Do subscribers of mobile networks care about Network Neutrality?*

By Augusto Espin† and Christoph Bauner‡

Draft: Thursday 1st September, 2022

Network neutrality mandates have been made out either as necessary to ensure a level playing field in online markets or, alternatively, as overly restrictive regulation preventing innovation and investment. However, there is little empirical research on the consequences of data throttling, which becomes legal without network neutrality regulations. Previous research has shown that internet service providers are applying policies to slow down the traffic from some content providers. We combine throughput levels measured for mobile ISPs in the United States with usage data to explore how sensitive users are to such practices. We find no evidence that users change their behavior when faced with throttled data rates.

I. Introduction

Network neutrality or net neutrality is a concept that was first coined by Wu (2003) and that has become one of the most discussed regulatory issues in the telecommunications industry. In Wu (2003)'s thinking, all packets traversing the internet network should be treated equally, without any blocking or prioritization regardless of the origin or content.[1]

* Author order determined by coin flip and not meant to indicate relative contribution.

† PhD Student. UMass, Amherst. Resource Economics. aespin@umass.edu.

‡ Assistant Professor: UMass, Amherst. Resource Economics. cbauner@umass.edu

[1] There is no single, universally accepted definition accepted for net neutrality. See Krämer, Wiewiorra and Weinhardt (2013) for definitions.

In economic terms, the case for net neutrality is that without it, internet service providers (ISPs) can contract with content providers to prioritize certain traffic, thus introducing inefficiencies by skewing competition at the content provider level. On the flipside, opponents of network neutrality rules have argued that allowing ISPs to accept payments for faster data transmission would provide them with additional funds for necessary investments. In addition, some applications are more sensitive to time delay than others. For instance, e-mail or browsing the web are not as time-sensitive as video or audio streams. Hence, efficient network management would entail prioritizing time-sensitive data over those less dependent on fast transmission.

In this paper, we combine information on data throughput by mobile ISPs with data on ISP market shares and content provider usage rates to empirically test consumers' reactions to data throttling, i.e., the intentional reduction of data transmission rates. While providers of fixed internet service in the U.S. have long been subject to network neutrality rules, mobile ISPs have been free to throttle data access, thus providing a natural laboratory for this question.

Regressing app usage on various measures of throttling, we find no significant effect of data throughput on app usage. This seemingly weakens the above-mentioned argument in favor of net neutrality rules. The most likely explanation, in our view, for this lack of a response is that mobile ISPs are mindful of potential consumer reactions and wary of loss of market share if they slow popular apps too much.

However, it is important to be aware of some caveats. First, and most obviously, lack of evidence of an effect is not, by itself, evidence of no effect. Second, our data is at the state level. Possibly, an effect would be visible with finer data. Third, as we are discussing firm and consumer behavior, there are important endogeneity concerns. We use an instrumental variable calculated from each ISP's coverage to deal with this problem, but it is possible that endogeneity persists, particularly in ISP behavior.

Due to the lack of easily accessible data, most of the economic literature on network neutrality is theoretical in nature, describing the impact of network neutrality regulations on market outcomes in two-sided market models using game theoretical analyses. The proposed models are analyzed with and without network neutrality rules, which are typically conceptualized as rules forbidding the ISPs to charge the content providers for prioritizing their content to the detriment of other CPs. Some authors, such as Choi, Jeon and Kim (2015) and Peitz and Schuett (2016), introduce, as an additional consideration, network congestion, and allow the ISP to engage in second-degree price discrimination based on quality.

The findings of this literature are ambiguous, depending on the exact model analyzed and often on parameter values. For instance, Economides and Hermalin (2012) find that network neutrality rules are welfare maximizing while in the models of Economides and Tåg (2012) and Jullien and Sand-Zantman (2014) the welfare consequences of net neutrality rules depend on the chosen parameters.

Peitz and Schuett (2016) find that under network neutrality there is an inefficiently large traffic volume. In Ma, Wang and Chiu (2017)'s model, abandoning net neutrality rules could solve this problem as it would provide ISPs with additional incentives to increase bandwidth. However, according to Choi and Kim (2010), enforcing net neutrality may *increase* ISPs' incentives for infrastructure investment and Gans (2015) finds that the existence of net neutrality rules may stimulate investments by content providers. Relatedly, ISPs ability to pay for additional investments may not increase when they are allowed to charge side payments as Boussion, Maillé and Tuffin (2012) argue side payments may not increase their revenues if they face competition.
Schuett (2010) and Greenstein, Peitz and Valletti (2016) provide a more comprehensive discussion of theoretical models.

The few empirical contributions to the understanding of net neutrality are mainly oriented on understanding the impact of network neutrality regulations on investments by ISPs. For instance, Hazlett and Wright (2017) evaluate the impact of the FCC's network neutrality rules of 2010 using capital investment at industry level. They find no evidence of changes in investment following the passing of these rules. Ford (2018) provides a good survey of empirical evaluations of the FCC's 2015 Open Internet Order[2] on investments in the industry, including a critique of the studies presented by the FCC. Briglauer et al. (2021) investigate the effect of net neutrality regulation on OECD countries, using industry panel data spanning 15 years and 32 countries. They find negative effects of regulation on investments. Lee and Kim (2014) use survey data of Korean internet users and computational experiments to evaluate the effect of changes in quality of service on application usage and willingness to pay of users. They find that ISPs have incentives to lower the quality of service of some content providers.

