estimate of λ_1 implies fixed costs of realizing the surplus of roughly 18.3 million USD; this fixed cost could reflect Netflix's fixed cost of installing the interconnection infrastructure, or fixed costs of bargaining itself. The estimate of λ_2 implies that every additional hundred million potential households in an ISP's market increases Netflix's average surplus from bargaining with that ISP by roughly 817 million USD, or 8.17 dollars per household. I also find that the coefficient on the variance shifter is much larger than the coefficient on the bargaining shock shifter, implying that asymmetric information—and not unobserved ISP-specific idiosyncratic bargaining shocks—is the primary friction responsible for bargaining delays.

Specification (2) has the same parameterization as (1), but assumes $\pi_{ft} = \underline{\pi}_{ft}$, so that ISPs expect all competitors to disagree indefinitely. I find that the fixed cost is higher in magnitude in model (1), which follows directly from the empirical regularity in my demand estimation that $\underline{\pi}_{ft}(a_{ft} = 1)/\underline{\pi}_{ft}(a_{ft} = 0) \leq \bar{\pi}_{ft}(a_{ft} = 1)/\bar{\pi}_{ft}(a_{ft} = 0)$ for all f and t. That is, the gain from agreement is greater when all other ISPs have not agreed, and higher when all other ISPs have already agreed.[33] Since it is therefore less costly to continue bargaining under $\underline{\pi}_{ft}$, ISPs will optimally take more time, and so a larger fixed bargaining cost is required to rationalize the observed delays compared to $\bar{\pi}_{ft}$. Because the estimated parameters are so similar in specifications (1) and (2), I will use (1) going forward.[34]

I provide two checks on whether the estimates of the surplus shifters are reasonable. First, I plot predicted disagreement times from model (1) against actual disagreement times and find that the fit is fairly close to the 45 degree line. Compared to model (3) with no constant, including the scale effect λ_1 substantially improves the ability of the model to replicate the variation in bargaining times across large and small ISPs, confirming that λ_1 is significant.

Second, I can check whether 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 the value of the agreement for Netflix with Comcast in particular, based on Netflix's projected cost savings from being able to avoid paying interconnection fees for third-party transit to Comcast's network.[35] They estimated Netflix's cost-savings at around 12 million USD per year, which implies a present value of $124 million at a yearly discount rate of $\beta^4 = 0.904$ if the savings were in perpetuity. From the model, I can compute the expected fee Comcast demands based on its optimal strategy given that the third period was reached, where the expectation is over the probability of each history of offers/non-offers and rejections that would lead the bargaining game to last to the third quarter. The lump

---

[33] Note that for supermodularity this would have to hold if any individual competitor agreed versus disagreed, not if all competitors agreed versus disagreed.

[34] 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.

[35] 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 for providing this reference.

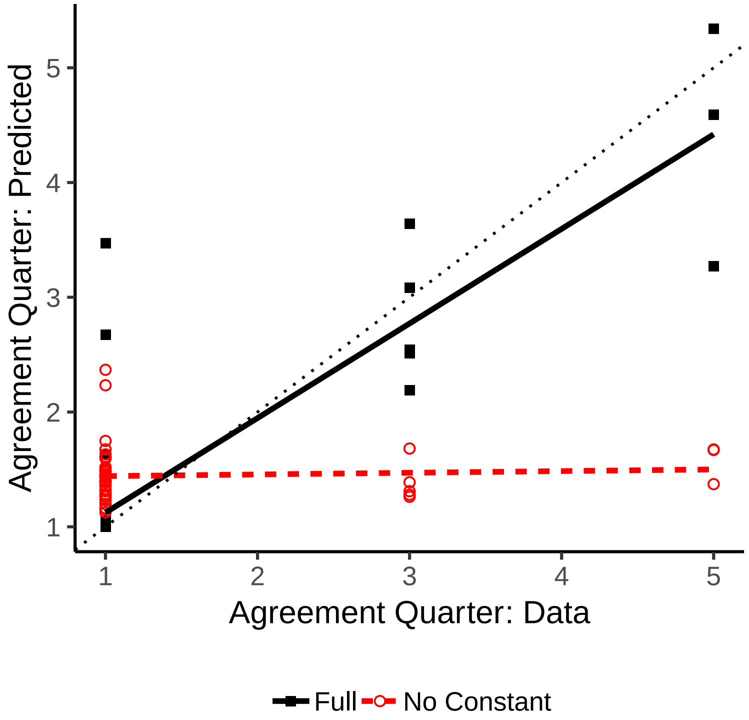

Figure 6: Supply Model Fit

Note: The full model predictions use specification (2) with $\underline{\pi}_{ft}$. There are 48 observations.

