THE INVISIBLE HAND: ALGORITHMIC CONTROL OF YOUTUBE CONSUMERS AND PROVIDERS

Mareike Möhlmann, Assistant Professor, Bentley University
Ola Henfridsson, Professor, Miami Herbert Business School
Working Paper: April 2021

Platforms are typically considered to have limited control over their participants, yet extensive use of machine learning is increasingly telling another story. Our qualitative study of YouTube participants shows how learning from data captured through behavioral tracking enables targeted personalization and behavioral nudging of consumers, and how data collected through behavioral performance monitoring builds a basis for automated censorship and subtle incentivization of content providers. Building on these observations, we present a grounded theory that captures algorithmic control of consumers and providers. We contribute to the platform literature by discussing how algorithmic control is subject to constant change and generativity, and how platforms collect not only input and output data but also process data, in order to develop controlling algorithms. We also demonstrate the ways by which algorithmic control blurs the traditional distinction between formal and informal control, where algorithmic control bears resemblance of formal control in terms of its coerciveness but still exhibits the fluidity of informal control. In this regard, algorithmic control can be interpreted as an 'invisible hand'.

INTRODUCTION

Research on platforms widely concludes that they provide an infrastructure allowing platform participants to undertake value-creating interactions (see Constantinides et al. 2018; Parker et al. 2016). Platform participants—consumers and providers—are actors who operate independently of the platform itself (Gawer 2014; Parker et al. 2016). For transaction platforms and innovation platforms alike (see Cusumano et al. 2019; Gawer 2020), this supposedly makes it difficult to exert direct control over platform actors. For instance, in research on transaction platforms, participants are viewed as external actors who use the platform as a marketplace for transactions (Hagiu and Spulber 2013; Hong et al. 2016; Parker and Van Alstyne 2005; Van Alstyne et al. 2016). Similarly, participants on innovation platforms are typically viewed as external actors who use the platform technology and its tools to extend the platform and create their own end-user applications (Ghazawneh and Henfridsson 2013; Tiwana 2015). In this regard, the platform owner's control over platform participants is traditionally seen as indirect. For instance, platforms such as YouTube can improve their matchmaking between providers and consumers by drawing on input data, such as viewers subscribing to a specific YouTube channel, and output data, such as the number of likes or comments received (e.g., Parker et al. 2016, Casadesus-Masanell and Halaburda 2014), but they are not typically regarded as coercing creators to create and consumers to consume particular content.

However, anecdotal evidence suggests that, in their quest to deliver platform user value (Gregory et al. 2020) and to address regulatory pressures (Aral 2020; Zuboff 2019), digital platforms are increasingly implementing measures that increase their control of both consumers and providers. In the example of YouTube, process data on consumers' browsing behavior and content providers' performance are collected on a large scale (Gregory et al. 2020; Schildt 2017). Machine learning algorithms are then employed by the platform to identify patterns in the data. For instance, learning from data captured through behavioral tracking enables targeted personalization and behavioral nudging of consumers (see Zuboff 2019). It also builds a basis for automated censorship and subtle incentivization of providers. In addition, platforms may employ computational tools to enable automatic analysis of millions of pieces of content uploaded to platforms (Burrell 2016; Faraj et al. 2018; Gregory et al. 2020; Schildt 2017), allowing them to filter and control by automatically deleting content.

Powered by the latest advances in machine learning technology, algorithmic control is qualitatively different from the types of control examined in the information systems literature (e.g., Kirsch 1997; Wiener et al. 2016). In particular, algorithmic control seems to blur traditional distinctions between formal and informal control, as it resembles formal control in terms of its coerciveness, yet exhibits the fluidity of informal control. These informal control elements are characteristic of the largely unpredictable, unseen and unobservable force of algorithmic control. We empirically investigate algorithmic control and its sub-dimensions through an in-depth study of YouTube, one of the largest and most influential digital platforms worldwide. We use grounded theory techniques (Charmaz 2014; Gioia et al. 2013; Glaser 1978) to allow new insights to emerge from the data. We detail how YouTube's algorithmic control operates on two dimensions: control of providers and control of consumers.

Our contributions to the literature on digital platforms are threefold. First, we develop a grounded view of platforms' control over their platform participants that goes beyond the extant literature. This grounded view enscapulates input, process and output data as integrated elements. Second, we extend existing work on control in information systems by proposing the notion of algorithmic control, which exhibits elements of both informal and formal control (see Wiener et al. 2016). Third, our findings challenge the assumption that control mechanisms tend to be inflexible and mechanistic (Kirsch 1997); rather, our findings suggest that algorithmic control is subject to constant change and generativity.

THEORETICAL BACKGROUND

Control Theory

Control can be seen as attempting to align individual behavior with organizational goals (Choudhury and Sabherwal 2003; Kirsch 1997; Wiener et al. 2019). It tends to assume a dyadic relationship between a controller as the source of control, and a controllee as the target of control. The former is viewed as implementing measures to regulate the latter (Wiener et al. 2016). In the context of platforms, the platform owner is the controller, while the controllees are platform participants, whether providers or consumers.

Drawing on Ouchi's (1977, 1978, 1979) seminal work, control typologies in IS research distinguish between formal and informal control mechanisms (Jaworski 1988; Kirsch 1997; Choudhury and Sabherwal 2003; Wiener et al. 2016). *Formal control* refers to explicit controller prescriptions and can be classified into input, behavior and outcome controls. Examples include specification, evaluation, reward and sanctioning mechanisms based on predefined criteria (Tiwana et al. 2010; Wiener et al. 2016). In contrast, *informal control* refers to "soft" mechanisms implemented to implicitly (rather than explicitly) influence the determinants of controllees' behaviors, for example by establishing and communicating shared values, norms and beliefs, or self-monitoring based on intrinsic motivation (Kirsch 1997; Tiwana et al. 2010; Wiener et al. 2016). Since formal rules and regulations, as well as informal social group settings, tend to be shaped over years of negotiation and interaction, previous research considers control mechanisms to be somewhat inflexible and mechanistic (Kirsch 1997).

Control on Digital Platforms

Although control is a theme within digital platform research (Ghazawneh and Henfridsson 2013; Tiwana et al. 2010), our review of the literature on transaction and innovation platforms (Cusumano et al. 2019; Gawer 2020), which are two main strands of this literature, reveals that platform participants are regarded largely as independent agents pursuing their own goals and desires.

Transaction Platforms

One stream of platform research is grounded in economics, viewing platforms essentially as marketplaces for transactions between platform participants (see e.g., Burtch et al. 2018; Gawer 2020; Hagiu and Spulber 2013; Hong et al. 2016; Parker and Van Alstyne 2005; Van Alstyne et al. 2016). Platform participants are treated as external actors using the platform for their own purposes. Platforms may draw on input and output data (e.g., Parker et al. 2016; Casadesus-Masanell and Halaburda 2014), including information provided by participants, and may use performance metrics such as "likes" or ratings to ensure efficient matching of demand and supply.