We build on research by Li et al. (2019) showing that ISPs limit the traffic speed for subscribers when accessing certain content.[3] Our aim is to understand whether subscribers are sensitive to such practice. Li et al. (2019) collect ISP-level throttling data using a crowd-sourcing scheme. We combine these data with market share estimates of the largest mobile ISPs[4] in the US and usage rates of three major applications[5], both provided by SimmonsLOCAL. We analyze the effect of throttling on usage rates making use of the variation in ISP market shares to estimate the extent to which subscribers are exposed to throttling. We employ instrumental variables based on each ISP's coverage to work around endogeneity concerns. Our findings – no significant effect of throttling on app usage – suggest that mobile ISPs may be hesitant to throttle rates too drastically.

[2]FCC (2015). In the Matter of Protecting and Promoting the Open Internet, Report and Order on Remand, Declaratory Ruling, and Order, Federal Communications Commission, FCC-15-24(March 12, 2015). 30 FCC Rcd 5601 (7)

[3]This practice is commonly referred as **throttling** in the industry

[4]AT&T, Verizon, T-Mobile, and Sprint

[5]YouTube,Netflix, and Skype

II. The Mobile Broadband Industry in the US

Since the deployment of mobile broadband in the US and the massive of smartphones around 2008, there has been a steady growth in the number of mobile connections. In 10 years, the number of connections grew tenfold, from around 30 million to 300 million. The introduction of higher speed technologies, in particular 4G LTE, allowed an important increase of available network throughput, while maintaining global compatibility. Since then, 4G LTE has increasingly been seen as standard in the U.S. and many other markets and the rollout of 5G technology is underway, promising a further increase in available throughput to consumers.

According to the FCC[6], approximately 99.8% of the American population live in areas with LTE coverage, available at a minimum speed of 5/1 Mbps.[7] According to such report, the coverage of LTE at 5/1 Mbps increased from 90% in 2013 to 99% in 2017 (table 1). At the same time, the availability of fixed terrestrial service at 25/3 Mbps reached 85.8% of the US population. However, rural area coverage is lagging behind urban centers, with fixed broadband access in the former reaching only 56.2% of the population, and mobile broadband reaching 69.3% of the population with a median speed of 10/3 Mbps and 99.1% with a median speed of 5/1 Mbps.

There are two types of operators in the U.S. market for mobile networks: Mobile Network Operators (MNOs) which own all necessary telecommunication infrastructure for managing mobile communication of their subscribers; and Mobile Virtual Network Operators (MVNOs) which resell wireless capacity of an MNO. In 2019, the US had 442.46MM mobile subscriptions reported,[8] of which approximately 62MM use MVNOs.[9] Around 86% of subscriptions were with one to

[6] 2019 Broadband Deployment Report. Bureau of Wireline Competition. Federal Communication Commission (FCC). FCC-19-44. 34 FCC Rcd 3857 (5)

[7] 5/1 Mbps means an asymmetric link with a downstream speed of 5 Mbps and an upstream speed of 1 Mbps.

[8] Source: Statista. https://www.statista.com

[9] Source: Bestmvno.com. https://bestmvno.com/mvnos

Table 1—Population Coverage with LTE

	LTE at 5/1 Mbps		LTE at 10/3 Mbps	
	2014	2017	2014	2017
United States	97.8%	99.8%	80.1%	89.0%
Rural Areas	90.2%	99.1%	70.3%	69.3%
Urban Areas	99.6%	100.0%	81.9%	92.6%
Pop. Evaluated (MM)	317.954	325.716	296.204	302.940

Data for 5/1 Mbps from Form 477. Data for 10/3 Mbps from Ookla data. 2019 Broadband Deployment Report. 34 FCC Rcd 3857 (5)

the four largest MNOs, i.e., AT&T Wireless, Sprint Corporation, T-Mobile and Verizon Wireless.[10]

The Mobile Market has a very heterogeneous offering, especially from MVNOs, which have very defined market niches, and very heterogeneous plans, including pre-paid service. However, the offerings of the four largest providers have evolved similarly and nowadays their mainstream product is what they call "unlimited plans," which are marketed as subscriptions that allow users to do unlimited texting and calls within the US as well as unlimited access to the internet. This may be a natural move since the capacity in mobile access networks has increased substantially with the deployment of LTE, and new innovative services over the internet now allow users to do calls, texting and video calls over the internet at no charge[11].

The move to unlimited plans started to be rolled out at affordable prices around 2016, with the main advantage of simplicity for subscribers. In table 2, we show the plans that were offered in 2018 under the unlimited plans. As we can see, there are limitations in the offering related mainly to video traffic, which accounts for the largest share of traffic by far.[12] In all cases, there are limitations that are imposed by providers both in download speed and monthly capacity. However,

[10]Following the merger of Sprint and T-Mobile in 2020, only three large MNOs remain in the U.S. at this time.

