



Learning from the Dark Web: leveraging conversational agents in the era of hyper-privacy to enhance marketing

Felipe Thomaz¹ · Carolina Salge² · Elena Karahanna³ · John Hulland³

Published online: 11 November 2019
© The Author(s) 2019

Abstract

The Web is a constantly evolving, complex system, with important implications for both marketers and consumers. In this paper, we contend that over the next five to ten years society will see a shift in the nature of the Web, as consumers, firms and regulators become increasingly concerned about privacy. In particular, we predict that, as a result of this privacy-focus, various information sharing and protection practices currently found on the Dark Web will be increasingly adapted in the overall Web, and in the process, firms will lose much of their ability to fuel a modern marketing machinery that relies on abundant, rich, and timely consumer data. In this type of controlled information-sharing environment, we foresee the emersion of two distinct types of consumers: (1) those generally willing to share their information with marketers (Bufs), and (2) those who generally deny access to their personal information (Ghosts). We argue that one way marketers can navigate this new environment is by effectively designing and deploying conversational agents (CAs), often referred to as “chatbots.” In particular, we propose that CAs may be used to understand and engage both types of consumers, while providing personalization, and serving both as a form of differentiation and as an important strategic asset for the firm—one capable of eliciting self-disclosure of otherwise private consumer information.

Keywords Web · Dark Web · Consumer privacy · Marketing strategy · Chatbots · Conversational agents · Personalization · Anthropomorphism

Introduction

Individuals talking about the Web often refer to it as a static entity. In reality, however, the Web is a constantly evolving,

complex system, with important implications for both firms and consumers. In this paper, we begin by reviewing the Web’s evolution over time, noting that these changes are driven by shifts in the relative market power dynamic that plays out between firms and consumers, based primarily on the issue of which party controls or owns information. We build on this review to suggest how the Web may potentially change in the next five to ten years. Our contention is that society will see a shift in the nature of the Web, as various stakeholders become increasingly concerned about privacy issues, away from a largely automatic “opt-in” culture (wherein consumers typically allow firms to collect, use, and share their personal information with other organizations) to one characterized as substantially “opt-out.” This shift will have profound implications for the practice of marketing.

In particular, we predict that various information sharing and protection practices currently found only on the Dark Web will be increasingly adapted in the overall Web, resulting in a hyper-private, more adversarial environment. In the process of this transformation, firms will lose the ability to create in-depth profiles of consumers, leading to eroded customer

Mark Houston served as accepting Editor for this article.

✉ Felipe Thomaz
felipe.thomaz@sbs.ox.ac.uk

Carolina Salge
salgeca@wfu.edu

Elena Karahanna
ekarah@uga.edu

John Hulland
jhulland@uga.edu

¹ Saïd Business School, University of Oxford, Oxford OX1 1HP, UK

² School of Business, Wake Forest University, Winston-Salem, NC 27106, USA

³ Terry College of Business, University of Georgia, Athens, GA 30601, USA

knowledge and the potential end of existing micro-targeting practices. Marketers will need to be nimble in order to survive this coming change. From the consumer perspective, this shift is likely to increase the costs of information search, and require new, more expensive, and more complex technological investments.

In this type of controlled information-sharing environment, we foresee the emersion of two broad types of consumers, with very different digital data footprints: (1) those consumers who are willing to give permission to firms to track, record, use, and share consequential information (e.g., purchase and site visit histories), rendering their digital essence “naked” to all, and (2) those who deny access to such information and thereby become digital “ghosts.” The first group of consumers (Bufs) will be similar to today’s digital consumers, whereas the second group (Ghosts) will be quite different.¹

Operating in such an environment will create problems for businesses familiar with current Web 2.0-based tools, designed to optimize marketing to the first group (Bufs). So how should firms respond? We argue that conversational agents (CAs)—often referred to as “chatbots”—will play an increasingly important role in helping firms market to both groups. We are already seeing signs of firms using artificial intelligence (AI) and machine learning techniques to provide value by replacing humans in social systems where the occupation or position is either “undesirable” or costly, and by augmenting and complementing humans in social systems where the task or job is either tedious or repetitive. For example, a growing number of chatbot service agents are replacing call center employees, but also freeing human claims processors to focus on resolving more complex cases (Juniper Research, “AI in Retail,” April 2019). Whereas all CAs will be “intelligent,” our prediction is that the underlying machine learning mechanisms required to engage the two consumer groups will be different, dictated by the data that they are willing to reveal. CAs interacting with Buff consumers will have supervised learning enabled, and greater personalization will be possible. In contrast, Ghost consumers will receive mass-personalized CA assistance, where aggregated data enabled by unsupervised learning algorithms will provide lower value for both the consumers and the firm. Furthermore, CAs may help firms elicit additional consumer information by nudging Ghost consumers to increase self-disclosure via the promise of more personalized and valuable CA interactions and