Platforms may also exercise control by providing escrow services and insurance to enable smooth interactions between platform participants (Hu et al. 2004; Pavlou and Gefen 2004; Rice 2012), or may intervene in the market, for example by implementing differentiated pricing strategies for consumers and providers (Parker et al. 2016; Parker and Van Alstyne 2005). However, the platform is a multi-sided marketplace in which platform participants remain distinct and independent entities (Hagiu and Spulber 2013; Pavlou and Gefen 2004; Parker and Van Alstyne 2005). Thus, the platform owner's main function is to coordinate consumers and providers (Parker et al. 2016; Tiwana et al. 2010; Wiener et al. 2019).

Innovation Platforms

Another research stream, grounded in engineering design (Gawer 2014), thinks of platforms as arenas for innovation (Cusumano et al. 2019), enabled by layered, modular architectures (Constantinides et al. 2018; Yoo et al. 2010). This strand of research views platforms as ecosystems, providing a foundation for other stakeholders to develop complementary products, applications and services (De Reuver et al. 2018; Gawer 2014). Platforms are conceptualized either as extensible codebases (see e.g., Tiwana et al. 2010), or as socio-technological artefacts comprising organizational and technical elements (De Reuver et al. 2018). While the platform's ecosystem provides a stable core (Gawer 2014; Tiwana et al. 2010), platform participants are independent entities (De Reuver et al. 2018), who are difficult to manipulate or influence. This research stream suggests that platform owners' control relates merely to managing participants' access to the platform environment (Ghazawneh and Henfridsson 2013; Karhu et al. 2018). In other words, control refers to the platform's strategic decisions on whether or not, and under what conditions, to grant external complementors access to the platform environment, for example through APIs or open-source codebases (Ghazawneh and Henfridsson 2013). Indeed, only through openness and by allowing outside contributions can platforms experience innovation and growth (Karhu et al. 2018); thus, overly tight control would hinder valuable third-party developments (Ghazawneh and Henfridsson 2013).

Identification of Research Gaps

Both research streams suggest that control on digital platforms is a salient issue, but such research is in its infancy. Indeed, Wiener et al. (2019) comment that control in the digital era is changing, and requires new approaches and methods. Signs of change can be seen in research on algorithm-mediated practices (Curchod et al. 2019; Kellogg et al. 2019; Zuboff 2019), which suggests that advanced digital technologies are accelerating large-scale monitoring and tracking of individuals, as platforms collect vast amounts of data on platform participants (e.g., Schildt 2017; Zuboff 2019). This line of research investigates how platforms use insights from data to perform tasks such as automated flagging of content (Faraj et al. 2018; Lustig et al. 2016), and to intelligently monitor well-defined tasks such as managing platform workers (Rosenblat and Stark 2016; Kellogg et al. 2019).

In view of emerging changes in the very nature of control (Wiener et al. 2019), empirical research is needed on whether and how platforms use machine learning algorithms to *control* platform participants. The notion of algorithmic control and its sub-dimensions warrant detailed attention and concise theorization. To this end, we embarked on a study of algorithmic control on the video streaming platform, YouTube.

RESEARCH DESIGN

Data Collection and Analysis

We collected data from a variety of sources within YouTube's platform ecosystem, including interviews, user comments and press releases. First, we conducted 64 semi-structured interviews with a variety of relevant stakeholders, including content creators (n=24) and viewers (n=26), as well as executives and corporate platform complementors (n=14), who included social marketers and corporate YouTube creator coaches. We also identified relevant forum posts by filtering YouTube content using the keywords "YouTube," "data" and "nudging." We harvested all user comments posted on videos associated with these key terms, providing us with 2,917 posts. Finally, we selected relevant material from official press releases published by the YouTube executive team using the key terms "algorithm" and "control" (n=35).

During our data collection we focused on a range of topics. When interviewing providers, we were particularly interested in the process of content creation and how providers interact with the platform. We focused on content creators' perceptions of the rules and regulations with which they must comply and how YouTube curates and controls content. From the consumers' perspective, we asked questions relating to individuals' video preferences, browsing and viewer behavior. We also focused on how viewers experience YouTube features, such as the recommender system which proposes content that may be of interest to them, and how they feel they are influenced by the content to which they are exposed. Finally, we interviewed YouTube executives and corporate platform complementors to gain insights into how the YouTube algorithm works and how YouTube potentially controls platform participants. We focused on similar topics when collecting data from user comments and press releases, filtering for relevant data using keywords such as "algorithm," "control" and "nudging" (see Table 1).

We adopted a grounded theory approach to our data analysis (Glaser 1978), and began this project with a broad research interest in the YouTube algorithm. Through an iterative process in which we constantly revised the coding scheme for our multi-method data and compared emerging findings with theory, we ultimately narrowed our research scope (Charmaz 2014; Glaser 1978) to focus on algorithmic control of consumers and providers.

The data analysis was conducted in two principal stages (Glaser 1978; Charmaz 2014). First, following the fundamental principles of grounded theory, we went through the data and labelled text pieces from all three data sources with open codes. Whenever a piece of text triggered a thought, we labelled it with keywords or a short sentence. These initial codes were our first-order indicators. In line with Gioia et al.'s (2013) approach to grounded theory, we ensured that the codes reported remained as close to the original text snippet as possible. In the second step, we abstracted and accumulated these codes into second-order themes, followed by third-order categories. In an iterative approach, we constantly revised our categories by reflecting on theory and literature. For example, only after several rounds of coding and consultation of relevant literature did we realize that consumers and providers on

YouTube are subject to different forms of algorithmic control. For example, *automated censorship* is a second-order theme that clearly describes an element of control over content creators but is less suited to capturing control over viewers, as only those producing content can be censored. Finally, as our multi-method data set comprised three different data sources, we paid attention to comparing data from users' comments, press releases and semi-structured interviews, and ensured that they were labelled coherently. Table 1 summarizes our first-order codes, second-order themes and third-order constructs. As we were interested in theorizing the core construct of algorithmic control rather than its nomological network, as in research by Curchod et al. (2019) and Sarker et al. (2002) we only implemented these two coding phases, waiving the axial coding phase.

