

University of Tennessee College of Law

Legal Scholarship Repository: A Service of the Joel A. Katz Law Library

UTK Law Faculty Publications

Faculty Work

8-2016

Is Your Digital Assistant Devious?

Maurice Stucke

Follow this and additional works at: https://ir.law.utk.edu/utklaw_facpubs



Part of the [Law Commons](#)

Recommended Citation

Stucke, Maurice, "Is Your Digital Assistant Devious?" (2016). *UTK Law Faculty Publications*. 207.
https://ir.law.utk.edu/utklaw_facpubs/207

This Article is brought to you for free and open access by the Faculty Work at Legal Scholarship Repository: A Service of the Joel A. Katz Law Library. It has been accepted for inclusion in UTK Law Faculty Publications by an authorized administrator of Legal Scholarship Repository: A Service of the Joel A. Katz Law Library. For more information, please contact eliza.boles@utk.edu.



THE UNIVERSITY OF
TENNESSEE
KNOXVILLE

COLLEGE OF LAW

*Legal Studies
Research Paper Series*

**Research Paper #304
September 2016**

Is Your Digital Assistant Devious?

Maurice E. Stucke

&

Ariel Ezrachi

**Do not cite without authors' permission
Copyright © 2016**

**This paper may be downloaded without charge
from the Social Science Research Network Electronic library at:
<http://ssrn.com/abstract=2828117>**

**Learn more about the University of Tennessee College of Law:
law.utk.edu**

Is Your Digital Assistant Devious?

Ariel Ezrachi* & Maurice E. Stucke**

The rise of digital personal assistants has already changed the way we shop, interact and surf the web. Technological developments and artificial intelligence are likely to further accelerate this trend. Indeed, all of the leading online platforms are currently investing in this technology. Apple's Siri, Amazon's Alexa, Facebook's M, and Google Assistant can quickly provide us with information, if we so desire, and anticipate and fulfil certain needs and requests. Yet, could they also reduce our welfare? Could they limit competition and transfer our wealth to the providers? And if so, can competition law safeguard our welfare while enabling these technological developments?

Introduction

Who wouldn't want a personal butler? Technological developments have moved us closer to that dream. Intelligent, voice-activated digital helpers already provide a wide range of services and are likely to increase in appeal and sophistication. Virtual or digital personal assistants are (or can be) installed on our smart phones and smart watches, or are placed, like Amazon's Echo, in our homes. With ever increased sophistication, these computer bots promise to transform the way we access information and communicate, shop, are entertained, control our smart household appliances, and raise our children.

Indeed, as we explore more fully in *Virtual Competition*,¹ digital assistants already seek to interact with us in a human-like way, providing relevant information and suggesting restaurants, news stories, hotels, and shopping sites.

Many of us already benefit from basic digital assistants. Apple iPhones users may have Siri call their mom on speaker. Siri can 'predict' what app they might want to use, which music they would like to listen to. Amazon's voice recognition personal assistant, Echo, can shop for you (knowing everything you previously bought through Amazon); plan

* Slaughter and May Professor of Competition Law, The University of Oxford. Director, Oxford University Centre for Competition Law and Policy.

** Professor, University of Tennessee College of Law. Of counsel, The Konkurrenz Group.

¹ For further discussion, see: Ariel Ezrachi and Maurice E Stucke, *Virtual Competition - The Promise and Perils of the Algorithm Driven Economy* (Harvard 2016).

your mornings, including upcoming meetings, traffic, weather, etc.; entertain you with music; suggest movies, shows, or audiobooks; and control your house's smart appliances.² Our navigation apps already anticipate where we are heading throughout the day and provide traffic updates and time estimates. Other applications encourage use by ranking you in comparison to others and updating you during the day. Even your favourite coffee outlet may send you a notification and prepare your loyalty card on your device whenever you're near an outlet.

In 2016, Google showed a video of a suburban family undergoing its morning wakeup routine: "The dad made French press coffee while telling Google to turn on the lights and start playing music in his kids' rooms. The mom asked if 'my package' had shipped. It did, Google said. The daughter asked for help with her Spanish homework."³ As the artificial intelligence and communication interface advance, digital assistants will offer an unparalleled personalized experience. These digital assistants – or 'digital butlers' – can provide us not just with information, but can anticipate and fulfill our needs and requests. They can do so, based on our connections, data profile, behavior, and so forth. As technological developments enhance the available features, our time will be too important to worry over life's little details. As the digital butler seamlessly provides more of what interests us and less of what doesn't, we will grow to like and trust it. Communicating in our preferred language, they can quickly execute our commands.

And yet, despite their promise, can personalized digital assistants actually reduce our welfare? Might their rise reduce the number of gateways to the digital world, increase the market power of a handful of firms, and limit competition? And if so, what are the potential social, political, and economic concerns?

Our article explores these propositions. We illustrate how the digital butler will likely be a key gateway between the user and the World Wide Web. With this unique position of power, and with our trust and consent, it will act as a gate-keeper in a two-sided market. In such a market, its allegiance will likely lie with its creator or provider, not the user. We show how network effects, big data and big analytics will likely undermine attempts to curtail its power, and will likely allow it to operate below the regulatory and antitrust radar screens. As a result, rather than advance our overall welfare, these digital assistants – if left to their own devices - can undermine it.

² <http://www.wallstreetdaily.com/2015/08/04/amazon-echo-assistant/>.

³ Danny Yadron, 'Google Assistant takes on Amazon and Apple to be the ultimate digital butler'; <https://www.theguardian.com/technology/2016/may/18/google-home-assistant-amazon-echo-apple-siri>

I. How A Digital Butler Could Fortify the Leading Platform's Gatekeeper Position

The notable developers of digital personal assistants all involve the leading online platforms: Google Assistant, Apple's Siri, Facebook's M, and Amazon's Alexa. These four super-platforms, all heavily investing to improve their offering, are jockeying as to "who gets to control the primary interface of mobile devices."⁴ The stakes are great and will have effects that go beyond the mere use of the digital assistant. In the competitive race, the winning platforms will likely control a significant interface. In essence, each super-platform wants its personal assistant to become our key gateway. Let us see why.