through evoking social responses through anthropomorphic design.²

Evolution of the Web

Each step in the Web’s evolution (summarized in Table 1) has been accompanied by a significant initial shift in the balance of power between firms and consumers, and a similar level of concern by managers seeking to leverage the available technology to maximize firm value. These initial threats to firm success are identified in the third column of Table 1. For example, in the early days of Web 1.0, there was a fear that by disseminating information widely (particularly price information), firms would lose their ability to effectively price discriminate across markets, segments, and purchase occasions, essentially collapsing markets to a lowest common denominator (e.g., Burke 1997). There was also the fear of engaging in competitive warfare, forcing firms to quickly “race to the bottom,” given the perception that consumers would have full information about competitive offerings (e.g., Peterson, Balasubramanian, & Bronnenberg 1997).

In a similar vein, the shift to Web 2.0 resulted in an increased ability for consumers to coordinate their activities, exchange information directly with one another, and further organize themselves into dedicated and vocal specific-interest communities. Version 2.0, in conjunction with the rise of social media, marked a substantial reduction in managerial control over the firm’s own messaging, as users created and disseminated their own content, with value-relevance to the firm (e.g., Edvardsson, Tronvoll & Gruber 2011).

However, in each of these first two evolutionary stages, over time managers were able to identify important firm advantages (listed in column four of Table 1) that resulted in a releveling of the balance of power between firms and consumers, in the form of increasing levels of data generated by consumers in these digital environments. Data-rich environments provide firms with greater amounts of consumer knowledge, and allow for less intrusive observation of actual behaviors and preferences. The resulting Digital, Social Media, and Mobile (DSMM) ecosystem has been considered as a source of intelligence (Lamberton & Stephen 2016) for observing, analyzing, and predicting behavior (Bucklin & Sismeiro 2003; Chatterjee, Hoffman, & Novak 2003; Montgomery et al. 2004). The advancements in this area have led to, among other developments, efficient behaviorally targeted ads (Lambrecht & Tucker 2013; Summers, Smith & Reczek 2016), intelligent product recommendation systems (Ghose

¹ This segmentation into only two types of consumers is clearly somewhat simplistic. However, focusing on Bufs and Ghosts throughout the balance of this paper allows us to draw important distinctions between the two groups, and to discuss how marketers will need to adopt very different strategies for identifying and engaging each. In reality, of course, many consumers are likely to fall somewhere in between the two extremes that we describe in detail here.

² In this paper, we use the term “nudge” to describe a process whereby bots have a perceived value that incentivizes consumers to share their data, rather than bots designed to extract data unwillingly. As we subsequently discuss, our perspective here is that marketers should not design CAs to “deliberately mislead users as to privacy features” (Leong & Selinger, 2019, p. 300).

Table 1 Evolution of the Web, by stage

Web Stage	Primary Focus	Threats to Firms	Advantages to Firms	Representative Description
Web 1.0	Focus on content and information presentation; “Publication logic”	Consumers have full access to information and can look for best deal	Gain access to rich demographic, transactional, search data	“Internet-related marketing can result in extreme price competition when products or services are incapable of significant differentiation. This can happen when they are perceived as commodities, partially because other factors that might moderate competition (e.g., store location) are absent and partially because of the relative efficiency of price searching engendered by the Internet.” Peterson, Balasubramanian & Bronnenberg (1997, p. 336)
Web 2.0	Focus on multiparty communication and information exchange; “Participation logic”	Consumers share content and opinions; firms have reduced control over messaging and branding	Gain access to relational (network) data, can undertake social listening, ability to foster eWom	“In the current Web 2.0 era, customers are not only passively using, but also actively creating and sharing, web content, and they thus not only co-create but also co-produce value and shape service as well as social systems. Organizations will have to adapt to—and ideally pro-actively influence—this new social reality if they want to be able to continue to understand and manage service exchanges and co-creation of value with their customers in the future.” Edvardsson, Tronvoll & Gruber (2011)
Web 3.0	Focus on privacy–personalization balance; “Privacy logic”	Consumers empowered to assert right to privacy; firms must find ways to recognize, understand, and engage consumers	AI applications can provide customer value and incentivize data disclosure to trusted platforms	“Evidence available from the recent past [suggests] that a market in privacy will emerge, with anonymity available to those who value and can afford it...” (Deighton 1997, p 250)

et al. 2012), and morphing banner advertising (Urban et al., 2013), which are made possible by leveraging user profiling techniques (Trusov, Ma, & Jamal 2016).