Table 1. Data Collection

<table>
<tr><th>Sources</th><th>Topics</th><th>Purpose</th></tr>
<tr><td>Semi-structured interviews with YouTube content creators (n=24), YouTube viewers (n=26) and YouTube executives and corporate platform complementors such as social marketers and YouTube creator coaches (n=14).</td><td rowspan="3">• Content creators' experience of YouTube, and particularly the rules and regulations with which they must comply
• Viewers' experiences of YouTube features (e.g. the recommender system) and how they feel influenced by content to which they are exposed
• Insights from YouTube executives and corporate platform complementors into the algorithm and how YouTube controls platform participants</td><td rowspan="3">• Identification of the key focus of the study: algorithmic control of consumers and providers
• Understanding and building grounded categories of the key constructs of algorithmic control, and how it may differ across consumers and providers</td></tr>
<tr><td>User comments posted below videos identified by filtering YouTube content using the keywords "YouTube," "data" and "nudging" (n=2,917 posts).</td></tr>
<tr><td>Official press releases by the YouTube executive team (n=35) using the key terms "algorithm" and "control."</td></tr>
</table>

CASE AND FINDINGS

Founded in 2005, YouTube is a leading global video streaming platform, ranked by Alexa Internet (2020) as the second most popular website worldwide. The YouTube platform facilitates uploading, viewing, commenting and sharing of video content, ranging from documentaries and educational videos to video blogging. About 500+ hours of video content are uploaded to YouTube every minute, which is consumed by as many as five billion viewers per day[1]. In a press statement, Susan Wojcicki (2017), CEO of YouTube, described the platform as follows:

> *As the CEO of YouTube, I've seen how our open platform has been a force for creativity, learning and access to information. I've seen how activists have used it to advocate for social change, mobilize protests, and document war crimes. I've seen how it serves as both an entertainment destination and a video library for the world. I've seen how it has expanded economic opportunity, allowing small businesses to market and sell their goods across borders. And I've seen how it has helped enlighten my children, giving them a bigger, broader understanding of our world and the billions who inhabit it[2].*

The YouTube platform facilitates the interaction between two major participant groups: viewers who consume content, and content creators who produce video content. Our grounded analysis reveals that algorithmic control operates differently on these two participant groups, resulting in the theorization of two distinct constructs: algorithmic control of consumers, and algorithmic control of providers.

Algorithmic Control of Consumers

YouTube leverages machine learning techniques to control consumers. Algorithmic control of consumers refers to the implementation of *behavioral tracking*, which allows platforms to collect vast amounts of data on browsing activities. They then employ machine learning algorithms to identify patterns in the data, allowing them to expose viewers to targeted and *personalized information*, and to employ *behavioral nudging techniques*.

Behavioral Tracking

YouTube attracts a variety of viewers, ranging from occasional visitors to extreme viewers who spend hours a day on the platform. The latter tend to create YouTube profiles, facilitating the platform's access to information about basic user demographics and the frequency of website visits. One YouTube executive reported:

> *…one sort of signal that the algorithm is using … is all about the user itself. Has the user been often on YouTube? Has the user been interested in that particular content*

[1] https://blog.youtube/press
[2] https://blog.youtube/news-and-events/expanding-our-work-against-abuse-of-our

> *or not, like how often this user goes on YouTube, if that user is interested in watching a long video versus a short video* (YouTube executive).

In addition, YouTube closely monitors viewers' every move by collecting vast amounts of data on their real-time website browsing behavior. This includes tracking their topic and video format preferences. YouTube also collects data on users' more active engagement, such as whether they subscribe to video channels, or like and dislike particular video content. As one YouTube executive explained:

> *That, second sort of big bucket—and it's a really important bucket—is everything around like the social aspects of how often that video is actually being watched, is there any likes, how many comments, like, you know, some of the feedback* (YouTube executive).

Data collected on viewers are not restricted to information gathered on YouTube itself. As the platform is a subsidiary of Google, individuals' data gathered on a variety of products, such as the Google search engine, Gmail and Chrome, can obviously be merged to gain even deeper insights into behaviors and preferences. In addition, once viewers have downloaded Google apps onto their smartphones, Google can constantly track their location, allowing the platform to add information about external variables, such as relevant weather data. As announced in the press release below, YouTube even strives to access information provided by external platforms such as Facebook and Yahoo!, for example to gain insights into individuals' personal networks, by allowing viewers to link their YouTube accounts with these platforms:

> *You might have noticed that YouTube's been getting a lot more social. We've launched several features in the last few months that let you better connect with the platforms that matter most to you (and discover new videos you're likely to love in the process). In addition to linking your YouTube account to social networks like Twitter and Facebook via AutoShare, friend suggestions and easier private sharing options, you can easily find the YouTube accounts belonging to your friends on Facebook, Yahoo! and Gmail* (YouTube press release).

Interestingly, we find that while many viewers reported that they are aware that YouTube collects data about them, to which they initially consented by agreeing to the terms and conditions of use, they tend to state that they have little detailed knowledge of the uses to which their personal data are put:

> *They use our data—but the security—you don't know where the data goes and what they do with it and how they use it. So it's scary when you understand this. But they don't make it obvious, that's the problem. … Regulations are very important too, you know. I'm not sure that everybody reads all their privileges. Everyone just clicks the terms and conditions box* (YouTube viewer).

Targeted Personalization

YouTube feeds vast amounts of the data gathered into its machine learning algorithms to identify patterns in the data. These allow it to expose viewers to personalized information in the form of video content, as well as advertising content. Indeed, YouTube's business model revolves around generating income by exposing users to targeted advertising, so YouTube viewers encounter ads at the beginning, middle or end of a video. As one content creator explained:

> *And if you compare that to YouTube, for example, which its whole business model is around ads that appear at the beginning, middle and end of a creator's videos. So the only way that YouTube is going to make more money is if creators' videos are seen more. So the more people spend time on the platform and the more people watch multiple videos and just spend more time on the platform essentially, the more money YouTube is going to make, which is very different from how Instagram works* (YouTube content creator).

From viewers' perspective, being exposed to personalized content is a double-edged sword. Although they receive information that is relevant and interesting to them, and thus discover videos through YouTube's recommendation system, the downsides are filter bubbles that result in users being logged into so-called "echo chambers." Personal views and beliefs are reinforced through repeated exposure to non-diverse information.

> *I think there's sort of a deep question here, which is about, like, what do we get when we get what we want? Because that happens on Twitter and on Facebook and on YouTube. I think from the outside, it could seem very utopian: you only see the things that you want to see. On the other hand, it can lead to things like filter bubbles or echo chambers. In the YouTube space, I think it's at risk of becoming sensationalist* (YouTube content creator).

The underlying mechanism is a non-linear cycle in which users develop preferences based on the content to which they are exposed, and end up watching even more of such content:

> *So my point being, like yeah, there's a horrible nonlinear cycle in the idea of viewer preference, where the viewers can't have a preference in a vacuum. So then the algorithm needs to show them something, and from those things that it shows them, then the viewers/users start to know what they can watch. From that, they develop their preferences, and their preferences strengthen. So essentially, anything that the algorithm shows is then like a positive feedback loop to get more of that* (YouTube content creator).

Behavioral Nudging

Exposure to targeted and personalized content opens up opportunities to employ behavioral nudging, which is the influencing and modifying of individuals' behavior through exposure to suggestions and positive reinforcement. As one YouTube executive put it:

> *The nudging … it's happening with YouTube. There's a lot of tech talk about it. It's super interesting phenomena. … For example, people are only looking at like right-wing content from the news based on their profile, just because of the algorithm it* [YouTube] *uses. You know, this is what people want, and this is what the algorithm is showing* (YouTube executive).