sum payment I estimate is \$155 million, which is close to the industry estimate.[36]

Scale effects, market power, delays and surplus

With the model estimated, it is possible to disentangle the role of scale effects and market power in explaining variation in bargaining delays. If scale effects are unimportant in prolonging disagreement—and the associated period of lower consumer welfare—then there is no reason not to allow mergers between non-overlapping ISPs, since market power for the merged firm should be unaffected. Moreover, the model will provide estimates of the expected surplus splits between each bargaining pair, which will have implications for the degree to which the fixed cost gives Netflix bargaining power.

---

[36] It is important to note that the industry estimate assumes Comcast extracts the entire surplus, which is not true in the model; my estimate of the surplus of \$312 million is substantially larger.

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 $\lambda_1 = 0$. These are plotted as the circles in Figure 7. For ISPs that were already bargaining, they now tend to bargain faster—the no-scale points mostly lie below their corresponding base case points in the left graph. However, these ISPs now tend to receive a higher share of the surplus, as the no-scale points mostly lie above the base case points for bargaining ISPs in the right graph. This result accords with the intuition that Netflix's bargaining position has weakened, even for these relatively large ISPs, as Netflix is more willing to accept less generous offers more quickly. For smaller ISPs that did not bargain because there was not enough chance of a positive surplus to make it worthwhile, bargaining now occurs, and indeed takes roughly as long and achieves roughly the same split as for large ISPs. That small ISPs exhibit similar bargaining behaviour to large ISPs after scale effects have been removed suggests that variations in market power cannot explain why small ISPs agree so quickly in the data—instead, scale effects play a dominant role. The reason why ISPs with less market power bargain more quickly is because they are also small, not because they have more to lose.

If the fixed cost λ_1 corresponds to an outlay by Netflix for 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 simulated above would change the amount and division of surplus.

Removing the fixed cost of installation also has an ambiguous effect on Netflix's retained surplus. Intuitively, removing the scale effect can make Netflix worse off with respect to an individual ISP in two ways. 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, if the scale effect leads to longer bargaining in expectation, the surplus Netflix retains is discounted more heavily. The model thus presents the possibility that subsidizing Neflix's investment may actually *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.