However, in the current evolution to Web 3.0, marketers face the very real possibility of losing these advantages if they no longer have the ability to identify consumers and connect them to their previous behaviors (e.g., Deighton 1997). Customer privacy is central to this potential problem (Martin & Murphy 2017, Stewart 2017).

Over the past two decades, accompanying the growth of the market on personal identity and behavior data described above, consumers, firms, and governments have increasingly engaged in discussions of the nature of ownership of this data, as well as potential mechanisms and controls to protect the different needs of these parties. For example, consumers increasingly use Adblock-like technologies to avoid advertising, cookies, and trackers online. Firms use this concern as additional points of differentiation (e.g., Apple refusing a government request to unlock a customer’s iPhone), and governments look for new ways to restrict access to customer data (e.g., the European Union’s General Data Protection Act (GDPR), California’s AB 375 bill).

Altogether, this shifting to a private-as-default behavior marks the emergence of a new Web 3.0 environment, wherein novel technologies and evolving consumer behaviors combine to create a new set of challenges and opportunities for marketers. Fortunately, there is a part of the Internet that we can examine to understand and learn more about this hyper-private future: the Dark Web.

What marketers can learn from the Dark Web

The Web is divided into three sub-components: the Surface Web (e.g., Clearnet), the Deep Web, and the Dark Web. The Surface Web is the component most people are immediately familiar with, as it incorporates all of the websites indexed by search engines, and represents all of the sites that a person can reasonably and easily navigate to. The Deep Web contains information that lies behind some sort of barrier (typically in the form of passwords) that inhibits easy, unapproved access. For example, individuals’ private bank account information resides in the Deep Web, secure behind a barrier, well away from any random access requests. Finally, the Dark Web (somewhat of a misnomer, in that it is not very web-like) is made up of non-indexed and disconnected websites that require specialized software (e.g., The Onion Router; usually referred to as TOR), as well as specific knowledge and authorization (i.e., a given URL or onion address) to gain access.

The Dark Web has gained recent popularity (and notoriety) in the press because of revelations about hidden criminal activity and black markets (such as the original *Silk Road* marketplace), whistleblowing websites (like *Wikileaks*), and

activist safe-havens (e.g., *Arab Spring*). However, the environment is not exactly new. For example, the TOR system has been around since 2002 (eleven years before the launch of *Silk Road*), and similar connecting environments even longer (e.g., *Napster* (started in 1998), *USENET* (1979)).

Research examining darknets and the Dark Web

Over the past quarter century, marketers have studied the impact of darknets—restricted and parallel or isolated network (e.g., file transfer P2P, like *Napster* or *Bittorrent*)—on a number of categories and industries (e.g., music, film, software). This research has predominantly focused on the effects of P2P file sharing (and other piracy behaviors) on price sensitivity (Jain 2008, Sinha & Mandel 2008), substitution effects (Danaher et al. 2010), diffusion (Givon, Mahajan & Muller 1995), and control mechanisms/policy (Sinha, Machado, & Sellman 2010). Similarly, other fields have explored traditional marketing questions, like the profitability of vendors and residual value to consumers (Holt, Smirnova, & Chua 2016), while utilizing traditional marketing approaches to guide research in this environment (Li, Chen, & Nunamaker 2016; Benjamin, Valacich, & Chen 2019). (See Table 2 for a brief summary of these papers). What is important to recognize is that these early (and often illegal) consumer practices have evolved over time from darknets to the Clearnet, and from services like *Napster* and *Pirate Bay* into *iTunes*, *Netflix*, and others.

In contrast, the Dark Web, as a massive World Wide Web-scaled “darknet” in terms of size and scope of activity, has not

been featured in the marketing literature. In the Dark Web, one can observe privacy-seeking customer behaviors, as well as the genesis of encrypted and private marketplaces that contain mundane items (e.g., books, services), in line with the existing “darknet” research. However, the Dark Web also entails a much broader set of both licit (e.g., Facebook, *The New York Times*) and illicit (e.g., trade in drugs, weapons, and human trafficking) commercial activities, as well as consumer-consumer, consumer-firm, and consumer-technology interactions that extend beyond individual preferences or singular marketplaces and platforms. Researchers outside of marketing (e.g., in criminology, information systems, and public policy) have started to explore a variety of topics relating to the Dark Web. Table 3 provides a summary of some of these papers; they are included here because they touch on topics that are familiar to or that might normally be considered the domain of marketing scholars.