[11]The business models developed by these application providers do not rely on direct payment from users

[12]Source: US Telecom Industry Metrics & Trends, 2020

TABLE 2—UNLIMITED PLANS OF MNOS IN 2018

Provider	**Plan Name**	**Cost/line (USD)**				**Limitations in streaming**
		1	**2**	**3**	**4**	
Verizon	Unlimited	75	130	150	160	480p
	Unlimited (Beyond)	85	160	180	200	720p up to 15 GB/mo
	Unlimited (Above)	95	180	210	240	720p up to 20 GB/mo
AT&T	Unlimited	70	125	145	160	480p
	Unlimited & More	80	150	170	190	720p up to 15 GB/mo
T-Mobile	One	70	120	141	160	480p (in 3G)
	One plus	80	140	171	200	1080p up to 10GB (LTE)
Sprint	Unlimited basic	60	100	120	140	480p (LTE up to 500MB)
	Unlimited plus	70	120	150	180	1080p (LTE up to 15GB)

Verizon: Additional $10 for streaming @1080p only available in Above and Beyond plans. **AT&T:** Slow downs are possible due to congestion. In Unlimited & More plans slow downs starts at 22GB of usage. **T-Mobile:** Slow downs start at 50GB of usage. **Sprint:** Restrictions for games and streaming. `theverge.com`: Unlimited data plans are a mess: here's how to pick the best one (July 12, 2018) there are slight differences in the plans that may allow savvy users to choose the more convenient plan to their requirements. Importantly, we could not find evidence that providers specify technical parameters under which they limit their offering. The information provided is somewhat qualitative and confusing and limitations are vague in all cases. Measuring those parameters provides information on how much the traffic is slowed down and under which circumstances, if any.

III. Data

Our data comprises two main components: throttling rates and usage data. The former comes from Li et al. (2019) who conduct a one-year study to find if content-based traffic differentiation policies were deployed by ISPs. They employ a crowd-sourced methodology, where people could download an application and run a test designed to find if their ISP is slowing down traffic for some of the most popular applications, accumulating around 1 million measurements conducted by more than 126 thousand users across the globe. Our interest focuses on ISPs located in the United States in 2018, where around 215 thousand tests were performed. Some of the applications tested include YouTube, Netflix, Amazon Prime Video, NBC Sports, Vimeo, Spotify and Skype. These applications were selected since they usually imply higher traffic usage and therefore are more likely targets for traffic differentiation practices.

Li et al. (2019)'s tests are performed by transmitting data packets twice: once using the original data and once using obscured data that cannot be detected by the ISP's Deep Packet Inspectors (DPI) systems and thus evading traffic controls in the provider's network. Comparing the throughput between the original and the obscured data then provides an estimate of the degree of data throttling. The data collected in Li et al. (2019) is available online as raw data that can be reprocessed, but additionally, the aggregated processed results are available in their website[13], where we scraped the data. The results of these tests show that consistent differentiation is being applied to subscribers of mobile networks in the US, while there is no evidence of such behavior in fixed providers.

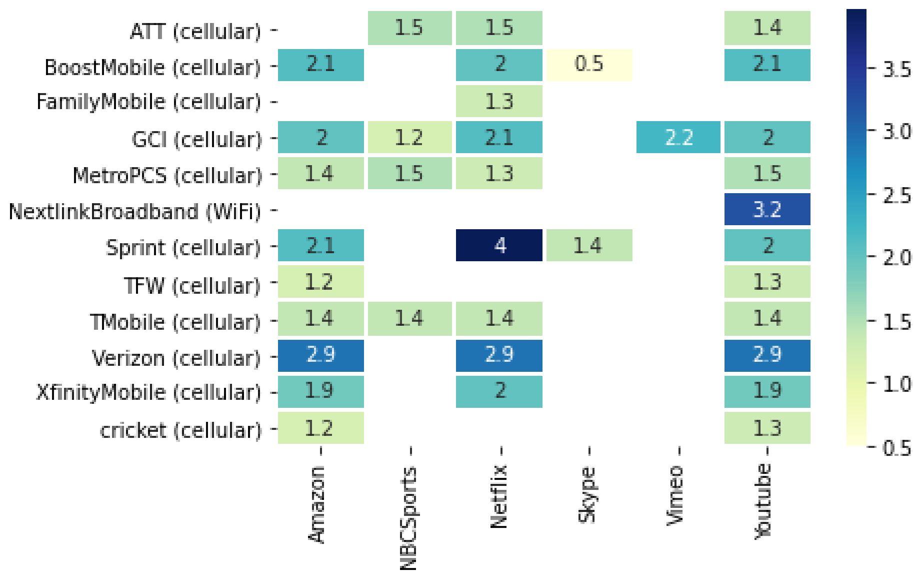

Figure 1. Average Throttling by Provider and Application in the US

Note: Figure represents throttling rates in Mbps. The throttling rate measures the data throughput for throttled traffic. An empty cell means that no throttling was detected.
Source: `https://wehe.meddle.mobi/USStats.html`

One of the most interesting findings in Li et al. (2019) is that the most common type of differentiation observed are fixed-rate bandwidth limits, known as *throttling*. In figure 1, we show a summary of the throttling rates found for the US ISPs on a set of common applications. All mobile ISPs in the dataset are exerting some level of throttling. However, the throttling rates differ significantly both across ISPs and across applications. However, the throttling rates differ significantly both across ISPs and across applications. In particular, data throughput for the same app frequently varies by more than a factor of two between the fastest throttled and the slowest throttled speeds, and for all apps providers exist

[13]`https://wehe.meddle.mobi/`

that do not throttle at all.

Our second dataset contains usage levels for mobile ISPs and applications in the US at the state level from Simmons LOCAL. Simmons LOCAL is based on survey data and uses demographic data to make predictions even at the census tract level. However, we use actual survey responses available via the crosstab feature. For these data, sufficient numbers of observations are only available at the state level. Our data obtained from Simmons LOCAL includes demographics, usage levels for applications, and usage of mobile ISPs at the state level. We will refer to the usage rate of ISPs as the market share of this ISP to simplify language. We drop states for which we observe fewer than 60 survey responses.