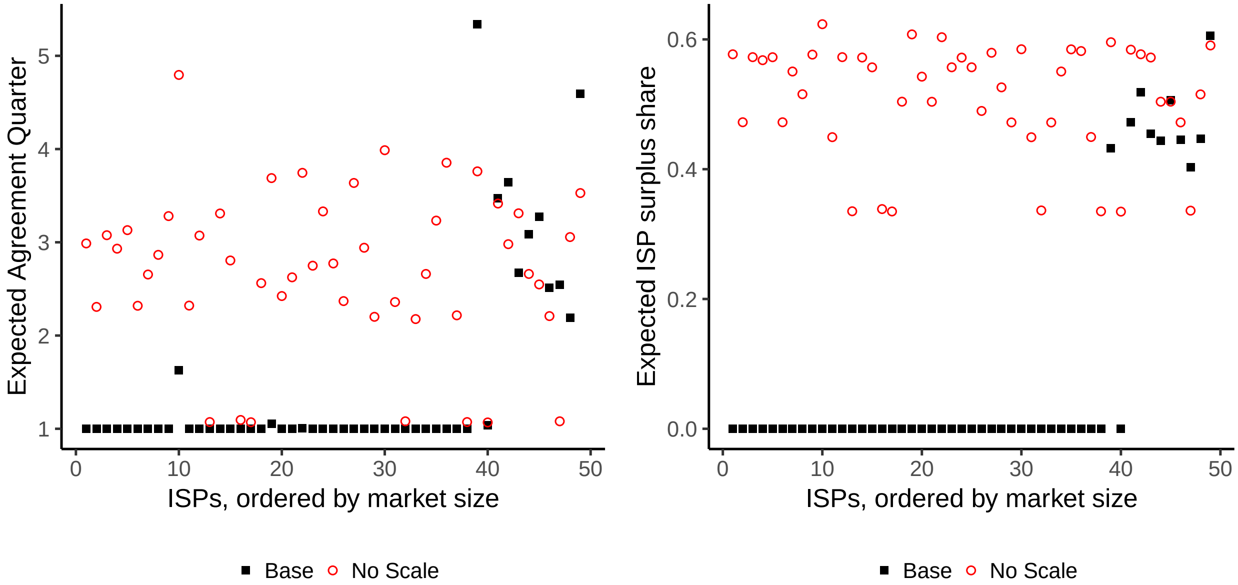

Figure 7: Scale effects, expected delays and surplus splits

Note: ISP market size for ordering purposes is taken in 2013Q3. The no-scale simulations are performed with $\lambda_1 = 0$ for the estimates from specification (1) in Table 3. The base simulations are the model predictions from specification (1).

Discounting all values to 2013Q3, the first period of bargaining, in the base case Netflix recovers 800 million USD in expectation from the 10 ISPs for which there is some positive surplus, and 866 million USD in surplus in the tax credit/no scale effect case. Thus while Netflix does recover extra surplus, much of the roughly 183 million USD in tax credits (18.3 million USD fixed cost × 10 ISPs) is appropriated by the bargaining ISPs. From the small ISPs Netflix captures 142 million USD in surplus where before there was no positive surplus to be gleaned. While it might be tempting to conclude that removing the scale effect always makes Netflix better off given the estimated parameters, in counterfactual analysis in Section 6 I show that some mergers—which remove duplicated scale effects through efficiencies—can actually reduce Netflix's expected surplus.

Lastly, for consumers, the welfare effect is ambiguous. Households with bargaining ISPs in their choice set typically benefit from a reduced period of disagreement, 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%. The net effect is a 0.029% increase in consumer welfare.

Robustness

In this section I evaluate two additional mechanisms that can explain why large ISPs take longer to bargain than smaller ISPs.

First, it may be that compared to small firms, large firms are more patient.[37] If so, their optimal strategy would be a more slowly decreasing sequence of offers and hence, longer expected bargaining durations. Large ISPs may be more patient due to lower costs of capital, leading to higher discount rates. In Table 3, I assume in the fourth model that the four largest ISPs (Comcast, AT&T, Verizon and TimeWarner) have a quarterly discount rate of 0.99 (yearly: 0.96) compared to 0.97 (yearly: 0.88) for the remaining ISPs. **To be updated**

Second, it may be that large ISPs are simply better at bargaining than small ISPs. Most estimates of bargaining games between upstream and downstream firms find that conditional on market power, large firms retain more surplus than small firms (Crawford and Yurukoglu, 2012; Grennan, 2013), perhaps due to better lawyers, more experience, or being able to credibly walk away from the bargaining table during a negotiating round. I re-estimate the model allowing for scale effects in the idiosyncratic bargaining parameter in the fifth column of Table 3. I estimate that large ISPs do have a relatively large variance on their idiosyncratic bargaining term, implying not only that they prefer to wait more often to make offers, but also that Netflix has relatively lower bargaining power since large ISPs have a credible commitment device for delay. However, scale effects in the surplus distribution are still significant and in fact larger than before.

---

[37] Collard-Wexler, Gowrisankaran, and Lee (Forthcoming) microfound the Nash-in-Nash static bargaining and show the larger bargaining power estimated for larger firms in the literature can map directly to greater patience.

6 Merger Counterfactuals

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 3, 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 positions and ISP's screening incentives.