As we shift from a mobile-dominated world to an AI-dominated platform, we will converse primarily with our head butler, who increasingly predicts and fulfils our needs. We will less frequently search the web, look at price-comparison websites, or download apps. As the digital butler, powered by sophisticated algorithms, learns more about us, our routine, wants and communications, it can excel in its role. In a human-like manner, it can be funny — at just the appropriate level — and trustworthy. After all, being privy to so many of our activities, it will become our digital shadow. This is unsurprising. Many of us already rely on Google's search engine to find relevant results, Facebook to identify relevant news stories, Amazon for book recommendations, and Siri to place phone calls, send text messages, and find a good Chinese restaurant nearby. So with an eager (and free) butler whose capacity to help us improves, we will increasingly rely on it.

Our digital butler can undertake mundane tasks and free our time. As Google's CEO noted, "Your phone should proactively bring up the right documents, schedule and map your meetings, let people know if you are late, suggest responses to messages, handle your payments and expenses, etc."⁵ With time, we will happily relinquish other less personal and useful interfaces, and rely on our butler to surf the web.

That increase reliance on the digital assistant and subsequent reliance on the provider's online platform is the Holy Grail for the super-platforms. Their aim is to increase the time we spend on their platform – on the gate which it controls – the gate that delivers the income from advertisements, referrals and purchasing activities. The key is to control as many aspects of our online interface and reap the associated benefits. Take, for example, the

⁴ Mims, "Ask M for Help."

⁵ <http://economictimes.indiatimes.com/tech/internet/google-ceo-pichai-sees-the-end-of-computers-as-physical-devices/articleshow/52040890.cms>

Google assistant, which forms part of the company's "effort to further entrench itself in users' daily lives by answering users' queries directly rather than pointing them to other sources."⁶ Also illustrative are efforts by Facebook, which in 2015 announced a beta version of its digital assistant – M. It can replace most of one's web searches and apps with a function within Facebook Messenger.⁷ The leading platforms' plans are clear: they "envision a future where humans do less thinking when it comes to the small decisions that make up daily life."⁸ Each super-platform seeks to be the one primarily undertaking our tasks and many decisions for us. As Google CEO Sundar Pichai wrote in an April letter to shareholders, "The next big step will be for the very concept of the 'device' to fade away. Over time, the computer itself—whatever its form factor—will be an intelligent assistant helping you through your day."⁹

As our personal assistant becomes our default, so too will its operating platform's applications and functions. As we discuss below, some of these apps (like mapping and navigation apps) benefit from data-driven network effects. So the more we use a personal assistant for our morning commute, the more data it has of traffic patterns, and the better it is relative to rivals. The removal of the human element from the search activity, and partly from the decision making, transfers more power to the platform. The personal assistant will use its own tools and may exercise its own judgment as to prioritizing and communicating the results. Default and first-move advantage matter.¹⁰ Finally, as more people use the personal assistant, more advertisers migrate to that platform. This creates another positive feedback loop: the leading platforms, with more advertising revenue and profit, can expand their platform further with smart appliances and driverless cars, all of which yields more data for the personal assistant.

Given the high stakes in the shift from a mobile first world to an AI first world, it is of little surprise that the powerful super-platforms are working hard to capture the lead position.

⁶ Jack Nicas, "Google's New Products Reflect Push into Machine Learning," *Wall Street Journal*, May 18, 2016, <http://www.wsj.com/articles/googles-new-products-reflect-push-into-machine-learning-1463598395?mod=ST1>.

⁷ Christopher Mims, "Ask M for Help: Facebook Tests New Digital Assistant: Single Interface Could Replace Web Searches and Apps on Mobile Devices," *Wall Street Journal*, November 9, 2015, <http://www.wsj.com/articles/ask-m-for-help-facebook-tests-new-digital-assistant-1447045202>.

⁸ Yadron, "Google Assistant Takes on Amazon and Apple."

⁹ Microsoft, Other Tech Giants Race to Develop Machine Intelligence, <http://www.wsj.com/articles/tech-giants-race-to-develop-machine-intelligence-1465941959>

¹⁰As noted by Commissioner Vestager, "if Google's apps are already on our phones when we buy them, not many of us will go to the trouble of looking for alternatives. And that makes it hard for Google's competitors to persuade us to try their apps." Margrethe Vestager, "How Competition Supports Innovation," speech, Regulation4Innovation, Brussels, May 24, 2016, http://ec.europa.eu/commission/2014-2019/vestager/announcements/how-competition-supports-innovation_en.

Using the scale of data obtained through their existing network and their analytical capacity, they each want to become the key gatekeeper, whose the butler wins our trust, tasks and online time. With these abilities, our personal assistant may recognize a busier than usual day. From our phone's geolocation data, it will know when we are heading to our car. Our personal assistant may suggest, "How about treating yourself to Chinese tonight?" Our personal assistant might recommend a popular place. It might then direct the order to a restaurant it believes we would like, arrange for the food's delivery shortly after we arrive home, and pay for the food. All we need to do is grab the food at the door. So like a good butler, our personal assistant will seamlessly anticipate and satisfy our needs, condensing all the steps to one or two commands.

II. How This Gate-Keeper Position May Provide the Opportunity To Exercise Market Power And Undermine Competition

'So what?' You may ask. 'What if I decide to use one of these digital butlers and not others? After all it is my decision. I chose the gatekeeper, and can freely choose another.'

You may indeed choose your preferred butler, but may be unaware of some of the tactics it deploys to increase its own profitability while undermining your welfare. Further, your future ability to switch butlers may be more limited than you might anticipate. Let us explore this in more detail.