In our analysis, we focus on three apps for which we observe both throttling rates from Wehe and usage rates in Simmons LOCAL: Netflix, YouTube, and Skype. Figure 2 shows application usage rates across states for these apps. We see large differences across apps and, more importantly for our purposes, significant variation across states.

Figure 3 provides an overview of ISP market share by states for the four largest mobile networks in the U.S.: AT&T Wireless, Verizon Wireless, T-Mobile and Sprint, accounting for 71% of the mobile market on average according to the survey data. We focus on these providers both because some smaller providers are not represented in the Wehe data and because we are concerned about the reliability of our market share data for small providers, in particular for states where we only observe small samples. Our average market share data is roughly in line with expectations based on national numbers. For each provider, observe substantial variation in market shares both across states, a requirement for our identification strategy.

Our last dataset comes from the mobile deployment FCC's form 477, and it contains coverage by ISP computed at US Census Bureau's block level using FCC's actual area methodology[14].

[14]FCC releases data on mobile broadband deployment as of December 31, 2015 collected through FCC

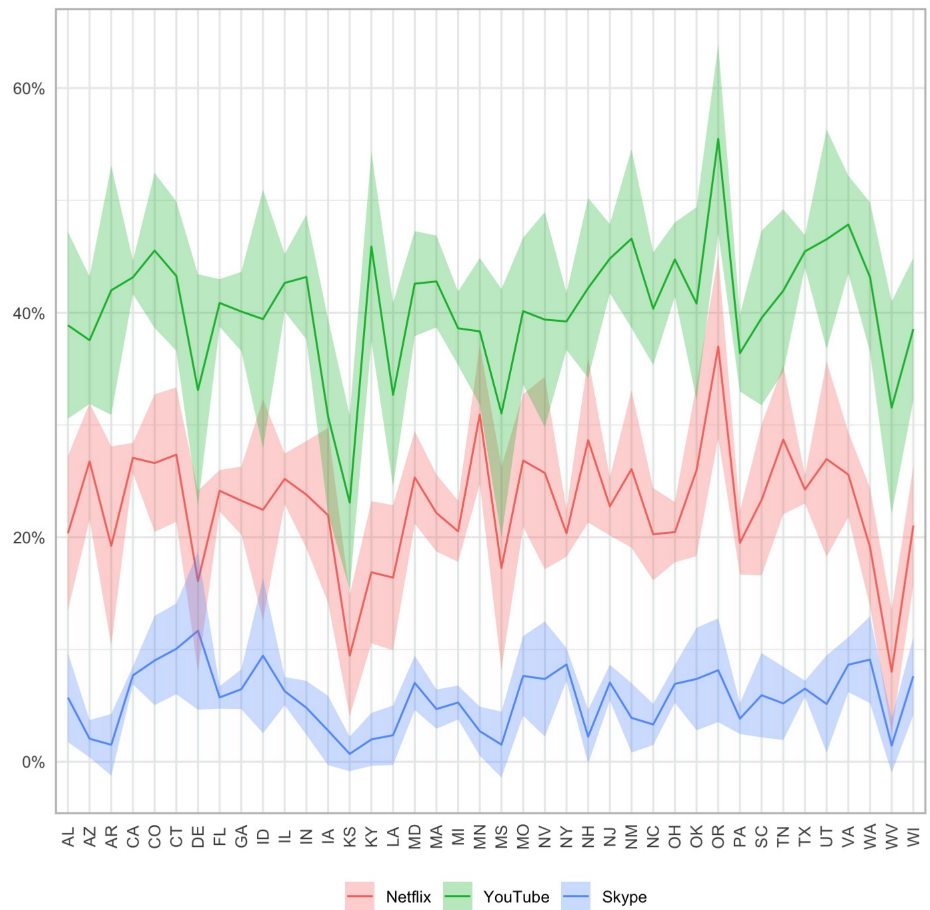

Figure 2. Application usage per State

Note: States with fewer than 60 survey responses and with missing data are omitted.
Source: Simmons Local Insights