Another YouTube executive specified that the company gains insights from sophisticated data analytics to identify customers who are on the verge of converting, and thus intend to buy a product or service. These customers are considered to be particularly valuable to the companies that run advertising campaigns on YouTube, as they just need a last "nudge" in order to buy. One YouTube creator shared the following example:

> *And so we know that if it's been one month that he is looking at YouTube videos of types of cars, etc., if he is looking at different websites, uh, of BMW, Mercedes or Audi etc., that maybe he's on the verge of converting. Okay, he's on the verge of actually buying a car. And so that's why we rely on that much data and rely on so many much more data points that just used to recall, uh, searching and videos that you watch. We also rely on what we call the intent, so the intention that you have that can be observed in your behavior online. So that's how we do it* (YouTube executive).

Two different forms of nudging can thus be distinguished on YouTube. First, viewers may be manipulated by being exposed to targeted display advertising embedded in YouTube videos, resembling more traditional forms of marketing. One viewer commented in the YouTube forum:

> *Taking a look at marketing today it has really gotten out of hand. It's straight up manipulation across the board now. The worst part is that people aren't aware and most don't understand what's being done to them. In my opinion banning all kinds of psychological manipulation to make someone buy things they don't want or need is the way to go* (YouTube viewer).

Second, viewers may be swayed by the opinions or recommendations of YouTube influencers, who represent a large proportion of content creators. Often, influencers may be trusted and admired, and what topics they address, and how, may nudge viewers into behaving differently from how they would have without exposure to that content. Influencers are often perceived to be more trustworthy than companies, and thus their role model function may be more profound:

> *There's a lot of influencers … who are sharing information in a less transparent way. Just like not making it completely transparent that they are being sponsored or got something for free, and they're actually doing whatever advertisement, whatever testing of the project, talking about it. … So I guess, you know, there's a difference between how clear they are and all of that, any kind of sponsorship I guess that they're getting. Uh, but yes, I think there's a difference, in the sense that I think there's research on that people are more likely to believe another person than a company. Um, so there's a difference in that way. 'Cause if it's marketing thoughts from that particular company or organization, then it can be viewed quite differently* (YouTube viewer).

Algorithmic Control of Providers

Algorithmic control of providers refers to *behavioral performance monitoring*, which allows vast amounts of data to be collected on the quality of a provider's work. Platforms employ machine learning algorithms to identify patterns in the data, enabling them to employ *automated censorship* and *subtle incentivization*.

Behavioral Performance Monitoring

For providers, YouTube offers a gateway to an audience. However, not all videos become successful. Similarly to search engine optimization (SEO), which marketers adopt in order to improve their Google search rankings, many YouTube creators seek to optimize their video content in order to appear more frequently in YouTube video recommendations. YouTube tracks the characteristics and performance of video content in real time to determine whether it is of interest to specific viewer demographics, and thus keep viewers active on the platform. One major tool is key content analytic metrics, which are shared with creators via "YouTube Studio." This displays each video's click-through rate, watch time and return rate. One content creator explained the key metrics as follows:

> *I think it's click-through rate, watch time, session watch time and then return rate. The delta of the return rate. So I think they're* [YouTube] *just trying essentially to keep people clicking on videos. They're trying to get people to watch the videos as long as possible. They're trying to get people to not just watch the video as long as possible, but to stay on YouTube as long as possible* (YouTube content creator).

Although YouTube creators' performance is constantly monitored, they are aware that only some key performance criteria are shared with them, helping them to create high-quality content. As the algorithm draws on "*hundreds if not thousands of data points*," they have little knowledge of exactly what metrics are used to evaluate their content, leaving them with a feeling that there is little transparency in precisely how the algorithm works:

> *YouTube does not give me probably what I guess to be hundreds if not thousands of data points that they use internally to determine how much distribution they are going to give you. So they are not giving me data to help me get more distribution; they are giving me data to help me improve the quality of my content. That's the distinction between what they show and what they don't show* (YouTube content creator).

As creators constantly try to figure out the underlying logic of the YouTube recommender algorithm, in addition to accessing information provided by YouTube, many seek to access information from external parties. One creator mentioned: "*I've heard about creator schools, where some YouTubers offer personal training.*" Indeed, many YouTubers pay for support from YouTube coaches, hoping to produce videos that the platform will classify as high-quality content.

Automated Censorship

YouTube employs machine learning technology, paired with human moderation, to identify content that does not comply with its community guidelines or with government regulations. YouTube CEO, Susan Wojcicki recently explained in a press release that millions of videos have been reviewed for violent, pornographic or extremist content, in order to train machine learning technology to automatically flag any such video content uploaded to the platform. The platform exercises control by reserving the right to take down such videos, and has even banned creators' accounts from the platform.

> *In the last year, we took actions to protect our community against violent or extremist content, testing new systems to combat emerging and evolving threats. We tightened our policies on what content can appear on our platform, or earn revenue for creators. We increased our enforcement teams. And we invested in powerful new machine learning technology to scale the efforts of our human moderators to take down videos and comments that violate our policies* (Susan Wojcicki, YouTube CEO).

While YouTube is transparent about filtering videos that explicitly violate community guidelines, some of its censorship is perceived to be somewhat arbitrary. YouTube publicly admits to reducing the distribution of content that comes close to, but does not quite cross the line into, violating its community guidelines, presenting this as a valuable means to fight the spread of fake news:

> *To that end, we'll begin reducing recommendations of borderline content and content that could misinform users in harmful ways—such as videos promoting a phony miracle cure for a serious illness, claiming the earth is flat, or making blatantly false claims about historic events like 9/11* (YouTube press release).

Content creators are generally critical of this rather opaque practice of censorship, and some blame the platform for attempts at "*getting rid of alternative political philosophies*" (YouTube content creator). One content creator reported:

> *The scary thing is that they can literally just get rid of your channel if they don't like what you're talking about. None of these guys* [content creators banned in the past] *were—they weren't trying to incite violence. I mean, there are people that are unpopular on YouTube that actually tried to incite violence … But the point being, they just didn't... If they don't like what you're talking about, they can get rid of your*

> *entire channel. But again, that's obvious censorship, right?* (YouTube content creator).

Subtle Incentivization

In addition to banning certain content through censorship, YouTube employs more subtle forms of control, by either monetizing or demonetizing video content. Monetization refers to creators' ability to earn advertising income from video content, while demonetization refers to content being classified as unsuitable for advertising. Some creators who depend on income generated from ads may be badly affected by demonetization, incentivising them to create videos that are unlikely to fall victim to this practice. One content creator reported on an incident he had observed:

> *What ended up happening was that she was creating weird content that YouTube deemed unsuitable for advertising, or they call it not advertiser-friendly, and so she got all of her ads or all of her videos marked not advertiser-friendly. And so she lost all the income that she was making, she was able to continue produce. … So what typically happens is YouTube comes out with a change, and almost overnight it affects the creators because there aren't a lot of warnings* (YouTube content creator).