For instance, within criminology, researchers have examined the effect of specific news on illicit sales volumes, finding counterintuitive increases in Dark Web transactions (Ladegaard 2019). Other criminology scholars have examined the structure of illicit digital markets, and the associated efficiency and resiliency they exhibit (Bakken, Moeller & Sandberg 2018; Duxbury & Haynie 2019). In a related IS study, Yue, Wang, & Hui (2019) conducted a user-generated content analysis of Dark Web hacker communities, and found evidence connecting increases in user chatter to a lower frequency of cyber-attacks. Closer to home (for marketing researchers), policy scholars have focused on aspects of consumer well-being in terms of satisfaction and safety (Barratt,

Table 2 Research examining darknets with a marketing focus

Paper	DV/Focus	Finding
Givon, Mahajan & Muller (1995)	Sales and Product Diffusion	Models the diffusion of a product along an official channel and a shadow (illicit, darknet) path, showing that piracy can positively impact the official diffusion through word of mouth
Jain (2008)	Price Competition	Piracy can enable lower price competition as price sensitive consumers move to the illicit (darknet) channel
Sinha & Mandel (2008)	Willingness to pay/Willingness to pirate	Improved functionality in legal channels significantly reduce willingness to pirate, while either increased likelihood of getting caught or severity of punishment can actually increase piracy in more risk tolerant consumers
Sinha, Machado, & Sellman (2010)	Digital Rights Management (DRM) as piracy deterrent	Removal of DRM has the potential to convert darknet consumers into paying consumers, while increasing the demand and willingness to pay for the product
Danaher et al. (2010)	Sales and Cannibalization	The temporary removal of a legal channel resulted in a disproportionate demand for pirated (darknet) content, in evidence of larger relative fixed/learning costs of piracy, but low marginal cost of additional consumption.
Holt, Smirnova, & Chua (2016)	Revenue of Vendors/ Value to Buyers	Data vendors profit from all transactions, whether the items being sold have value or not, earning and estimated \$326,000 over the observation period. Buyers bear more risk, but profit from exploitation of data, earning as much as \$2 Million in return.
Li, Chen, & Nunamaker (2016)	Identification of key vendors	Utilization of social media analytics, sentiment scoring, and topic modeling to identify key vendors in data markets
Benjamin, Valacich, & Chen (2019)	Darknet Research Framework	Development of a framework for research on darknets, including the identification of data sources, data collection strategies, data evaluation, and ethical concerns

Table 3 Recent (non-marketing) Dark Web research

Paper	DV/Focus	Finding
Criminology		
Bakken, Moeller & Sandberg (2018)	Market Efficiency	Dark Web black markets are formally organized, with hierarchical structures, rules, and limitations. These markets reduce the costs and risks associated with traditional drug marketplaces and allows for more efficient transactions
Duxbury & Haynie (2019)	Criminal Network Resilience	Brokers (diverse connectivity) within a criminal network represent better targets for removal in law enforcement intervention when compared to high connectivity individuals. Recovery time is conditional on market strategies (secrecy/efficiency).
Ladegaard (2019)	Illicit trade volume	Media coverage of law enforcement activity on dark web marketplaces does not hinder trade, but rather increases the volume of transaction, including immediately following arrests and sentencing
Information Systems		
Yue, Wang, & Hui (2019)	Cybersecurity	Increasing the number of discussions in Dark Web hacker forums about denial of service (DDOS) attacks actually reduces the number of successful attacks executed
Public Policy		
Barratt, Ferris, & Winstock (2016)	Consumer Safety	Dark Web marketplaces are associated with significantly fewer threats and violence than alternative drug markets, even compared to relatively safe private/closed markets
Caudevilla et al. (2016)	Drug Purity	Hundreds of samples of drugs from Dark Web were tested in a laboratory, finding that the samples contained the advertised substances, and the majority was of high purity.
Van Buskirk et al. (2016)	Consumer Characteristics	Dark web drug customers tended to be younger, involved in property crime, and engaging with a growing set of drugs

Ferris, & Winstock 2016; Caudevilla et al. 2016; Van Buskirk et al. 2016). Despite these early efforts, however, little is known about consumer Dark Web behavior, and virtually no marketing implications have been explored.