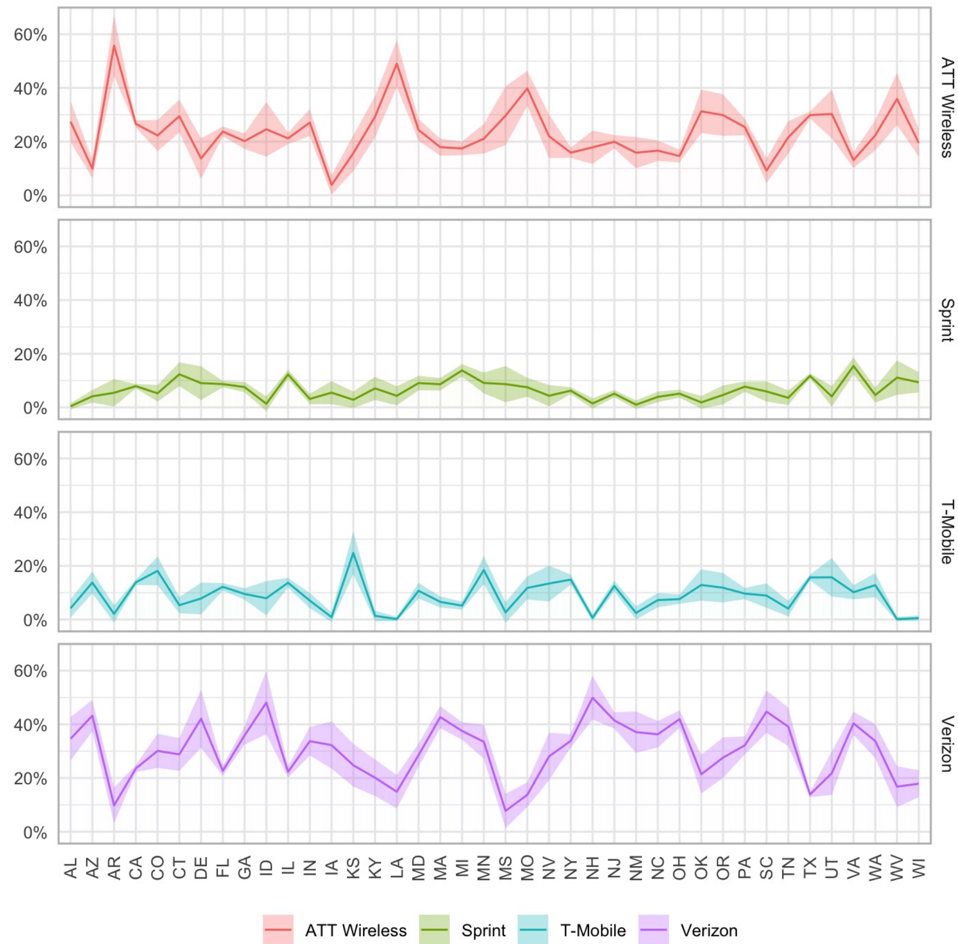

Figure 3. Mobile Provider's Market Share per State

Note: States with fewer than 60 survey responses and with missing data are omitted.

Source: Simmons Local Insights

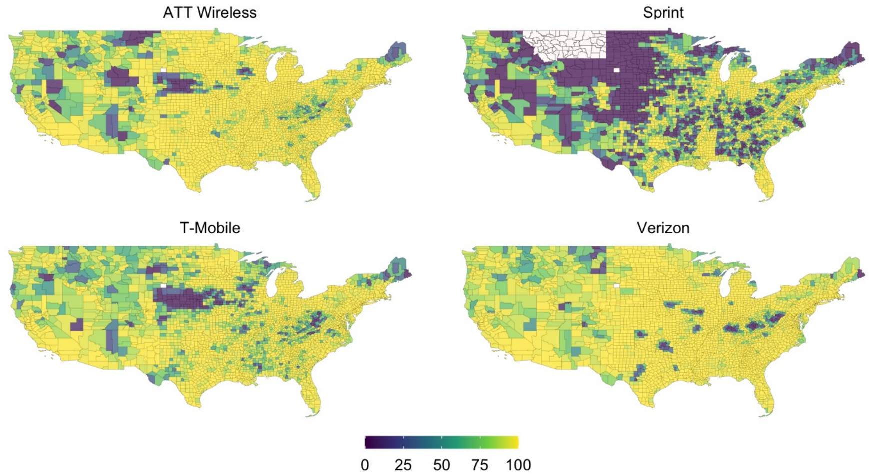

Figure 4. Geographic 4G coverage at County level in 2018

Source: FCC Mobile Deployment Form 477 (https://www.fcc.gov/mobile-deployment-form-477-data)

The dataset contains computed coverage at each block for each technology available at any given year. Since our interest is in broadband, we focus on 4G technologies. All ISP's networks had different deployment schedules, because they started from non-compatible 3G technologies. Because AT&T Wireless and T-Mobile used GSM technology, they deployed HSPA+ before deploying LTE, whereas Sprint and Verizon had CDMA technology and jumped directly to LTE. We account for such deployment strategies in our analysis by considering HSPA+ as part of the 4G network for ATT Wireless and T-Mobile. Thus, to determine 4G coverage we apply the best coverage available among HSPA+ and LTE. One could argue that LTE provides better bandwidth, but given the usual deployment schedule in mobile networks, where the best technology is rolled out first in high demand sites, while areas with less demand are left for later deployment, the available bandwidth per subscriber ends up being relatively similar. In figure 5, we show the geographic coverage at the state level for the year 2018. Substantial variation is evident both among providers and geographically.

form 477. DA 16-1107. Sep 30, 2016

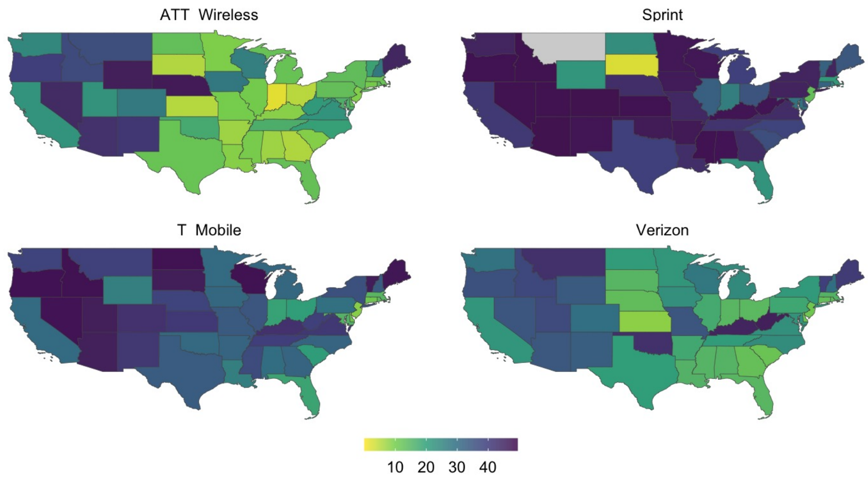

Figure 5. Mobile Provider's Coverage per State in 2015

Note: 4G Coverage shown in States where market share data exists.
Source: `https://www.fcc.gov/mobile-deployment-form-477-data`