Table 4: Cable Mergers

	Pre-Merger		Post-Merger		Netflix ΔSurplus	
ISPs	Length	ISP Share	Length	ISP Share	mil. USD	Percent
Comcast, Timewarner	2.31	0.45	2.22	0.46	4.38	1.34
Charter, Timewarner	2.69	0.44	2.12	0.46	7.22	4.32
Cablevision, Timewarner	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.

In Table 4, I simulate several bilateral mergers between Cable ISPs. In the first row, the Comcast-TimeWarner 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 TimeWarner 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 scale up.

In the second row, the Charter-TimeWarner 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 TimeWarner were both large at the time of the merger, and so the scale 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. With two smaller ISPs, as with Cablevision and Metrocast both ISPs' strategies are strongly affected by the scale effect: Netflix has a strong bargaining position with both since it does not have much to lose in case of disagreement, and ISPs face the incentive to make quick, generous offers to avoid losing subscriber revenue for very little expected surplus. When they merge, Netflix's combined surplus increases very quickly with the removal of the duplicated fixed cost, and the ISPs thereby experience a substantial increase in bargaining power.

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 scale 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 scale effect my paper allows for internally consistent counterfactuals to quantify the consequences of firm size separate from market power.

7 Conclusion

In this paper I analyze the consequences of horizontal mergers in the broadband internet industry for content provider investment and consumer welfare. Using a unique dataset of bargaining durations between Netflix and U.S. ISPs, I build and estimate a structural model of dynamic bargaining where scale and market power can both play a role in predicting bargaining outcomes. I find that scale is far more important than market power in determining the length of bargaining and the surplus split, implying that a merger even between ISPs serving geographically distinct markets can have a effect on consumer welfare and Netflix's incentive to invest in the internet backbone. Mergers between large ISPs largely have no effect, while mergers between small ISPs can decrease Netflix's retained surplus and reduce consumer welfare through longer bargaining.

How to increase investment in the infrastructure that makes the transmission of data possible is a pressing issue (Wu and Yoo, 2007; Cambini and Jiang, 2009). One idea is allowing ISPs to capture content provider surplus by providing tiered quality of service for different content types (Krämer and Wiewiorra, 2012), which has recently been implemented with the repeal of network neutrality. Through analysis of the case of Netflix, this paper raises the possibility that by providing additional bargaining power and flexibility to ISPs, content providers' incentives to invest in the internet backbone may be diminished, presenting a new tradeoff to consider in the debate over how to govern the internet going forward.

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.

Becker, G. S., D. W. Carlton, and H. S. Sider (2010): “Net Neutrality and Consumer Welfare,” *Journal of Competition Law and Economics*, 6(3), 497–519.

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.

Cambini, C., and Y. Jiang (2009): “Broadband investment and regulation: A literature review,” *Telecommunications Policy*, 33(10-11), 559–574.

Carare, O., C. McGovern, R. Noriega, and J. Schwarz (2015): “The willingness to pay for broadband of non-adopters in the US: Estimates from a multi-state survey,” *Information Economics and Policy*, 30, 19–35.

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 (Forthcoming): *Journal of Political Economy.*

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.

Dutz, M., J. Orszag, and R. Willig (2012): “The Liftoff of Consumer Benefits from the Broadband Revolution,” *Review of Network Economics*, 11(4).

Evans, D. S. (2003): “The Antitrust Economics of Multi-Sided Platform Markets,” *Yale Journal on Regulation*, 20(2), 325–381.

——— (2010): “The Web Economy, Two-Sided Markets, and Competition Policy,” *Working Paper, University of Chicago.*

Evans, D. S., and R. Schmalensee (2013): “The Antitrust Analysis of Multi-sided Platform Bussiness,” 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.

Krämer, J., and L. Wiewiorra (2012): "Network neutrality and congestion sensitive content providers: Implications for content variety, broadband investment, and regulation," *Information Systems Research*, 23(4), 1303–1321.

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. (2000a): "Mergers with Differentiated Products: the Case of the Ready-to-Eat Cereal Industry," *RAND Journal of Economics*, 31(3), 395–421.

——— (2000b): "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.