The Dark Web as an unregulated, adversarial testbed

The Surface Internet has always lagged somewhat behind the Dark Web in terms of the available technology it utilizes. In contrast, the Dark Web functions as an unregulated testbed for new ideas and technologies, with its successes often later migrating to the Surface Web. In the Dark Web, unpolished user interfaces, unstable services, and higher levels of user involvement appear to be fairly standard.

Despite the illegal (and immoral) activities often associated with the Dark Web, it continues to embrace a libertarian, hacker ethos, especially in terms of its respect for experimentation and associated freedoms. In return, the beta-like environment of the Dark Web is often forgiving of errors; developments do not have to satisfy governmental (or other institutional) standards, and they are not beholden to other stakeholder concerns. As a result, trial-and-error is more rampant in the Dark Web. For example, whereas the Surface Web is currently grappling with the high volatility in the value of Bitcoin (the most familiar example of cryptocurrencies), the Dark Web trades in a variety of different cryptocurrencies, some of which are market-specific. Similarly, whereas the Surface Web is concerned with identifiable data leakage, misuse, and manipulation, the Dark Web operates with encrypted

communications (PGP) and distributed signals to aid in anonymity (TOR).

Over time, we can expect that Surface Web consumers as well as digital/digitized consumption environments will become more similar to what is currently observed in the Dark Web, with commensurate impact on marketers and firms. What's the basis for this claim? We already see some adoption of Dark Web technology and behaviors on the Surface. For example, in 2019 the Firefox web browser adopted an anti-fingerprinting measure to increase user privacy and circumvent advertising relying on tracking. This technology was originally developed for TOR, the Dark Web browser. Along similar lines, Google announced in 2019 anti-fingerprinting actions on their Chrome browser, as well as an entire set of open industry standards to safeguard user privacy for the entire web, dubbed the 'Privacy Sandbox.' Lastly, as perhaps the best-known example, the *WhatsApp* application allows for end-to-end encrypted communication for individuals and communities, using established cell and data networks to ensure user privacy. Thus, *WhatsApp* usage creates a Dark Web-like experience, wherein consumers employ hyper-private communications of content hidden even from the service provider. As a result, it is impossible for the provider to pass on information to third parties; the type of information that has fueled many firms' marketing successes in the Web 2.0 environment. Nor is *WhatsApp* alone in this space; encrypted messengers include *Telegram* and *Signal*, among others. New encrypted browsers that route all traffic through encrypted virtual private networks to mask user identity and location are also being introduced (e.g., *Epic*).

The primary aim of Dark Web innovation is to maintain complete user privacy (e.g., total anonymity). Users can utilize “anonymity-granting technologies” to protect their privacy from government agencies, political opponents, trolls, data-hungry organizations, and even Internet service providers (Jardine, 2018). In this adversarial environment, individuals view every other entity as a potential “enemy,” eager to acquire useful information and ready to deploy it against them. As a result, Dark Web participants use every possible measure, from technological to behavioral, to minimize (or eliminate) their digital footprints. The end-result of this is that very little information is visible about any individuals operating in the Dark Web, unless they choose to disclose it. (Due to the high costs of exposure, particularly in illegal markets, such disclosures are quite rare.)

This focus on privacy enables Dark Web users to personally control (and thus limit) access to information about themselves (Altman 1975; Westin, 1967). This view of privacy as selective control represents a common perspective on privacy that originated in Westin’s (1967) and Altman’s (1975) theories of general privacy. For example, Altman defined privacy as “the selective control of access to the self.” The control-based definition of privacy is broadly accepted and has been used as the foundation of most information privacy research (Smith et al. 2011). But the ability to completely limit access to the self in order to protect the self comes at a relatively high cost in terms of having to accept lower computer performance, slower internet browsing, and greater inconveniences. In fact, the desire to increase transaction efficiency while remaining anonymous drives much of the Dark Web innovation we described earlier.

The role of privacy in the emerging dark surface

The privacy-focused behaviors described above and enacted by Dark Web users represent a significant threat to marketing practices in the current Surface Web, which rely on easy capture of digitally-based consumer information from data rich environments that facilitate precise targeting, re-targeting (abandoned baskets), behavioral advertising, lookalike modeling, etc. Furthermore, many firms currently benefit from sharing or selling this information to third parties. However, in a system where consumers have control over their information and act in ways to look like unknown new visitors (i.e., no part of their digital character is revealed, but instead protected), the value of firms’ existing data and models will be vastly diminished. For example, lacking the ability to connect users to previously collected information (e.g., click-stream data, which is fairly common today), firms will be forced to resort to “average consumer” profiles to predict consumer behavior (only updating when consumers are willing to disclose

specific information about themselves). This data-impooverished environment will result in firms adopting more traditional mass-market approaches (with attendant lower profit potential, loss of efficiency, and eroded effectiveness).