In table 3, we show the summary statistics of geographical coverage for the year 2015. As we can see, all ISPs have 100% coverage as the maximum, which corresponds to Washington, DC. Otherwise there is variation across all providers. For all providers, there are states with less than 50% coverage. In particular, Sprint and T-Mobile do not provide coverage for large regions in some states. In our analyses, we use coverage to instrument for market share, making substantial variation of coverage crucial.

IV. Analysis

To determine the effect of throttling on user behavior, we regress app usage rates on various measures of network speed. This regression suffers from an obvious endogeneity problem as users interested in a specific app may select their network based on the access speed of that app. For instance, a user interested in watching movies on Netflix is less likely to select a network providing only slow download speeds from Netflix's servers. To address this issue, we employ the instrumental

Table 3—4G Geographical Coverage by ISP in year 2015

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
ATT Wireless	51	91.9	12.3	41.0	91.7	99.0	100.0
Verizon	51	89.2	13.1	23.3	85.0	97.2	100.0
Sprint	49	60.7	26.0	0.1	43.8	80.0	100.0
T-Mobile	50	74.3	24.5	7.6	63.7	89.3	100.0

Note: Table shows for each provider the number of observations (N), mean, standard deviation, minimum, 25th percentile, 75th percentile, and maximum of geographic coverage calculated by state (including Washington, DC).
Source: Computed from `https://www.fcc.gov/mobile-deployment-form-477-data`

variable approach using network coverage as our instrument for network usage. Coverage turns out to be highly predictive of our instrumented variable and is plausible exogenous. While it is theoretically possible that network providers alter coverage based on the apps their subscribers use, this seems unlikely to be a major factor given the significant financial investments and set-up time required to make large-scale changes to the network.

$$A_{ai} = \gamma_a + \beta_1 SlowProp_{ai} + \beta_2 X_i + \epsilon_{ai} \tag{1}$$

where A_{ai} is the usage share of app a in state i, $SlowProp_{ai}$ is the percentage of users for whom app a's traffic is throttled to slow levels, and control variables X_i. γ_a denotes a fixed effect for app a and ϵ_{ai} the i.i.d. error term. The coefficient of interest is β_1.

We use various specifications for X_i. The covariates considered are the average household income, the percentage of residents with college degree, and the percentage of residents born abroad. The last variable deserves some explanation: we hypothesize that affiliation with a foreign country may affect the degree to which residents make use of video calling apps such as Skype and possibly of video streaming, if they are unable to find content in their native languages or about their native countries in regular TV services.

Table 4 shows the results. The coefficient associated with our variable of interest, the share of customers with slow data throughput, is near zero in all our specifications, indicating an increase of app usage between 0.21 and 0.34 percentage points if the share of users with slow access increases by 10 percentage points. Overall, these coefficients are economically and statistically insignificant in all specifications.[15]

Since it is possible that effect varies by app, we repeat the regressions separately for each app in our data, including only household income as a control because of the reduced number of observations.[16] Thus, the equation or our separate regressions is:

$$A_i = \beta_0 + \beta_1 SlowProp_i + \beta_2 hh\_inc + \epsilon_i \tag{2}$$

where β_0 is the constant and hh_inc is household income in \$10,000. We also suppress the app identifier since each regression now contains data for only a single app. We report the results in table 5. Our estimated coefficients of interest are now larger in absolute value, indicating for a 10 percentage point increase of users with slow access a 1.04 percentage point decrease in usage of Netflix and increases of 1.09 or 2.81 percentage points, respectively, for YouTube and Skype. However, each of these coefficients is similar in magnitude to the estimated standard error and therefore insignificant.

[15] Our results are qualitatively comparable if we drop the most extreme values.
[16] Including all covariates leaves our results qualitatively comparable.

Table 4—Results of Pooled Regression

	(1)	(2)	(3)
SlowProp	0.021	0.028	0.034
	(0.072)	(0.072)	(0.074)
Netflix	0.220***	0.159***	0.140***
	(0.033)	(0.043)	(0.053)
Skype	0.394***	0.333***	0.315***
	(0.034)	(0.043)	(0.052)
YouTube	0.054***	-0.006	-0.020
	(0.008)	(0.021)	(0.030)
HH income		0.009***	-0.008
		(0.003)	(0.006)
College			0.388
			(0.174)
Foreign			0.175
			(0.064)
Observations	122	122	122
1st Stage F Stat	43.10	44.67	43.21

HH income: Avg. household income in $10,000s. *College*: Percentage of residents with college degree. *Foreign*: Percentage of foreign-born residents.
Standard errors in parentheses. *, ** and, *** indicate significance at the 90%, 95%, and 99% levels, respectively.