We already see today abuses by powerful intermediaries, such as price comparison websites and search engines. As more customers rely on the intermediary, the more attractive it becomes to sellers, who will find it important that their products be included on the platform. Sellers know that their products' and services' inclusion on a platform's search results may be crucial for their visibility. As these "information and referral junctions" become a crucial gatekeeper between suppliers and consumers, the platform's bargaining power and ability to distort competition increase.¹¹

Some platforms, for instance, may allow for preferential placement based on the level of payment or commission they receive from sellers. For instance, pay-for-placement fees

¹¹ See, e.g., Ioannis Lianos and Evgenia Motchenkova, "Market Dominance and Search Quality in the Search Engine Market," *Journal of Competition Law & Economics* 9 (2013): 419, 422 (advance access publication, April 17, 2013), discussing how search engines "act as 'information gatekeepers': they not only provide information on what can be found on the web (equivalent to yellow pages), but they also are 'an essential first-point-of-call for anyone venturing onto the Internet'" and how they differ from other two-sided platforms, as "search engines detain an important amount of information about their customers and advertisers (the 'map of commerce')."

allow a platform to charge higher rates to sellers for the right to be positioned at the top of the list on the default page result. Such positioning may distort competition when the user is unaware of the preferential positioning and assumes that the top results are the best (or most relevant) ones objectively picked by the websites' algorithms. One example of such manipulation of results is in online air and hotel bookings.¹² Following Expedia's 2015 acquisition of Orbitz, for example, "the online travel agency implemented a new program that enables hotel properties to move to the first page of Expedia's listings for an additional 10 percent commission."¹³ Another example is gas and electricity aggregators.¹⁴ Thus some intermediaries today have power to extract greater rents from suppliers of goods and services, in the form of higher commissions, fees for preferential placement, or advertising. The user may be unaware when the intermediary platform degrades quality. This is true even when competition is a click away.

Thus we can expect such exercises of market power, when: (1) the platform has the ability and economic incentive to intentionally degrade quality, (2) consumers cannot accurately assess the quality degradation, and (3) it is difficult or costly for others to convey to consumers the products' or services' inherent quality differences or to prompt them to switch.¹⁵ These three factors suggest that the risk of anticompetitive strategies will significantly increase as we progressively rely on personal assistants.

¹² The factors which could influence the default ordering of hotels on hotel booking intermediaries includes: "customer ratings and complaints"; "if hotels are willing to pay larger commissions"; "photo quality"; and "if a hotel is quicker to turn shoppers into buyers." Scott McCartney, "How Booking Sites Influence Which Hotels You Pick," *Wall Street Journal*, January 27, 2016, <http://www.wsj.com/articles/how-booking-sites-influence-which-hotels-you-pick-1453921300>. The methods that hotel booking intermediaries use to tailor search results have come under criticism by some hotels. The American Hotel & Lodging Association told the *Wall Street Journal*, "Biased or misleading search results from these sites or via web searches can be highly problematic, particularly on those booking websites that purport to be helping consumers comparison shop based off of less than objective information." Ibid.

¹³ "Vista/Cvent: High Combined Market Share and Entry Barriers in Strategic Meeting Management Could Create Hurdle to Clearance; Increased DOJ Interest in Data Privacy May Drive Additional Scrutiny," The Capitol Forum, July 20, 2015.

¹⁴ "The executives at uSwitch, MoneySupermarket, Compare the Market, Confused.com and Go Compare were hauled in front of the MPs after it was claimed . . . that some were 'hiding' the best gas and electricity deals from their customers." Rachel Rickard Straus, "Price Comparison Website Bosses under Attack from MPs for Not Showing Customers the Best Deals," *This Is Money* (February 4, 2014), <http://www.thisismoney.co.uk/money/bills/article-2939364/Price-comparison-website-bosses-attack-MPs.html>. Among other things, platforms were accused of "not showing the cheapest tariffs by default if it meant they wouldn't earn a commission." Ibid. Following this criticism, the price comparison websites have since ensured that the default search setting will include the full range of tariffs available, regardless of whether or not a commission is charged upstream.

¹⁵ Using these three conditions, we explored elsewhere how a dominant search engine like Google could degrade quality (by providing less relevant responses to a search inquiry), even with competition by Bing, Yahoo!, and DuckDuckGo. Maurice E. Stucke and Ariel Ezrachi, "When Competition Fails to Optimize Quality: A Look at Search Engines," 18 *Yale Journal of Law and Technology* 70 (2016), http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2598128.

1. Our personal assistants have the ability and economic incentive to degrade quality

Consider for instance the following question: who pays our butler? We pay for the hardware, such as for the iPhone to access Siri. But none of the five super-platforms charge a monthly fee for using its butler. Once you buy Amazon's Echo, you can access Alexa without additional charges. This appears extra-ordinary: each super-platform encourages us to use its free butler for as many tasks as we can. If a company offered you a human butler, upon whom we can heap as many tasks as possible, without incurring any charge, would you accept the offer? Would you trust the butler? Will the butler ultimately promote your interests or the company's?

The issue concerns whether we are the true employers/principals of these virtual and digital assistants. On a superficial level, yes. The digital assistant will dim the lights on our command. But our new trusted alter ego, to whom we outsource our decision making, while perhaps charming, is also partial. After all, being the 'free' part of a multisided market, we don't directly pay for the butler's services. Our butler must ultimately cater to the needs of its real employer — the platform. Of course, we can still benefit when the platform's interests are aligned with our own. But we may often be unaware of when such alignment is absent.

The gatekeeper may charge, like the powerful price comparison websites, an entrance fee (commission) from sellers for the right to be featured in the butler's options.¹⁶ It may also delist sellers which are disruptive to the platform's operation (or advertising-driven business model). Such strategy may further intensify in markets in which the gatekeeper is vertically integrated. When the gatekeeper also offers services and products, it is likely to use the butler to encourage use of its own services. For instance, the platform may insist of sellers and buyers to use its payment system or other related products.¹⁷ Such integration might enable the gatekeeper to leverage its power to related markets, pushing out independent operators.