As demonetization is considered to be a relatively informal mechanism that YouTube can implement very suddenly and without previous warning or explanation, information about new policies tends to be informally dispersed among the YouTube community.

> *So we have to kind of hear it. The information gets dispersed among the community, but it takes a few days, and by the time that we realize it, we're already seeing a decrease in views, or subscribers have gone backwards, or our income has gone down. And I think that being on YouTube five years, I've seen a lot of ups and downs and wild swings in the algorithm* (YouTube content creator).

As YouTube does not provide detailed information about the process and the factors taken into consideration, other than referring to the fact that the videos address controversial topics, content creators make various assumptions about the reasons for demonetization. One creator explained that he had experienced demonetization of content he had posted, finding that topics such as mental health, depression and suicide were eventually demonetized:

> *I made a video this year about CBD oil, and that's just a topic that advertisers aren't happy to be in front of … It was demonetized because they say, "Oh. It's a*

controversial topic" … It [the video] *doesn't get promoted. Or if I have made videos about mental health, talking about suicide, talking about depression, you can't monetize it, but therefore the algorithm does not care to promote it whatsoever. … It's a computer-generated thing that listens to the words spoken. But then, even so much so, as if the comment section is very inappropriate, that will demonetize a video, which I totally am for. I support that completely* (YouTube content creator).

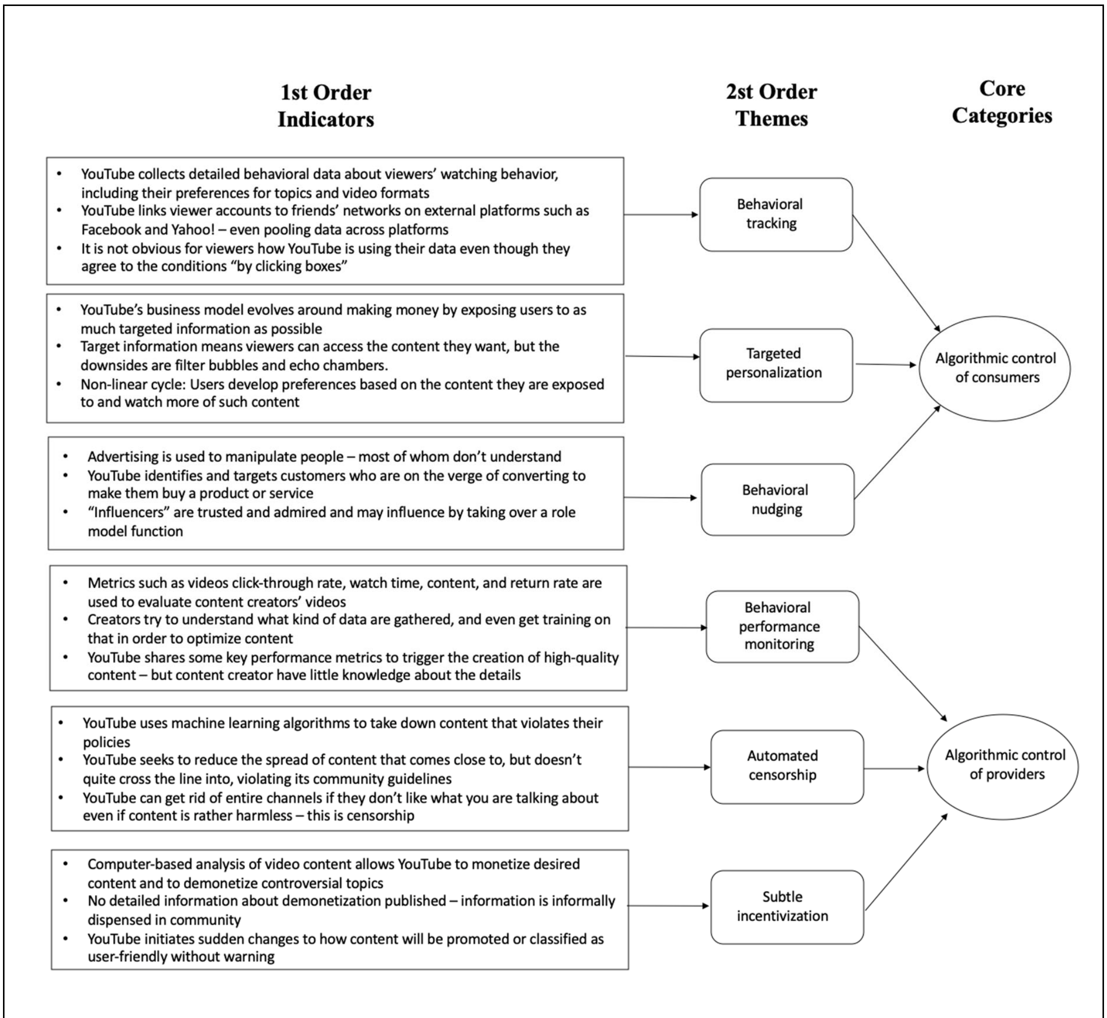

Figure 1. Coding scheme for core categories of algorithmic control of consumers and algorithmic control of providers

INTEGRATION OF FINDINGS, AND IMPLICATIONS FOR RESEARCH AND PRACTICE

Two distinct constructs emerge from our grounded conceptualization. *Algorithmic control of consumers* refers to *behavioral tracking*, which allows platforms to collect vast amounts of data on browsing activities. They then employ machine learning algorithms to identify patterns in the data, enabling them to expose viewers to targeted and *personalized information*, and to employ *behavioral nudging* techniques. *Algorithmic control of providers* refers to how data which is collected through *behavioral performance monitoring* builds a basis for *automated censorship* and *subtle incentivization* of content providers.

Our findings on *behavioral tracking* and *behavioral performance monitoring* resonate with research on how digital platforms monitor and track their consumers' and providers' behavior (Aral 2020; Curchod et al. 2019; Schildt 2017), and how new advances in digital and mobile technologies allow companies to capture data on individuals' clicking and browsing behavior (Kellogg et al. 2019; Zuboff 2019). However, in examining control in the digital era (Wiener et al. 2019), we extend this previous research by zooming into how platforms leverage these data and advances in AI by employing machine learning technology to *control* consumers and providers.

Although it is widely accepted that participants on transaction and innovation platforms operate independently of the platform firm itself (Gawer 2014; Parker et al. 2016), our findings suggests otherwise. Through the collection of vast amounts of data enabled by *behavioral tracking* and *behavioral performance monitoring,* consumers' and producers' behavior is "absorbed" in real time, and these data are used to train the platform's algorithm. A large proportion, for example in the form of information on their behavioral preferences, will remain permanently inscribed in the platform's algorithm, even if they decide to leave the platform environment. Thus, our findings question the assumption that participants are independent entities, as the lines between the platform firm and participants are becoming increasingly blurred (see Gawer 2020).