If the Dark Web is indeed the unstable precursor of the future Surface Web, we can expect the Surface to go “dark,” and that browsing will become incredibly private in a user-friendly way with minimal to no-cost. The incentive for greater privacy, allowing consumers to secure control over and limiting access to their own personal information in all levels of the Web, will come from a combination of: (1) a continued increase in the value of an individual consumer’s “Identity Graph” (the aggregated total digital and analog data footprint of an individual), (2) improved software, and (3) growing government and quasi-government concerns with privacy violations of individuals.

With respect to the last of these three drivers, there is increasing momentum towards the view of privacy as a fundamental “human right,” and is recognized as such under Article 12 of the 1948 UN Universal Declaration of Human Rights³ as well as by the constitutions of many countries. In the past, a distinction has been made (Smith 2001, pp. 1000-1001) between countries that viewed privacy as a human right and passed “sweeping privacy bills that address all the instances of data collection, use, and sharing (Bennett & Raab 2006; Dholakia & Zwick 2001)” versus those that viewed privacy as a commodity, and enacted a “patchwork of sector-specific privacy laws that apply to certain forms of data or specific industry sectors (Bennett & Raab 2006; Dholakia & Zwick 2001).” It follows that countries seeing privacy as a right generally adopt practices where the default is privacy, whereas the privacy-as-commodity countries instead consider the process of requiring opt-ins to disclosure to be an “undue burden.” However, the momentum towards a view of privacy as a right does shift the market towards an opt-in environment, where companies only have access to data about consumers who choose to make that specific information available (Smith 2001). Privacy that was once described as “the right to be let alone” (Warren & Brandeis, 1890) will be best described in a few years as “the right and ability to control information about the self.”

So, what is inhibiting this shift to a Dark Surface Web? First, Dark Web privacy currently comes with a significant performance loss (e.g., slow page loading speeds within TOR), and safe encryption requires managing long public keys. However, technological advances deployed in the Dark Web are increasingly making privacy protection technologically feasible at scale, and financially viable. Easy, single-click privacy and encryption software allowing consumers to minimize (or eliminate) their digital footprints will facilitate this shift in the longer term by minimizing the effort required

³ https://en.wikipedia.org/wiki/Universal_Declaration_of_Human_Rights

to assure privacy. (For example, consider Google’s activities described earlier, and note that the fingerprinting protection enables greater user privacy at no cost to the user, and with a potential advertising revenue loss to Google).

Second, many current consumers are not fully aware of the value of their personal data, and tend to make it available to firms at little or no cost. Towards this end, states are passing privacy acts that protect consumers (e.g., see California’s AB 375 bill), raising awareness about the costs and risks of freely sharing information with firms and thus exacerbating the concerns consumers already have. However, even those consumers who are fully aware of the costs, and are concerned about their privacy, often choose to disclose identifiable personal data (Adjerid et al., 2018). This discrepancy between attitudes and behaviors is what scholars refer to as the “privacy paradox” (Acquisti et al., 2015). Thus, in addition to advances that make safeguarding privacy viable at the individual level, consumer-based behavioral changes on the Surface Web will be necessary in order to ensure that consumers do not “give away” information, rendering privacy software irrelevant.

In this new, Dark Surface Web environment, we believe that the default consumer behavior with respect to granting data usage permission to firms will be minimal, since customers are becoming more reluctant to opt-in and less predisposed to share information unless given strong incentives to do so (e.g., the “privacy calculus” phenomenon (Dinev & Hart, 2006). In this emerging dark surface environment, the standard consumer will be a Ghost, with the firm having very little insight into the nature of the individual. In order to overcome this, firms will need to provide some value, or incentive, for users to engage in a mental privacy calculus that may lead them to opt-in and provide the firm with additional, consumer-specific information. However, at the other end of the spectrum, it seems likely that a second general group of consumers (Buffs)—those who are readily willing to share their personal information with firms—may also exist.

How can firms, accustomed to having access to digital footprints of customers and other profile information to personalize offerings and interactions, operate in a hyper-private opt-in rather than a naked-to-all opt-out world? While Buff consumers will share personal information enabling firms to continue using their extant methods, firms will have little information on the digital footprint and preferences of Ghost consumers. One way to entice Ghost consumers to disclose personal information is to provide them with financial incentives. This strategy would result in an information market where firms could purchase personal information from consumers willing to sell it, but also where consumers could purchase this information ‘back’, or even sell it to other consumers. Studies on privacy calculus already show preliminary evidence of this dynamic, whereby consumers weigh privacy concerns and related risks against the benefits of information

disclosure, and sometimes end up trading privacy for monetary rewards (e.g., Caudill & Murphy 2000; Hann et al. 2008; Phelps et al. 2000; Xu et al. 2010).