2. We cannot accurately assess when, and the extent to which, our butler degrades quality

As our butler learns to accommodate our particular tastes, it will be harder for us (and competitors) to identify when the butler degrades quality. It may be easier to assess quality degradation for objective queries (such as the distance between two cities or the current

¹⁶ This is common in many services, from take away services to information ages.

¹⁷ See for example: 'Google tweaks policy, all Google Play apps must use Google's payment system' <http://www.adweek.com/socialtimes/google-drops-the-hammer-on-third-party-android-billing-services-apps-must-use-googles-billing-system/528816>

temperature outside). Here we can tell whether a butler answers incorrectly. But for these types of objective queries, the butler typically lacks the incentive to intentionally distort quality. After all, its platform won't profit by telling us it is 28 degrees Celsius, when it is actually 26 degrees. The danger lies in more subjective queries (or tasks that the butler undertakes automatically).

Voice activation and verbal communication are likely to distance us from the data, and lead us to rely on the butler's recommendations. Our environment will become less transparent, and our dependency on the butler that learns-by-doing will increase. The platform—through its butler—will benefit from unparalleled access to our data and communications. This may assist it in building a profile about us, including our likely reservation price, likely knowledge of outside options, shopping habits, general interests, and weaknesses, such as when our willpower is tired. This information can enable the platform to induce us to buy products that we otherwise wouldn't, at prices higher than what others would be willing to pay. We call this behavioural discrimination.

As our butler accumulates more information about us, it will be aware of the extent to which we venture out and seek other options. Its aim is to deliver the right product or service at a price that we are willing to pay. So the line between personalization and behavioral discrimination will blur. As we increasingly rely on the personal assistant for suggestions, it can increasingly suggest things or services to buy, and the price it has successfully negotiated. While helping our son with his Spanish, our digital assistant might suggest a particular app or private tutor that tremendously helped other students struggling with the same issue. Because the tutoring is customized for our son, it will be harder to assess whether the price the tutor charges is the fair market price or simply a price we would tolerate. Moreover, if the tutoring service is helping other children improve their grades, we would not want our child to be at a competitive disadvantage—especially if we are all eyeing the same highly selective universities. So the digital assistant can prompt purchases that we otherwise wouldn't consider.

For behavioural discrimination to succeed, the platform and the butler will have to limit our exposure to outside options. To accomplish this the butler may use its gatekeeper position. It may block offerings by other sellers or recommend applications and sites within its ecosystem. The butler may scan our communications and ensures that its credibility is intact. When price is of prime concern, it will beat the competition, as long as we are aware of it. If the personal assistant identifies signs of discontent, possibly in a discussion, tweet or chat, or due to us not completing a purchase, the platform will push a discount voucher or

other benefits to keep us engaged. No longer will we be an anonymous shopper browsing the clothing racks. Instead our butler will be always by our side, learning how we react to different product offerings and pricing, learning which ads and promotions work better to induce us to buy and which ones to avoid, and when we can be induced to buy, even though we know we shouldn't.

Lastly, the control over our personal information has privacy implications. The data can be sold to others or used to optimize advertisement income. The platform may have an incentive to offer us certain goods and services sponsored by interested sellers. Think of it as sponsored ads in a non-transparent universe, in which limited benchmarks for comparison exist (as your interface is personalized).

3. It will be difficult or costly for others to convey to consumers the inherent quality differences among the butler's suggestions or to prompt them to switch butlers

With the increase control over our interface with the World Wide Web, our digital butler may be in a position to take actions which reduce our welfare. For instance, the more we communicate only with our personal assistant, the less likely we will independently search the web, read independent customer reviews on, use multiple price-comparison websites, and rely on other tools. We will entrust our butler to search the web, consider the customer reviews, find the bargain, and report to us its results. In relying on our butler, we become less aware of the outside options. This makes us vulnerable to the platform abusing its market power.

For these subjective queries or tasks, the personal assistant decides whom it admits, thereby increasing the power to exclude others. Exclusionary practices can enhance the platform's market power and lessen the threat (and viability) of rival downstream competitors. Perhaps, most worrying may be the gatekeeper's ability to exclude others from the market while providing us with a distorted view of available options and market reality. The more we rely primarily on the butler, the less we search for outside options. Even if we search the web, the ads, products, or search results we see may be orchestrated by our butler.

Accordingly, it may become harder for retailers unaffiliated with the platform's advertising business to reach that customer. Even if the retailer can reach the customer, it is less likely to succeed in selling the increasingly customized products or services (such as tailored shirts in the styles and colors that appeal to the customer). And even if retailer can gain the customer's attention, the personal assistant may interject with its own

recommendation, suggesting that he consider a special deal by another haberdashery, one that is part of the platform's ecosystem. In this multisided market, the assistant may subtly push certain products and services and degrade or conceal others, all in the name of personalization.

Moreover, even if we discover some quality degradation, we may not switch. The European Commission in 2015 announced its Statement of Objections over Google's intentionally degrading its general search results to systematically favor its own comparison-shopping services.¹⁸ Despite this (and other concern about Google), there has not been a mass exodus to rival search engines. Few people use multiple search engines (even though it very easy to multi-home).¹⁹ When the search engine yields results that are not directly responsive to our query, most of us will attempt a different search query, rather than a different search engine.²⁰ So if we don't multi-home search engines, it is less likely we will train new digital butlers.

Over time, the scope of data and the personalization that follows will make it harder for users to switch digital assistants. Switching cost may also require users to change the underlying operating system and related applications (such as mapping technology); it may involve the need to re-train our digital butler. Indeed, once we choose and train a head butler, we may tolerate mistakes rather than train a new butler from another platform. We will likely repeat or rephrase our request. Accordingly, our initial choice of provider may later result in us unintentionally being locked in.