Although we theorize *algorithmic control of consumers* and *algorithmic control of providers* as two distinct constructs, we identify some common sub-themes. Both theorized constructs capture the platform's collection of vast amounts of data (*behavioral tracking* or *behavioral performance monitoring*), which provides necessary inputs to leverage the efficiency of sophisticated machine learning technology. Furthermore, the two constructs capture how the algorithms identify patterns in the data, which are used to carefully curate and filter content to manipulate and trigger desired behaviors among platform participants (*targeted personalization* and *behavioral nudging*, or *automated censorship* and *subtle incentivization*).

However, despite these common sub-themes, we also extend previous research by theorizing substantial differences between these two forms of algorithmic control. Previous work on innovation platforms (Cusumano et al. 2019; Gawer 2014) has focused on control of platform providers, in relation to external complementors' access to the platform environment (Ghazawneh and Henfridsson 2013; Karhu et al. 2018), but remains relatively silent on the potential control of consumers. While control on transaction platforms tends to be regarded as a market intervention enabling the platform intermediary to balance the needs of consumers and providers (Hagiu and Spulber 2013; Parker et al. 2016; Parker and Van Alstyne 2005), related research has not focused sufficiently on how the intermediary may apply very different control mechanisms to the two sides. Unlike previous studies, we offer evidence that the two sides of the platform are subject to different forms of control. For example, our grounded theory captures that control of consumers, in the form of *behavioral nudging*, differs from control of providers, in the form of *subtle incentivization*. Unlike YouTube viewers, content creators are driven by a more deliberate desire to understand YouTube's control mechanism in order to maximize their content views or revenue generation. YouTube controls providers through *subtle incentivization*, influencing their activities on the platform such as the kinds of content they create, whereas control through *behavioral nudging* is likely to effect consumers' buying or voting behavior outside the platform ecosystem.

Contributions to Research on Control on Digital Platforms

Previous studies widely accept that platform owners tend to coordinate (rather than control) interactions between platform participants (Parker et al. 2016; Tiwana et al. 2010; Wiener et al. 2019; Yoo et al. 2010). In contrast to this dominant view, we present a grounded theory of two different forms of algorithmic control. Our findings make several key contributions.

First, while previous research has tended to stress that platforms access *input* and *output* data in order to optimize matching efficiency (e.g., Parker et al. 2016; Casadesus-Masanell and Halaburda 2014), our findings indicate that platforms like YouTube control platform participants by also collecting detailed *process* data, through behavioral tracking versus performance monitoring. We extend the literature by offering a grounded view of platforms' control over their participants, in which input, process and output data are integrated elements. Previous research focuses on output data such as "likes" or reputation scores to measure performance (e.g., Pavlou and Gefen 2004; Kellogg et al. 2019) as a key metric of the matching algorithm. However, "likes" and reputation scores present accumulated data as a proxy for determining whether many others may favor a specific YouTube posting, Uber driver or Airbnb host, but fail to provide information on whether the content, service or material may be of interest to a particular individual, who may even have niche preferences or interests. Gathering fine-grained, individual behavioral process data allows platforms to personalize and tailor content and implement subtle reward systems that may nudge users into desirable behaviors through algorithmic control.

Second, we extend previous research (Choudhury and Sabherwal 2003; Kirsch 1997; Wiener et al. 2016) by presenting algorithmic control as qualitatively different from the types of control studied in existing information systems literature (e.g., Kirsch 1997; Wiener et al.

2016). In particular, algorithmic control seems to blur the traditional distinction between formal and informal control. It resembles formal control in its coerciveness, yet exhibits the fluidity of informal control. While some aspects of algorithmic control seem to reflect formal control elements (e.g., YouTube's community guidelines on content, which also comply with government regulations), other aspects of our constructs show overlaps with the definition of informal control. Indeed, the subtle modifying and influencing of individuals' behavior through targeted personalization and behavioral nudging, as theorized in our construct, are rather implicit. Furthermore, content creators criticize YouTube's automated censorship practices, which are perceived to be highly opaque and somewhat arbitrary. However, unlike the forms of informal control identified in previous literature, such as clan control or self-control, the informality of our constructs stems not from social interactions or beliefs shared by specific social groups (Kirsch 1997; Tiwana et al. 2010; Wiener et al. 2016), but purely from the technological and computational capacity to informally manipulate and influence individuals based on making personalized and targeted information available to them. Thus, algorithmic control can be considered to be a new and distinct form of control that transcends the distinction between formal and informal control. It can be interpreted as an "invisible hand," a metaphor used by Adam Smith to describe the free-market economy, but also appropriate to the context discussed here because it conveys the largely unseen and unobservable force of informal control elements.

Third, our research challenges the assumption that control mechanisms are inflexible and mechanistic (Kirsch 1997). Algorithmic control embodies characteristics of much greater levels of generativity. Since digital platforms profit from relatively easy access to data, and advances in digital and mobile technologies allow them to capture billions of data points on individuals' clicking and browsing behavior, location, personal networks and social interaction patterns, the platforms' control mechanisms can adjust dynamically to changing environments (Constantiou and Kallinikos 2015; Dourish 2016; Faraj et al. 2018; Gregory et al. 2020). Data on user behavior are fed into the algorithm, which then automatically implements control measures to shape user behavior, and user behavior is in turn traced and fed back into the algorithm. Therefore, human/algorithm interactions can be described as an endless, dynamic loop of mutual interdependence (see Faraj et al. 2018). One major implication is that, unlike other forms of control, algorithmic control can adjust to each individual consumer and provider, allowing the implementation of much more personalized and situational control measures, and potentially opening up opportunities to shape consumers' and providers' behavior in real time. This reaches far beyond access control through boundary resources, as addressed in previous research on digital platforms (Eaton et al. 2015; Ghazawneh and Henfridsson 2013; Karhu et al. 2018).

Managerial Implications

Policy makers are increasingly concerned about large tech platform companies' commercial practices that leverage their vast access to user data and new advances in digital technologies (Lustig et al. 2016; Zuboff 2019). While some formal control elements are applied, the platform companies' power stems from the relatively hidden, informal control exerted over platform participants through "algorithmic control." Informal control mechanisms may be much more powerful than formal control mechanisms, as they may remain unnoticed by most platform participants. We urge policy makers and platform participants to implement measures to make some of these platforms' "invisible" practices "visible," for instance by encouraging them to share more detailed, more understandable and more frequent information with their users about the extent to which they are monitored and how their personal data are being used. Rather than deleting or demonetizing content without sharing detailed insights into the reasons, regulations might potentially demand public disclosure of such information. Algorithmic control yields manifold benefits, ranging from automatically deleting inappropriate content in order to maintain child-friendly online environments, to valuing artists' contributions by implementing algorithmic filters that flag copyright breaches. Maintaining transparency about platform activities will increase the legitimacy of algorithmic control and enable platforms to leverage its benefits and full potential.