Another possibility is to utilize technology to nudge consumers toward self-disclosure in exchange for hyper-personalization. Hyper-personalization is a significant benefit, separate from those provided by financial rewards, and consumers are frequently willing to exchange their own privacy for personalized offerings. This complex trade-off between personalization and privacy is known as the “personalization paradox” (Aguirre et al. 2016; Bleier & Eisenbeiss 2015). Technology can play a prominent role in this trade-off. For example, anthropomorphized technology has the potential to nudge consumers towards greater self-disclosure by transmitting social cues that activate social scripts and through conversations that invoke norms of reciprocity (Moon 2000). Such anthropomorphization can evoke social responses that encourage greater self-disclosure, even by Ghost consumers.

Ultimately, firms will need to create strategies to personalize interactions and provide value to both Buffs and Ghosts. Though there has been extensive research on personalization, the emerging Dark Surface environment that makes consumer profiling less accessible creates new challenges. Firms adapting to this new environment will need to understand the ways in which they are affected by the “personalization paradox” and which consumer-facing technologies will generate the greatest value for consumers in a way to tip the trade-off towards data sharing. Among the set of candidate technologies that can provide a lever in this trade-off, the increasing shift towards conversational-commerce (a term coined by Uber’s Chris Messina) provides one of the most compelling candidates. Conversational commerce refers to the use of natural language interfaces (such as chats and messaging) by consumers to interact real-time with organizations (humans and bots). Gartner predicts that by 2020, 85% of all consumer interactions with a firm will occur via conversational agents (CAs). As such, given the expected ubiquity of the technology in consumer interactions, we see CAs as one critical technological facilitator of firm-consumer exchange in the emerging Dark Surface Web 3.0 environment. In the balance of this paper, we explore in greater detail the role that can be played by CAs to nudge consumers towards greater self-disclosure through anthropomorphization, and to provide varying personalization value to each of the two consumer groups described above.

Conversational agents

Conversational agents (also called chatbots, conversational AI-bots, virtual assistants, and dialogue systems) are natural language computer programs designed to approximate human speech (written or oral) and interact with people via a digital interface. Although they have existed since the 1960s (e.g.,

ELIZA developed by Joseph Weizenbaum in 1966), conversational agents (CAs) have recently garnered substantial industry attention. They are becoming the new front-office face of many companies, representing a shift from “clicks to conversations” (Daugherty & Wilson, 2018) and from e-commerce to conversational-commerce. CAs are also becoming critical components of the customer service infrastructure, by replacing or augmenting tasks traditionally performed by sales employees (Larivière et al., 2017; Verhagen, van Nes, Feldberg, & van Dolen, 2014) and by providing consumers with successful service encounters (Larivière et al., 2017; van Doorn et al., 2017). The recent availability of conversations-as-a-platform (CAAS) tools is making it easier for firms to develop and deploy such CAs.

Examples of CAs abound, and range from *Alexa*, which allows people to execute a variety of mundane tasks such as ordering food and tracking flight statuses, to *Dressipi's Amiya*, which helps customers find and purchase products they want based on style preferences. CAs can be entirely digital and exist online (e.g., Bank of America's *Erica*), or can have physical embodiments and exist offline in organizational settings, stores (e.g., *LoweBot*), or one's home (e.g., *Alexa*). Given the range of CAs, one way in which they have been classified is based on whether they are (a) general-purpose CAs, such as *Siri* and *Alexa*, or domain-specific CAs, such as *IKEA's Anna*, and (b) whether their primary mode of communication is text-based or speech-based (Gnewuch, Morana, & Maedche, 2017).⁴

The major aim of CAs is to enhance both the experience and the outcomes of consumer interactions with the organization across sales, marketing, and customer service (Daugherty & Wilson, 2018). For example, *Hello Hipmunk* is a CA that makes it easier and more convenient for people to search and book vacation trips. As Adam Goldstein, the CEO and co-founder of Hipmunk, noted:

The average traveller runs 20 searches when planning a trip. Hello Hipmunk shrinks that process to one simple conversation. It can process tons of information from flight pricing to room availability and synthesize it instantly (Staff, 2016, para. 3).