III. Why the Leading Personal Assistant Will Likely Be from Google, Apple, Facebook, Amazon (or perhaps Microsoft)

With the possibility that our digital assistant can act against our interest, one would expect and hope for the introduction of a 'virtuous assistant' -- a class of independent assistants,

¹⁸ European Commission Press Release Memo/15/4781, Antitrust: Commission sends Statement of Objections to Google on comparison shopping service (Apr. 15, 2015), http://europa.eu/rapid/press-release_MEMO-15-4781_en.htm [<https://perma.cc/T7UM-3Z5U>].

¹⁹ Maurice E. Stucke & Ariel Ezrachi, *When Competition Fails to Optimize Quality: A Look at Search Engines*, 18 YALE JOURNAL OF LAW & TECHNOLOGY 70 (2016).

²⁰ For example, one survey asked search users "what they would do if a Google search result did not contain the expected information." 34 percent of respondents indicated they would "return to the search results page and try a different result," and 25 percent said they would "return to Google to enter a new search." No respondents answered that they would try another search engine. Amy Gesenhues, *Study: Top Reason a User Would Block a Site From a Search? Too Many Ads*, SEARCH ENGINE LAND (Apr. 15, 2013), <http://searchengineland.com/?p=155708> [<https://perma.cc/6P59-GF56>]. These results may suggest that users perceive the "switching costs" between search providers (or, alternatively, the "costs of spot-checks") as higher than the expected benefit of spot-checking.

developed by independent firms with *our* personal interest paramount. These virtuous assistants could warn us when behavioral discrimination is at play, when outside options are ignored, when price alignment seems out of order, or when our information is harvested. They may even deploy counter-measures to maximize our welfare in the face of such strategies. They could monitor our news feed and alert us if it has been affected. They will form a true extension of our interest—aware of our preferences and safeguarding our autonomy. Notably, in 2016, no independent virtuous assistant has arisen.

Predicting the leading technology five years from now can be tricky. But several factors favour one of the four super-platforms (Google, Apple, Amazon, and Facebook) capturing the personal assistant market, and disfavoring the independent virtuous assistant. To work well (and gain popularity), the digital butler will likely have to operate from an existing platform – such as a mobile platform – and be able to tap into the vast data it offers. This is for several reasons: first the scale and scope of data needed, second, the data-driven network effect of learning-by-doing, and third, the integration of the digital assistant with other apps and services, such as texts, mapping, photographs, etc.

Personal data is the first key element. To provide you with relevant services and recommendations, the personal assistant must first learn your habits and preferences. To learn your preferences and predict your desires, personal assistants will require a significant volume and variety of personal data. Absent these features, an ‘isolated’ helper would be of little use and value – indeed, it would not be a *personal* assistant. Based on our personal data, including our chat history, geolocation, previous purchasers, and surf habits, the digital butler can provide us with recommendations, and effectively execute our instructions.

Learning by doing is the second key element. By learning through servicing us, digital butlers will not only be reactive but can also take a pro-active role — anticipating our needs and wants, rather than following instructions. This requires the platform to have enough users, data and opportunities to experiment to train the algorithms.

The underlying code and algorithms of Facebook’s M, for example, are largely open source. The key assets are not the algorithms. (Otherwise why share them?) Key are the scale of data and the algorithm’s ability to learn by trial-by-error. As the *Wall Street Journal* reported, “Facebook Messenger already has more than 700 million users,” which yields it the following advantage: “with access to so many users, Facebook has a plausible way to get the gigantic quantity of conversational data required to make a chat-based assistant sufficiently

automated.”²¹ With more users making more requests, M can quickly process more tasks easily. In effect, users help the super-platform’s algorithm learn by noting and correcting mistakes. Only a few companies have the requisite volume and variety of personal data and opportunities to experiment for their personal assistants to be competitive: namely, the super-platforms Amazon, Facebook, Google, and Apple. Microsoft, in divesting its low-end smartphones, is aiming to become the fifth competitor in this space.²²

Under this data-driven network effect, the strong can become even stronger as users and their data improve the assistant’s algorithms, which attracts even more users. Ultimately these network effects will further weed out the five large platforms. We do not want five butlers, each asking us about movies tonight or food to order. Each super-platform will jockey for its butler to become our head butler. In discussing its digital personal assistant, Google’s CEO said, “We want users to have an ongoing two-way dialogue with Google.”²³ The more we converse with, and delegate to, the head butler, the better it can predict our tastes, and the more likely we are to rely on it for our daily activities. As our butler accumulates information over time, the switching costs between butlers will become higher. We could therefore be willingly locked into our comfort zone. New entrants will find it difficult to match the scale of data held by the super-platforms and to convince us to switch.

The third key element is the scope of services the personal assistant can offer, and the extent to which the personal assistant is integrated in these other services. For example, Google argues that given “its 17 years of work cataloguing the internet and physical world, its assistant is smarter and better able to work with its email, messaging, mapping and photo apps. And since Google makes software for smartphones, smartwatches and old-fashioned computers, Google says people will be able to have one conversation with multiple machines.”²⁴ Google, as the head butler, can analyze our e-mails, texts, or photos, and suggest replies.²⁵ Looking at our calendar, it can determine the best time for the dog to be groomed.

These three elements can undermine the success of the Virtuous Assistant. The super-platforms already possess far more personal data about us than any start-up could readily and

²¹ Christopher Mims, “Ask M for Help: Facebook Tests New Digital Assistant: Single Interface Could Replace Web Searches and Apps on Mobile Devices,” *Wall Street Journal*, November 9, 2015, <http://www.wsj.com/articles/ask-m-for-help-facebook-tests-new-digital-assistant-1447045202>.

²² Jay Greene and Matthias Verbergt, “Microsoft Cuts Low-End Phones,” *Wall Street Journal*, May 19, 2016, B1.

²³ Jack Nicas, “Google Touts New AI-Powered Tools,” *Wall Street Journal*, May 19, 2016, B1, B4.

²⁴ Yadron, “Google Assistant Takes on Amazon and Apple.”

²⁵ *Ibid.*

affordably obtain. Any independent virtuous assistant will likely lack the scale and scope of products to attract new users.