Table 5—Results of Separate Regressions

	Netflix	YouTube	Skype
SlowProp	-0.104	0.109	0.281
	(0.091)	(0.113)	(0.237)
HH income	0.005	0.011**	0.007**
	(0.006)	(0.005)	(0.003)
Contant	0.243***	0.278***	-0.016
	(0.073)	(0.071)	(0.028)
Observations	41	41	40
1st Stage F Stat	23.84	23.84	9.41

If consumers reacted strongly to the levels of throttling prevalent in the market, we would expect to see significantly negative coefficients on *SlowProp*. We fail to find evidence of such an effect. However, it is important to be aware of potential endogeneity issues. Our instrumental variable controls for endogeneity in consumer behavior. Another potential source of endogeneity is ISP behavior. ISPs may strategically throttle widely used apps to preserve bandwidth for other apps. Unfortunately, we have no way of controlling for this kind of endogeneity. However, we find it unlikely that this effect is strong. Most consumers have access to at least two mobile ISPs. Hence, throttling apps based on their popularity would provide an incentive for consumers to switch providers, hence leading to a reduction in market share.

The explanatory variable used so far is somewhat coarse as it uses a cutoff to distinguish fast from slow access speeds. It is possible that a more flexibly defined variable will be more able to capture effects of data throttling on app usage. To investigate this we define *wgt_speed* as the market-share-weighted average download speed:

$$wgt\_speed_{ai} = \sum_j MaxSpeed_{aij} s_{ij} \tag{3}$$

where wgt_speed_{ai} is the weighted average speed in for app a in state i, $MaxSpeed_{aij}$ is the observed maximally available download speed for app a's data with provider j in state i, and s_{ij} is provider j's market share in state i.

Using wgt_speed_{ai} directly in our regression would make our results liable to the same endogeneity concerns that before we were able to sidestep by the application of the 2SLS procedure. However, we cannot use a standard 2SLS approach with this independent variable because we have an instrument only for s_{ij}, not for wgt_speed_{ai}. To circumvent this problem, we run the two steps of 2SLS separately by first regressing s_{ij} on network coverage and the relevant exogenous variables; and then, based on the results of this regression, using the predicted market shares

$\hat{s}_{ij}$ to calculate predicted weighted download speeds following the definition in (3):

$$\widehat{wgt\_speed}_{ai} = \sum_j MaxSpeed_{aij}\hat{s}_{ij} \tag{4}$$

Now we can run the second stage by using $\widehat{wgt\_speed}_{ai}$ in the following regression which, except for the adjusted variable of interest, is akin to (1):

$$A_{ai} = \gamma_a + \beta_1 \widehat{wgt\_speed}_{ai} + \beta_2 X_i + \epsilon_{ai} \tag{5}$$

A complicating factor with this procedure is the significant difficulty of finding an analytical solution for the standard error. We employ clustered bootstrapping with 1,000 iterations to estimate standard errors.[17]

Tables 6 and 7 show the results for pooled and separate regressions, respectively. The coefficients of interest in the pooled regressions indicate that an increase of the weighted average download speed by 1 Mbit/s is associated with a decrease of app usage between 0.8 and 3.2 percentage points. However, they are largely insignificant and the exceptions become insignificant if we use sharpened q-values (Anderson, 2008) to account for multiple regressions.

[17]Test with different numbers of iterations produce similar estimates, indicating that our results are not sensitive to this choice.

Table 6—Results of Pooled Regression on Average Speed

	(1)	(2)	(3)
wgt_speed	-0.008	-0.032**	-0.028*
	(0.017)	(0.016)	(0.016)
Netflix	23.278***	18.323***	16.826***
	(1.055)	(2.211)	(2.342)
Skype	6.721***	4.110	2.404
	(2.358)	(2.839)	(3.199)
YouTube	40.665***	35.563***	34.095***
	(0.970)	(2.162)	(2.272)
HH income		0.888***	-0.853
		(0.264)	(0.662)
College			39.306**
			(18.107)
Foreign			19.057**
			(6.531)
Observations	122	122	122
1st Stage F Stat	54.72	28.34	14.94

HH income: Avg. household income in $10,000s. *College*: Percentage of residents with college degree. *Foreign*: Percentage of foreign-born residents.
Bootstrapped standard errors in parentheses. *, ** and, *** indicate significance at the 90%, 95%, and 99% levels, respectively. When using sharpened q-values to adjust for multiple regressions, coefficient on *wgt_speed* is insignificant for all models.

Table 7—Results of Separate Regressions on Average Speed

	Netflix	YouTube	Skype
wgt_speed	-0.020	-0.023	-0.001
	(0.054)	(0.051)	(0.005)
HH income	0.782*	0.680**	0.775***
	(0.388)	(0.323)	(0.249)
Constant	19.990***	38.893***	1.053
	(7.172)	(6.292)	(2.946)
1st Stage Observations	156	156	156
2nd Stage Observations	39	39	39
1st Stage F Stat	14.42	14.42	14.42

HH income: Avg. household income in $10,000s.
Bootstrapped standard errors in parentheses. *, ** and, *** indicate significance at the 90%, 95%, and 99% levels, respectively.

Overall, our results when using user-weighted average speeds as our explanatory variable are similar to those when using the share of users with slow data access: Our data provide no evidence that a connection between download speeds and app usage exists.