Wu, T., and C. Yoo (2007): "Keeping the Internet Neutral?: Tim Wu and Christopher Yoo Debate," *Federal Communications Law Journal*, 59, 575.

Online Appendix

A Scale Effects in Static Bargaining

In this section I show that in the bilateral Nash bargaining solution, introducing a scale effect into Netflix's surplus implies that as market size increases but ISP market power 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 TimeWarner. 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.[38] 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

---

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

[Netflix Throughput, NYC] [Non-Netflix Throughput, NYC]

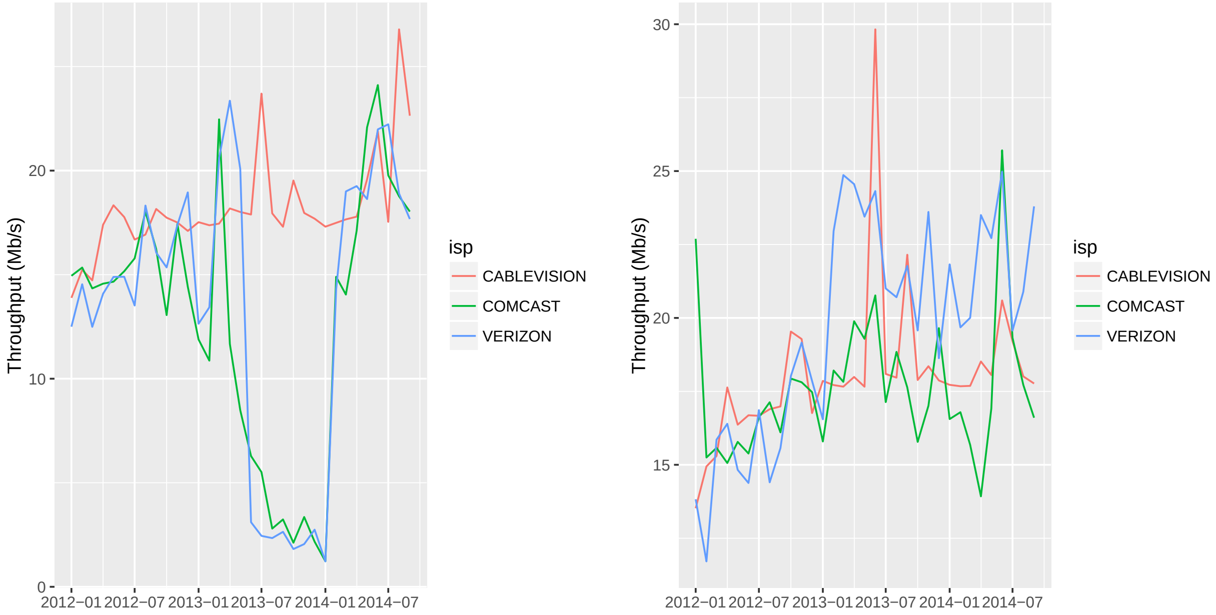

[Netflix Throughput, LA] [Non-Netflix Throughput, DC]