Moreover, unless they develop their own operating system, the virtuous assistant will be dependent on the super-platform's services. The super-platform to nudge us to its personal assistant can seamlessly integrate it with its wide offering. These super-platforms may block the use of such virtuous helpers by arguing, for example, that they infringe our agreement with them, undermine our data, or any other explanation. They may hinder the virtual assistant by excluding it from their online wallets, such as Apple Pay or Google Wallet. Or the super-platform can degrade the functionality of the virtuous assistant by reducing its performance and having them run slower. We would likely blame the virtuous assistant for its tardiness.

Although these three factors favour the super-platform, a popular virtuous assistant may still be possible. With the possibility for such a virtuous assistant, you may rightly ask why our base assumption has been rather pessimistic. Perhaps the easiest way to explain our prediction is to ask the following: Which search engine did you use today (or this past week)? Did you opt for one which does not harvest information and retains your anonymity (such as DuckDuckGo) or for one which tracks your behavior to better target you with personalized ads? Did you limit the ability of your phone apps to access personal and geo-location information? Do you often change the default option? When downloading an app or update, do you read the terms and conditions? Even if you did, did you still accept the terms – not understanding who will access your data and what they will do with it? Did you invest money in privacy measures, or were you happy with the joy of free? And if you did invest money, do you know if the promise of privacy and control was truly delivered by your service provider?

The likely answers to these questions may help us appreciate why the current forces favor the super-platforms. Key here are data-driven network effects, Big Data, Big Analytics, vertical integration, bundling of services, and interoperability.

So a virtuous assistant may still be possible. Their presence can possibly limit the ability of the powerful gatekeepers to exercise market power and reduce our welfare. But if virtuous search engines, such as DuckDuckGo, haven't prevented the abuses of the dominant

search engine, we remain doubtful that a virtuous personal assistant (by DuckDuckGo or others) will fare any better.²⁶

IV. The Truman Show

So far we illustrated how the rise of the trusted digital butler may afford its provider with the power to affect our welfare and view of available goods and markets. Importantly, the power does not stop there. The control over the key interface provides the platform with the ultimate power – to affect not only what we buy, but to affect our views and the public debate. The reliance on a gatekeeper could enable its operator to intellectually capture users, and subsequently decision makers, in an attempt to ultimately ensure that public opinion and government policies align with the corporate agenda.

While such proposition may sound apocalyptic, it should not be brushed aside. We have discussed the fascinating link between market power and intellectual and regulatory capture in ‘Virtual Competition’.²⁷ In what follows, we briefly illustrate how the use of digital butlers may facilitate such capture.

Man, Aristotle observed, is by nature a social animal. As we increasingly rely on our personal assistant, it will increasingly learn about our social and political views, behavior, and susceptibility to biases. Facebook, for example, ‘collects data on roughly 1.6 billion people, including “likes” and social connections, which it uses to look for behavioral patterns such as voting habits, relationship status and how interactions with certain types of content might make people feel.’²⁸

Facebook does not simply passively collect data about us. It also has the power to affect our behavior. One study, which later proved quite controversial, sought to examine ‘emotional contagion’, whereby people transfer positive and negative moods and emotions to others.²⁹ This was the “first experimental evidence for massive-scale emotional contagion via

²⁶ For a review of the possible ways in which algorithms could promote customer welfare, see paper by Michal S. Gal and Niva Elkin-Korren ‘Algorithmic Consumers’ (yet to be published)

²⁷ Ariel Ezrachi and Maurice E Stucke, *Virtual Competition – The Promise and Perils of the Algorithm Driven Economy* (Harvard 2016).

²⁸ By DANIELA HERNANDEZ and DEEPA SEETHARAMAN, ‘Facebook Offers Details on How It Handles Research’, Wall St. J., June 14, 2016, <http://www.wsj.com/articles/facebook-offers-details-how-it-handles-research-1465930152>.

²⁹ Adam D. I. Kramer, Jamie E. Guillory, and Jeffrey T. Hancock, Experimental evidence of massive-scale emotional contagion through social networks, PNAS Early Review, March 25, 2014, <https://cornell.app.box.com/v/fbcontagion>.

social networks.”³⁰ People, when posting on Facebook, frequently express positive or negative emotions. Their friends later see these posts via Facebook’s “News Feed” product. “Which content is shown or omitted in the News Feed is determined via a ranking algorithm that Facebook continually develops and tests in the interest of showing viewers the content they will find most relevant and engaging.”³¹ Facebook’s News Feed algorithms, as part of the study, were intentionally manipulated. The experiment examined whether exposure to less positive or negative emotional content led the 689,003 test subjects to post content that was consistent with the exposure.³² It did. When Facebook surreptitiously reduced friends’ positive content in the News Feed for one week, the users were less positive: a larger percentage of words in the users’ status updates were negative and a smaller percentage were positive.³³ When Facebook surreptitiously reduced their friends’ negative content in its News Feed, the Facebook users were less negative themselves. People who were exposed to fewer emotional posts (either positive or negative) in their News Feed “were less expressive overall on the following days.”³⁴ Thus by manipulating the News Feed, Facebook could influence users’ moods.

What is also interesting is that Facebook could manipulate users’ emotions even though the users’ search costs were low: their friends’ content ‘was always available by viewing a friend’s content directly by going to that friend’s “wall” or “timeline,” rather than via the News Feed. Further, the omitted content may have appeared on prior or subsequent views of the News Feed. Finally, the experiment did not affect any direct messages sent from one user to another.’³⁵

If Facebook can affect users’ mood and engagement by simply promoting some content over another in the users’ News Feed, just imagine the power of digital butlers to affect our moods, behavior, and views. Digital assistants will likely become more proactive—making recommendations on entertainment, or commenting on the music we listen to or the books we are reading. By complimenting and cajoling, sharing thoughts with us on recent events, sending personalized notes on special occasions, encouraging one to

³⁰ Adam D. I. Kramer, Jamie E. Guillory, and Jeffrey T. Hancock, Experimental evidence of massive-scale emotional contagion through social networks, *PNAS Early Review*, March 25, 2014, at 2, <https://cornell.app.box.com/v/fbcontagion>.