Agenda for Future Research

Research on algorithmic control is still in its infancy. In this research note, we provide a grounded conceptualization of "algorithmic control of consumers" and "algorithmic control of providers." We encourage future research on this construct in relation to other stakeholder groups, different platforms, non-platform contexts and other industries. Furthermore, the nomological net in which they are embedded requires significant attention. Future research might unveil the antecedents and consequences of algorithmic control, and address relationships between its sub-elements. While we seek to contribute first insights to this emerging research stream, we encourage future research to build on and extend our work.

REFERENCES

Adner, A., and Kapoor, R. 2010. "Value Creation in Innovation Ecosystems: How the Structure of Technological Interdependence Affects Firm Performance in New Technology Generations," *Strategic Management Journal* (31:3), pp. 306-333 (doi: 10.1002/smj.821).

Alexa Internet. 2020. "The Top 500 Sites on the Web," *Alexa Internet*, dynamic web page (https://www.alexa.com/topsites).

Andersen, J. V., Lindberg, A., Lindgren, R., and Selander, L. 2016. "Algorithmic Agency in Information Systems: Research Opportunities for Data Analytics of Digital Traces," in *Proceedings of the 49th Annual Hawaii International Conference on System Sciences*, T. X. Bui and R. H. Sprague (eds.), Piscataway, NJ: IEEE, pp. 4597-4605.

Aral, S. 2020. *The Hype Machine: How Social Media Disrupts Our Elections, Our Economy, and Our Health—and How We Must Adapt*, New York, NY: Random House.

Burrell, J. 2016. "How the Machine 'Thinks': Understanding Opacity in Machine Learning Algorithms," *Big Data & Society* (3:1), pp. 1-12 (doi: 10.1177/2053951715622512).

Burtch, G., Carnahan, S., and Greenwood, B. N. 2018. "Can You Gig It? An Empirical Examination of the Gig Economy and Entrepreneurial Activity," *Management Science* (64:12), pp. 5497-5520 (doi: 10.1287/mnsc.2017.2916).

Casadesus-Masanell, R., and Halaburda, H. 2014. "When Does a Platform Create Value by Limiting Choice?" *Journal of Economics & Management Strategy* (23:2), pp. 258-292 (doi: 10.1111/jems.12052).

Charmaz, K. 2014. *Constructing Grounded Theory*, 2nd edition, London: SAGE.

Choudhury, V., and Sabherwal, R. 2003. "Portfolios of Control in Outsourced Software

Development Projects," *Information Systems Research* (14:3), pp. 291-314 (doi: 10.1287/isre.14.3.291.16563).
Constantinides, P., Henfridsson, O., and Parker, G. G. 2018. "Introduction: Platforms and Infrastructures in the Digital Age," *Information Systems Research* (29:2), pp. 381-400 (doi: 10.1287/isre.2018.0794).
Constantiou, I. D., and Kallinikos, J. 2015. "New Games, New Rules: Big Data and the Changing Context of Strategy," *Journal of Information Technology* (30:1), pp. 44-57 (doi: 10.1057/jit.2014.17).
Curchod, C., Patriotta, G., Cohen, L., and Neysen, N. 2019. "Working for an Algorithm: Power Asymmetries and Agency in Online Work Settings," *Administrative Science Quarterly* (65:3), pp. 644-676 (doi: 10.1177/0001839219867024).
Cusumano, M. A., Gawer, A., and Yoffie, D. B. 2019. *The Business of Platforms: Strategy in the Age of Digital Competition, Innovation, and Power*, New York, NY: Harper Business.
De Reuver, M., Sørensen, C., and Basole, R. 2018. "The Digital Platform: A Research Agenda," *Journal of Information Technology* (33:2), pp. 124-135 (doi: 10.1057/s41265-016-0033-3).
Dourish, P. 2016. "Algorithms and their Others: Algorithmic Culture in Context," *Big Data & Society* (3:2), pp. 1-11 (doi: 10.1177/2053951716665128).
Eaton, B., Elaluf-Calderwood, S., Sørensen, C., and Yoo, Y. 2015. "Distributed Tuning of Boundary Resources: The Case of Apple's iOS Service System," *MIS Quarterly* (39:1), pp. 217-243 (doi: 10.25300/MISQ/2015/39.1.10).
Faraj, S., Pachidi, S., and Sayegh, K. 2018. "Working and Organizing in the Age of the Learning Algorithm," *Information and Organization* (28:1), pp. 62-70 (doi: 10.1016/j.infoandorg.2018.02.005).
Galliers, R. D., Newell, S., Shanks, G., and Topi, H. 2017. "Datification and its Human, Organizational and Societal Effects: The MARK Strategic Opportunities and Challenges of Algorithmic Decision-Making," *Journal of Strategic Information Systems* (26:3), pp. 187-190 (doi: 10.1016/j.jsis.2017.08.002).
Gawer, A. 2014. "Bridging Differing Perspectives on Technological Platforms: Toward an Integrative Framework," *Research Policy* (43:7), pp.1239-1249 (doi: 10.1016/j.respol.2014.03.006).
Gawer, A. 2020. "Digital Platforms' Boundaries: The Interplay of Firm Scope, Platform Sides, and Digital Interfaces," *Long Range Planning*, in press (doi: 10.1016/j.lrp.2020.102045).
Ghazawneh, A., and Henfridsson, O. 2013. "Balancing Platform Control and External Contribution in Third-Party Development: The Boundary Resources Model," *Information Systems Journal* (23:2), pp. 173-192 (doi: 10.1111/j.1365-2575.2012.00406.x).
Gioia, D. A., Corley, K. G., and Hamilton, A. L. 2013. "Seeking Qualitative Rigor in Inductive Research: Notes on the Gioia Methodology," *Organizational Research Methods* (16:1), pp. 15-31 (doi: 10.1177/1094428112452151).
Glaser, B. 1978. *Theoretical Sensitivity*, Mill Valley, CA: Sociology Press.
Gregory, R. W., Henfridsson, O., Kaganer, E., and Kyriakou, H. 2020. "The Role of Artificial Intelligence and Data Network Effects for Creating User Value," *Academy of Management Review*, in press (doi: 10.5465/amr.2019.0178).
Hagiu, A., and Spulber, D. 2013. "First-Party Content and Coordination in Two-Sided Markets," *Management Science* (59:4), pp. 933-949 (doi: 10.1287/mnsc.1120.1577).
Hong, Y., Wang, C., and Pavlou, P. A. 2016. "Comparing Open and Sealed Bid Auctions: Evidence from Online Labor Markets," *Information Systems Research* (27:1), pp. 49-69 (doi: 10.1287/isre.2015.0606).
Horton, J. J. 2017. "The Effects of Algorithmic Labor Market Recommendations: Evidence