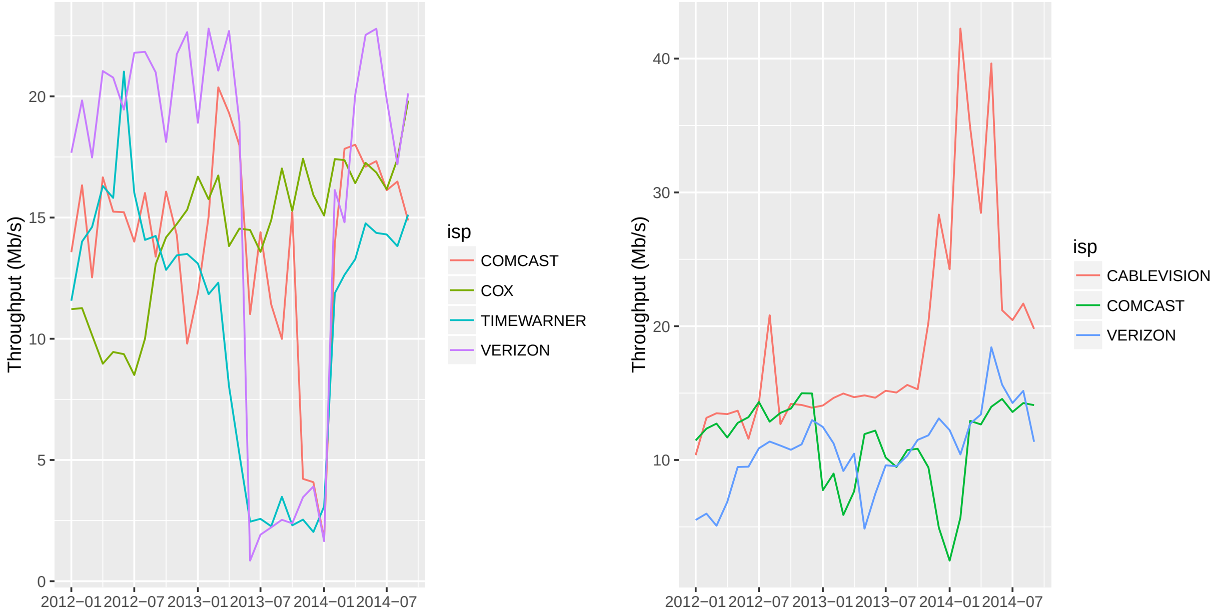

Figure B.1: ISP-Netflix Pairwise Throughput in NYC, LA and DC greater than the present discounted value of the costs." [39]

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.[40]

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, TimeWarner, and Verizon reach agreement in 2014Q1, while AT&T does not. No other ISPs are mentioned.)[41]

---

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

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

[41] 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 Time Warner Cable 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

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.[42] 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.

B.2 Data Availability

---

could be alleviated as soon as Netflix and AT&T could settle on a price. From this report it is clear that AT&T actually took longer to resolve negotiations, so that the slowdown truly indicates negotiation time and not just greater technological difficulties implementing interconnection with Open Connect.

[42] "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	✓	✓	✓	✓	✓
Timewarner	✓	✓	✓	✓	✓
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 Imputating Mean Utility for Omitted ISPs

The model parameters are fitted with the 30 ISPs for which I have data on shares. However, Netflix potentially bargained with the 49 ISPs for which I have choice set data and which they have reported in their throughput data, and even if they did not, that they chose not to bargain is informative about the parameters of the bargaining model. To compute the revenue elasticities due to disagreement for these ISPs using the estimated parameters, I require the 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.

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	

D.2 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.

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.

D.3 Submodular Profit Functions

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

ISP	Elasticity	Marginal Cost	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
Timewarner	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