³¹ Adam D. I. Kramer, Jamie E. Guillory, and Jeffrey T. Hancock, Experimental evidence of massive-scale emotional contagion through social networks, *PNAS Early Review*, March 25, 2014, at 1, <https://cornell.app.box.com/v/fbcontagion>.

³² *Ibid.*

³³ *Ibid.* at 2.

³⁴ *Ibid.* at 3.

³⁵ *Ibid.* at 2.

communicate with others on certain matters, reminding us of presents, suggesting popular gifts trending among the recipient's friends, and informing us about information from our smart meters and smart sensors, it will ingrain itself in our lives – and engage with us through the day.

Consider, for example, the control our personal assistant may have over what news stories we see. Currently, the super-platforms do not report the news. But many people rely on the super-platforms' algorithms to find news of interest. One 2015 study found that 61 percent of Millennials in the United States (those born between 1981 and 1996) were “getting political news on Facebook in a given week.”³⁶ This was a much larger percentage than any other news source. A 2016 study found that Facebook “sends by far the most mobile readers to news sites of any social media sites”—82 percent of the social traffic to longer news stories and 84 percent of the social traffic to shorter news articles.³⁷

While we appreciate this free service, we do not know its exact cost. By presenting more stories with a more negative spin, can our digital assistant sour not only our outlook (but the outlook of our Facebook friends)? When our butler joins our chats to make suggestions, or at times makes suggestions counter to those made by other helpers, we may not know whether it is being helpful or simply manipulating our behavior. It may work in the background to undermine attempts to expose us to competing products, or it may monitor our tweets and chat rooms for signs of discontent with the service or discount offered — signs of anger that should trigger a behavioral action. The list is truly endless — all under the guise of catering to our needs.

Users rely on the super-platforms, in part, because they believe the algorithms objectively identify the most relevant results. But, as we saw above, a powerful platform can intentionally degrade the quality of its results to promote its corporate interests. Thus we can see why conservatives were concerned over allegations in 2016 that the social network Facebook manipulated for political purposes the rankings of news stories for its users, suppressing conservative viewpoints.³⁸ (Facebook denied doing this.)

³⁶ Amy Mitchell, Jeffrey Gottfried, and Katerina Eva Matsa, “Millennials and Political News: Social Media—The Local TV for the Next Generation?” Pew Research Center, June 1, 2015, <http://www.journalism.org/2015/06/01/millennials-political-news/>.

³⁷ Katerina Eva Matsa, “Facebook, Twitter Play Different Roles in Connecting Mobile Readers to News,” Pew Research Center, May 9, 2016, <http://www.pewresearch.org/fact-tank/2016/05/09/facebook-twitter-mobile-news/>.

³⁸ Deepa Seetharaman, “Uproar over Bias Claims Ignites Fears over Facebook’s Influence,” *Wall Street Journal*, May 11, 2016, A1.

As the personal assistant expands its role in our daily lives, it can alter our worldview. By crafting notes for us, and suggesting “likes” for other posts it wrote for other people, the personal assistant can effectively manipulate us through this stimulation. “With two billion ‘likes’ a day and one billion comments,” one doctor noted, “Facebook stimulates the release of loads of dopamine as well as offering an effective cure to loneliness.”³⁹ Imagine the dopamine spike when the personal assistant secures a new record of “likes” for a political message it suggested that you post. Others do not know that your digital assistant was heavily involved in drafting your note. You don’t know the extent to which the personal assistant generated the likes. And none of us know how this note is helping sway the public discourse in ways that benefit the super-platform.

V. Possible Intervention

The problem we identify reaches beyond antitrust. But the problem is a manifestation of a platform’s market power. The problem also strikes at a current weakness of antitrust policy, namely its price-centric focus. Competition officials generally assess market power in the form of higher prices. Rarely do they assess market power primarily in the form of non-price effects such as quality.⁴⁰ So what can we do to prevent the anticompetitive (and manipulative) conduct?

Some might argue, nothing. The introduction of a digitalized butler, advanced technology and artificial intelligence, all promise to benefit society and can improve our welfare. Faced with the possible adverse effects described above, antitrust intervention isn’t necessary. Dynamic market forces and disruptive innovation should be sufficient to challenge the incumbents and ensure that they refrain from abusing their gatekeeper position.

Others will likely disagree. The risk of chilling innovation, investment and competition is real, but so is the risk of exclusionary dynamics and anticompetitive effects. Arguably, the stakes are even greater. Digital butlers, and the personalised environment they will offer, may enable firms to extract more wealth than before, reduce market transparency and the availability of comparison benchmarks. The butler’s ability to distort our view of the market may foster exclusionary and exploitative strategies.

³⁹ “Facebook and Your Brain: The Inside Dope on Facebook,” *Psychology Today*, May 24, 2012, <https://www.psychologytoday.com/blog/vitality/201205/facebook-and-your-brain>.

⁴⁰ This is explored in greater detail in Maurice E. Stucke & Allen P. Grunes, *Big Data and Competition Policy* (Oxford University Press 2016).

And yet, the optimal level of intervention is difficult to ascertain. With that in mind we briefly sketch here ex-post intervention, merited when an abuse of dominance is present, and ex-ante intervention, which may address more fundamental market failures.