It seems likely that this is because mobile ISPs are cautious in their approach to throttling and do not slow data throughput to a degree that would severely affect user experience. In other words, market forces may be putting sufficient constraints on ISPs to limit the effect of the presence or absence of network neutrality rules. It is entirely possible that our results would be quite different, if we were to observe throttling in a monopoly setting.

V. Conclusion

A major worry of proponents of network neutrality rules, backed by some theoretical literature, is that abandoning such rules can lead to discriminatory behavior and skew competition among content providers toward the most solvent and powerful companies. However, to date there is scant empirical evidence for such effects.

We combine measured throughput rates with usage surveys to analyze how users react to discriminatory throttling by mobile ISPs. In multiple specifications we find no effect of throttling on app usage rates.

We employ an instrumental variable approach to control for the obvious endogeneity problem that consumers can switch to a provider offering fast access to data they care about. Another source of endogeneity is that ISPs could reduce data throughput for the most popular apps in their networks. We have no direct way of controlling for this behavior. However, ultimately we do not believe this effect to be too important. With consumers in most local markets being able to choose among multiple ISPs, any provider throttling popular content too drastically would risk losing market share. In other words, we interpret the lack of significant effects as attributable to ISPs showing restraint. ISPs could, but choose not to affect relative data transmission rates too much.

While, to our knowledge, our study represents the first effort of testing the effect of net neutrality rules on consumers and content providers empirically, it suffers from having limited data which varies only at the state level. As such, it is only a starting point and future studies should try to find finer data to get a more detailed picture of consumer behavior in the light of throttling.

REFERENCES

Anderson, Michael L. 2008. "Multiple Inference and Gender Differences in the Effects of Early Intervention: A Reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects." *Journal of the American Statistical Association*, 103(484): 1481–1495.

Boussion, François, Patrick Maillé, and Bruno Tuffin. 2012. "Net neutrality debate: Impact of competition among ISPs." 1–8, IEEE.

Briglauer, Wolfgang, Carlo Cambini, Klaus Gugler, and Volker Stocker. 2021. "Net Neutrality and High Speed Broadband Networks: Evidence from OECD Countries." *23rd Biennial Conference of the International Telecommunications Society (ITS): "Digital societies and industrial transformations: Policies, markets, and technologies in a post-Covid world", Online Conference / Gothenburg, Sweden, 21st-23rd June, 2021.* Calgary:International Telecommunications Society (ITS).

Choi, Jay Pil, and Byung-Cheol Kim. 2010. "Net Neutrality and Investment Incentives." *RAND Journal of Economics*, 41(3): 446–471.

Choi, Jay Pil, Doh-Shin Jeon, and Byung-Cheol Kim. 2015. "Net Neutrality, Business Models, and Internet Interconnection." *American Economic Journal: Microeconomics*, 7(3): 104–41.

Economides, Nicholas, and Benjamin E. Hermalin. 2012. "The Economics of Network Neutrality." *RAND Journal of Economics*, 43(4): 602–229.

Economides, Nicholas, and Joacim Tåg. 2012. "Network neutrality on the Internet: A two-sided market analysis." *Information Economics and Policy*, 24(2): 91–104.

Ford, George S. 2018. "Net Neutrality and Investment in the US: A Review of Evidence from the 2018 Restoring Internet Freedom Order." *Review of Network Economics*, 17(3): 175–205.

Gans, Joshua S. 2015. "Weak versus Strong Net Neutrality." *Journal of Regulatory Economics*, 47: 183–200.

Greenstein, Shane, Martin Peitz, and Tommaso Valletti. 2016. "Net Neutrality: A Fast Lane to Understanding the Trade-Offs." *Journal of Economic Perspectives*, 30(2): 127–50.

Hazlett, Thomas W., and Joshua D. Wright. 2017. "The Effect of Regulation on Broadband Markets: Evaluating the Empirical Evidence in the FCC's 2015 "Open Internet"Order." *Review of Industrial Organization*, 50(4): 487–507.

Jullien, Bruno, and Wilfried Sand-Zantman. 2014. "Pricing Internet Traffic: Exclusion, Signalling and Screening." *CESifo Working Paper Series No. 4709.*

Krämer, Jan, Lukas Wiewiorra, and Christof Weinhardt. 2013. "Net neutrality: A progress report." *Telecommunications Policy*, 37(9): 794–813.

Lee, Daeho, and Yong-Hwa Kim. 2014. "Empirical evidence of network neutrality – The incentives for discrimination." *Information Economics and Policy*, 29: 1–9.

Li, Fangfan, Arian Akhavan Niaki, David Choffnes, Phillipa Gill, and Alan Mislove. 2019. "A Large-Scale Analysis of Deployed Traffic Differentiation Practices." *SIGCOMM '19*, 130–144. New York, NY, USA:Association for Computing Machinery.

Ma, Richard TB, Jingjing Wang, and Dah Ming Chiu. 2017. "Paid prioritization and its impact on net neutrality." *IEEE Journal on Selected Areas in Communications*, 35(2): 367–379.

Peitz, Martin, and Florian Schuett. 2016. "Net Neutrality Inflation of Traffic." *International Journal of Industrial Organization*, 46: 16–62.

Schuett, Florian. 2010. "Network Neutrality: A Survey of the Economic Literature." *Review of Network Economics*, 9(2).

Wu, Tim. 2003. "Network neutrality, broadband discrimination." *J. on Telecomm. & High Tech. L.*, 2: 141.