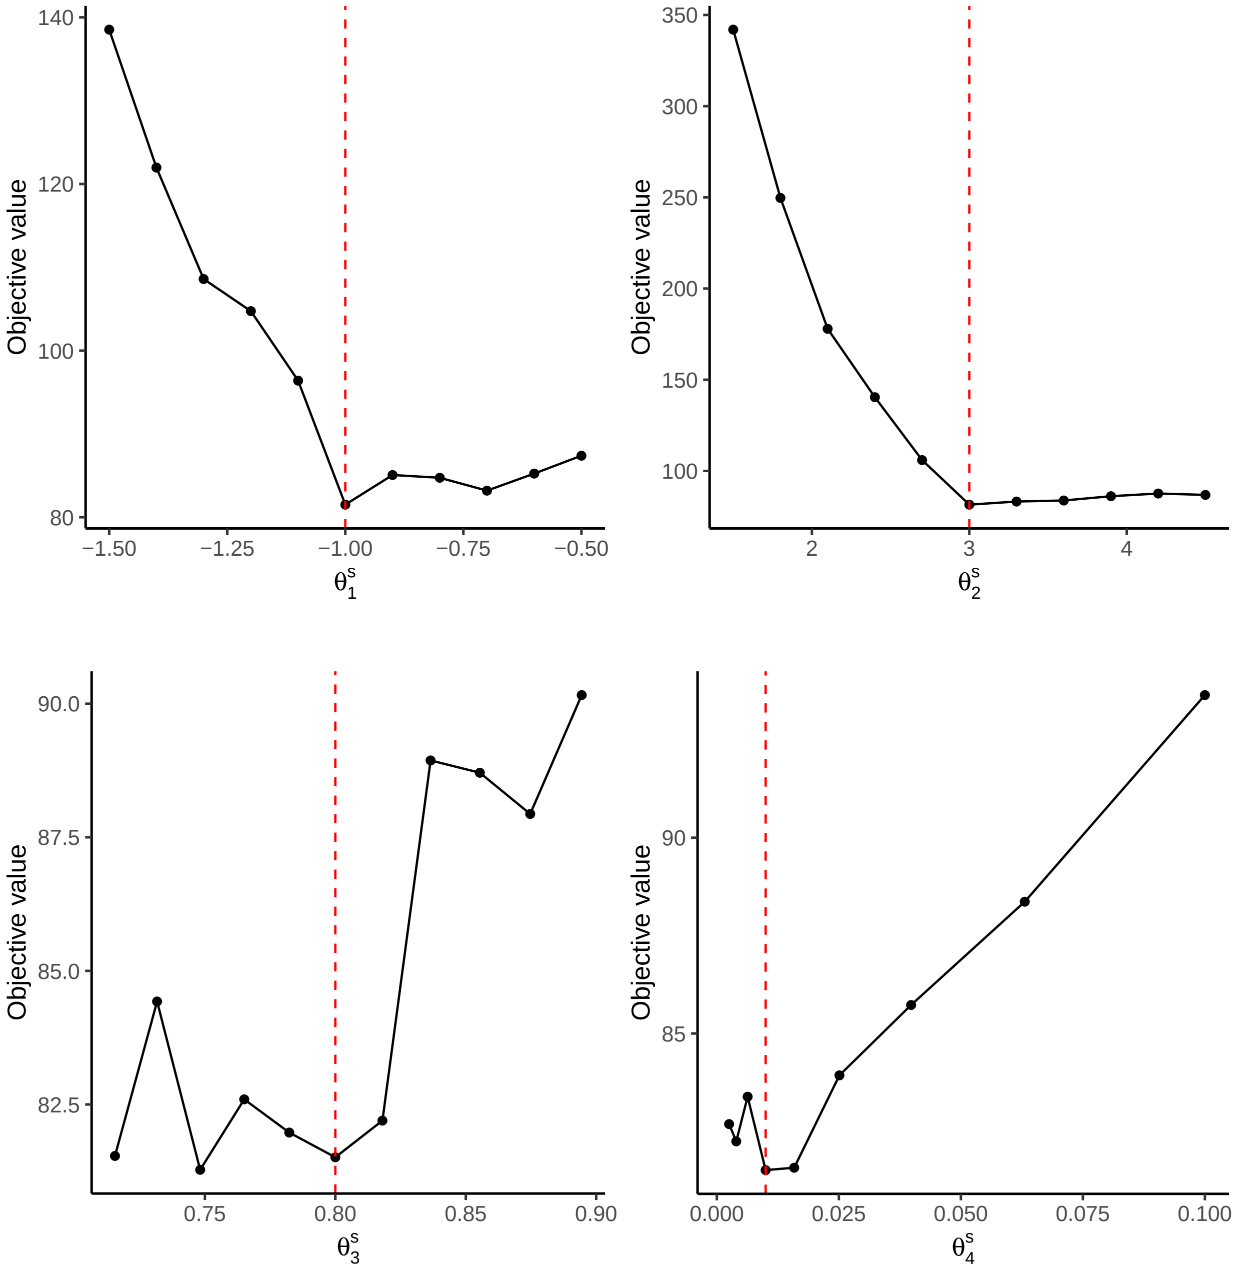

Figure E.1: Objective Function Behaviour, 50 simulated observations

Algorithm 1 Bargaining Model Estimation Algorithm

```
1: Choose $K \in \mathbf{N}$ gridpoints and discount rate $\delta$
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 $\theta^s$
```