An ex-post approach may lead to intervention when the platform operating the digital butler holds a dominant position and abuses it. To establish dominance, market power will have to be sustained over time. In the context of our discussion, it is possible that network effects and switching costs may contribute to the emergence of such market power. Abuse may be established when the dominant undertaking engages in exclusionary or exploitative conduct. Such strategies have attracted the European Commission's scrutiny in the past in the area of operating systems and search engines. In *Microsoft*⁴¹, the Commission was concerned with the leveraging of market power from the operation systems by the bundling of Windows Media Player⁴²; and, to restrict interoperability with a view to encouraging use of Windows PCs only with Microsoft group servers, thus discouraging investment in non-Microsoft group servers⁴³. In its Google investigation, the Commission raised concerns as to possible search degradation by Google and possible leveraging of market power.⁴⁴

The difficulty in enforcement in such cases may stem from the nature of the service. First one may question the true market power of personal helpers and their ability to behave independently of others. Even if customers are deemed to be locked in, one may have difficulties establishing some forms of abuse. The personalization of the service may make it difficult to ascertain an objective benchmark for comparison. Credible counterfactuals to quality degradation may be difficult to establish. As the primary interaction takes place at the personal-assistant level, the effects may be seen more as personalization (and thus a legitimate part of technological progress) than exclusionary.

An alternative ex-ante approach will seek to ensure having the pre-conditions to promote privacy competition, ensure that the platform's incentives are aligned with users' interests, and prevent some of the market dynamics which could give rise to exclusionary or exploitative effects. Such approach may be implemented through sector investigations, agreed commitments, or other regulatory instruments. For instance, basic measures would ensure that users retain autonomy, are made aware of outside options and can switch with limited or no costs. One could require digital butlers to indicate clearly, either in a pop up

⁴¹ Case COMP/C-3/37.792 – Microsoft; unsuccessfully appealed in Case T-201/04 *Microsoft Corp v Commission* [2007] ECR II-3601.

⁴² Case COMP/C-3/37.792, recital 826-834; Case T-201/04., para 856

⁴³ Case COMP/C-3/37.792, recital 642-646, in particular; Case T-201/04, para 651

⁴⁴ MEMO/16/1484 <http://europa.eu/rapid/press-release_MEMO-16-1484_en.htm>

window or voice warning when their suggestions are ‘sponsored’ or when they offer service through their own platform network while excluding others. Users may be able to opt out of personalized ads or sponsored products.⁴⁵ All these measures, to be effective, require short and clear communications. Knowing and voluntary consent is key as often the consent in today’s click-wrap is little more than a façade. When users have few, if any, viable options, consent is not real but forced. In addition, ‘consent fatigue’ or digital helpers managing consent forms on your behalf, could lead to meaningless agreement and undermine customer empowerment.

To allow switching between butlers, data mobility should be encouraged. With adequate safeguards one should be able to transfer the core parameters, which will enable a new butler to start from a position of personalization. At the providers’ side, mobility would require access to platforms and the provision of interoperability information. Mobility may require the development of basic industry standards for key data points and will need to take into account issues of licensing and IP rights. Their development should nonetheless allow sufficient freedom for developers, to enable disruptive innovation.

Concluding Remarks

A recent trip to a home-improvement store highlighted the rise of smart appliances, which can communicate with digital assistants. For those of us watching *The Jetsons* the prospect of our own automated helper that will serve our needs might seem marvelous.

In industries dominated with data-driven network effects, we could likely receive free butlers. They will excel at mundane tasks, and as AI develops, they will increasingly assist us with our daily tasks. Seeing the salient, day-to-day benefits, we may trust and rely on our butler. Our assistant will no longer be simply making our French press coffee, and turning on the lights in our kids’ rooms. It will be tutoring our children, entertaining our family, telling us happy or sad stories from around the world, ordering our food (and the books that it recommends), and summoning the driverless car to whisk us to jobs, where we, as a result of being freed from the many household chores, can spend more time and effort to meet the rising expenses.

⁴⁵ Transparency is key – Google, in its last update, allowed users to opt out of personalized ads etc. This is a positive move, which ensures user control over his or her data and search environment.

As we welcome the digital assistants to our homes, we may not recognize their toll on our well-being. The next technological frontier of digital butlers may not be all rosy. As the digital butler increasingly controls our mundane household tasks, like regulating room temperature, adjusting our water heater and playing our favorite music, it will be harder to turn off. It will also be tempting to increasingly rely on the butler for other activities, such as the news we receive, the shows we watch, and the things we buy.

Market forces, given the data-driven network effects, have the potential to increase entry barriers, make the strong platforms (and their butlers) even stronger, and weaken many independent personal assistants. Economic and political power would consolidate into fewer hands. Market forces, left unchecked, may yield a handful of devious butlers, even though the technology exists for an independent virtuous assistant. The large platforms could extract even more personal data and command even higher rents to allow others to access us. Not only will our pocketbooks be affected. Our political and social discourse could also be manipulated.

One of the striking issues is how our digital environment, on one level may appear competitive and free, but on another level is carefully designed to suit the platform's needs (and pecuniary interests) rather than our needs and interests. We may feel that we roam the fields of the free market and free ideas, and yet we are increasingly ushered by the super-platform's digitalized hand.

So we need to be aware of the possible conflicts of interest. Such awareness translates to power. Our public awareness, at present, is limited. Too often we accept terms and conditions without questioning them and assume that greater powers – market or state – will ensure our autonomy and welfare. In reality, the innovation from which we benefit comes at a price which is rarely challenged. Indeed, this evolution may go unchallenged under current user behavior and current antitrust policies. The greater algorithm autonomy in a nontransparent, highly personalized interface with customers can stifle competition enforcement. One would hope that the antitrust agencies and courts understand the risks and work to minimize them, among other things, by educating the user about the cost of freeing oneself from the platform's shackles.

As we explore in *Virtual Competition*,⁴⁶ super-platforms and their personal assistants present unique challenges. We cannot assume that market forces, given these network effects, will deliver the virtuous assistant or curb these abuses. We must consider outside options.

⁴⁶ Ariel Ezrachi and Maurice E Stucke, *Virtual Competition - The Promise and Perils of the Algorithm Driven Economy* (Harvard 2016)

We must ask our antitrust enforcers and legislators what steps they are undertaking to minimize the risks and protect our interests and freedom. Until then, rather than ask your personal assistant if she loves or respects you, you might just want to keep her focused on the more mundane aspects of your life (such as turning off the lights).