from a Field Experiment," *Journal of Labor Economics* (35:2), pp. 345-385 (doi: 10.1086/689213).
Hu, X., Lin, Z., Whinston, A. B., and Zhang, H. 2004. "Hope or Hype: On the Viability of Escrow Services as Trusted Third Parties in Online Auction Environments," *Information Systems Research* (15:3), 236-249 (doi: 10.1287/isre.1040.0027).
Jaworski, B. J. 1988. "Toward a Theory of Marketing Control: Environmental Context, Control Types, and Consequences," *Journal of Marketing* (52:3), pp. 23-39 (doi: 10.1177/002224298805200303).
Karhu, K., Gustafsson, R., and Lyytinen, K. 2018. "Exploiting and Defending Open Digital Platforms with Boundary Resources: Android's Five Platform Forks," *Information Systems Research* (29:2), pp. 479-497 (doi: 10.1287/isre.2018.0786).
Kellogg, K., Valentine, M., and Christin, A. 2019. "Algorithms at Work: The New Contested Terrain of Control," *Academy of Management Annals* (14:1), online (doi: 10.5465/annals.2018.0174).
Kirsch, L. J. 1997. "Portfolios of Control Modes and IS Project Management," *Information Systems Research* (8:3), pp. 215-239 (doi: 10.1287/isre.8.3.215).
Lee, M. K., Kusbit, D., Metsky, E., and Dabbish, L. 2015. "Working with Machines: The Impact of Algorithmic, Data-Driven Management on Human Workers," in *Proceedings of the 33rd Annual ACM SIGCHI Conference*, B. Begole, J. Kim, K. Inkpen and W. Wood (eds.), New York, NY: ACM Press, pp. 1603-1612 (doi: 10.1145/2702123.2702548).
Lustig, C., Pine, K., Nardi, B., Irani, L., Lee, M. K., Nafus, D., and Sandvig, C. 2016. "Algorithmic Authority: The Ethics, Politics, and Economics of Algorithms that Interpret, Decide, and Manage," in *Proceedings of the 34th Annual CHI Conference on Human Factors in Computing Systems*, New York, NY: ACM, pp. 1057-1062.
Möhlmann, M., and Zalmanson, L. 2017. "Hands on the Wheel: Navigating Algorithmic Management and Uber Drivers' Autonomy," paper presented at the International Conference on Information Systems, December 10-13, Seoul, South Korea.
Newell, S., and Marabelli, M. 2015. "Strategic Opportunities (and Challenges) of Algorithmic Decision-Making: A Call for Action on the Long-Term Societal Effects of 'Datification,'" *Journal of Strategic Information Systems* (24:1), pp. 3-14 (doi: 10.1016/j.jsis.2015.02.001).
Ouchi, W. G. 1977. "The Relationship Between Organizational Structure and Organizational Control," *Administrative Science Quarterly* (22:1), pp. 95-113 (doi: 10.2307/2391748).
Ouchi, W. G. 1978. "The Transmission of Control through Organizational Hierarchy," *Academy of Management Journal* (21:2), pp. 173-192 (doi: 10.5465/255753).
Ouchi, W. G. 1979. "A Conceptual Framework for the Design of Organizational Control Mechanisms," *Management Science* (25:9), pp. 833-848 (doi: 10.1287/mnsc.25.9.833).
Parker, G. G., and Van Alstyne, M. W. 2005. "Two-Sided Network Effects: A Theory of Information Product Design," *Management Science* (51:10), pp. 1494-1504 (doi: 10.1287/mnsc.1050.0400).
Parker, J., Van Alystyne, M. W., and Choudary, S. P. 2016. *Platform Revolution: How Networked Markets Are Transforming the Economy and How to Make Them Work for You*, New York, NY: Norton and Company.
Pavlou, P. A., and Gefen, D. 2004. "Building Effective Online Marketplaces with Institution-Based Trust," *Information Systems Research* (15:1), pp. 37-59 (doi: 10.1287/isre.1040.0015).
Rice, S. C. 2012. "Reputation and Uncertainty in Online Markets: An Experimental Study," *Information Systems Research* (23:2), pp. 436-452 (doi: 10.1287/isre.1110.0362).
Rosenblat, A., and Stark, L. 2016. "Algorithmic Labor and Information Asymmetries: A Case Study of Uber's Drivers," *International Journal of Communication* (10), pp. 3758-

3784.
Sarker, S., Sahay, S., and Lau, F. 2002. "Teaching Information Systems Development Using 'Virtual Team' Projects," *Journal of Informatics Education and Research* (4:1), pp. 35-46.
Schildt, H. 2017. "Big Data and Organizational Design: The Brave New World of Algorithmic Management and Computer Augmented Transparency," *Innovation* (19:1), pp. 23-30 (doi: 10.1080/14479338.2016.1252043).
Tiwana, A. 2015. "Evolutionary Competition in Platform Ecosystems," *Information Systems Research* (26:2), pp. 266-281 (doi: 10.1287/isre.2015.0573).
Tiwana, A., Konsynski, B., and Bush, A. A. 2010. "Research Commentary—Platform Evolution: Coevolution of Platform Architecture, Governance, and Environmental Dynamics," *Information Systems Research* (21:4), pp. 675-687 (doi: 10.1287/isre.1100.0323).
Van Alstyne, M., Parker, G., and Choudary, S. 2016. "Pipelines, Platforms and the New Rules of Strategy," *Harvard Business Review* (94:4), pp. 56-60.
Wiener, M., Mähring, M., Remus, U., and Saunders, C. 2016. "Control Configuration and Control Enactment in Information Systems Projects: Review and Expanded Theoretical Framework," *MIS Quarterly* (40:3), pp. 741-774 (doi: 10.25300/MISQ/2016/40.3.11).
Wiener, M., Mähring, M., Remus, U., Saunders, C., and Cram, A. 2019. "Moving IS Project Control Research into the Digital Era: The 'Why' of Control and the Concept of Control Purpose," *Information Systems Research* (30:4), pp. 1387-1401 (doi: 10.1287/isre.2019.0867).
Yoo, Y., Boland, R., Lyytinen, K., and Majchrzak, A. 2012. "Organizing for Innovation in the Digitized World," *Organization Science* (23:5), pp. 1398-1408 (doi: 10.1287/orsc.1120.0771).
Yoo, Y., Henfridsson, O., Lyytinen, K. 2010. "Research Commentary—The New Organizing Logic of Digital Innovation: An Agenda for Information Systems Research," *Information Systems Research* (21:4), pp. 724-735 (doi: 10.1287/isre.1100.0322).
Zuboff, S. 2019. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*, London, UK: Profile Books.