Conversational agents: Competitive assets in an increasingly dark surface web

CAs that interact with people as useful private assistants or effective customer service representatives are likely to be

⁴ Many other classifications also exist. For example, Gartner describes CAs based on (a) engagement levels, ranging from the provision of explicit information at one end to interactive conversation at the other end; and (b) task complexity, ranging from informational at one end to transactional at the other end.

major assets for companies. For example, Juniper Research predicts that by 2022 CA use for customer service will save companies \$8 billion. But beyond being just another customer interaction tool, CAs can also become a way firms differentiate themselves.

Because they make it more convenient for people to rapidly access data, evaluate information, and execute tasks (Sankar & Balakrishnan, 2016; Shum, He, & Li, 2018), in addition to providing more enjoyable experiences (Brandtzaeg & Følstad, 2018) and a sense of companionship (Turkle 2017; Brandtzaeg & Følstad, 2017), CAs can nudge consumers to voluntarily share personal data with companies. This sharing of data is clearly important for online interactions, where CAs can prod people to disclose identifiable, rather than de-identified, data (see section below on “Personalizing Interactions for Buff and Ghost Consumers” where we elaborate further on this). Furthermore, CAs provide a means for companies to collect offline consumer data, a task that has traditionally proven more challenging. For example, they can track what clothing items people bring into fitting rooms and answer questions related to size availability, color options, matching accessories, etc. (Daugherty & Wilson, 2018). In the process, they not only collect information on popular items in general, but can also incentivize consumers to share identifiable personal data in order to receive more personalized recommendations.

By directly collecting private data from consumers (both online and offline), CAs can enable the generation of more accurate identity graphs that allow companies to market products and services more efficiently and effectively by, for example, targeting people with the right content at the right time (McAfee & Brynjolfsson, 2012; Schumann, von Wangenheim, & Groene, 2014; Spangler, Hartzel, & Gal-Or, 2006). Collected personal data can also be leveraged by CAs in future interactions to further personalize, at great scale, conversations with people. Mattel's *Hello Barbie*, the world's first Barbie CA, represents an example. Not only can *Hello Barbie* engage in meaningful conversations with children, but she can also capitalize on the details of prior interactions, such as a child's favorite color and beloved pet, to quickly become a close friend (Vlahos, 2018).

Moreover, CAs can be a source of differentiation and competitive advantage when they become orchestrators of customer interactions, not just within the company, but across other companies as well—for example, people can use Amazon's *Alexa* to both order pizza from Domino's and get flight status updates from Delta (Daugherty & Wilson, 2018). As Daugherty and Wilson (2018, p. 95) point out: “In the past, companies like Domino's, Capital One, and Delta owned the entire customer experience, but now, with *Alexa*, Amazon owns part of the information exchange as well as the fundamental interface between the companies and the customer, and it can use the data to improve its own services.” Consequently,

companies owning the most popular CA interfaces will be advantaged.

The ability to use CAs to collect identifiable data from people, both online and offline, and both within and across firms, is going to become especially important in a world of web platforms where browsing activity is increasingly private (Bursztein, 2017). As a result, companies that motivate and nudge consumers to self-disclose private data, in interactions with CAs or otherwise, will have an edge.

Furthermore, firms will have to rise to the challenge of meaningfully personalizing consumer interactions in such an information impoverished environment. Personalization has been identified as one of the most successful relationship-building mechanisms used by firms (Claycomb & Martin, 2001), since it increases sales' leads, customer acquisition and retention (Bojei et al., 2013; Sahni, Wheeler, & Chintagunta, 2018), firm profit, customer satisfaction, and enables the discovery of novel consumer needs and preferences (Arora et al., 2008; Huang & Rust, 2017). While some CAs may provide an impersonal experience, the more successful CAs will be designed to engage and personalize the experience even for Ghost consumers.

Our discussion of CAs in the next sections focuses on these two issues: ethically nudging consumers towards voluntary self-disclosure of personal data, and designing CAs to engage consumers through personalized interactions.

Encouraging voluntary self-disclosure with conversational agents

Developing CAs to nudge people, especially Ghost consumers, to self-disclose private data requires research to inform which design features and personality traits may result in the creation of engaging, trustworthy, and ethical CAs.

Ethical anthropomorphism One way to foster trust, increase engagement, and encourage self-disclosure is to ethically anthropomorphize⁵ CAs, as individuals often feel less inhibited when interacting with anthropomorphic computers, sharing private information (Leong & Selinger, 2019; Turkle, 2017), and even developing personal relationships (Moon, 2000) (see Appendix 1, Table 7 for a review of prior studies). The process of anthropomorphising CAs can occur in a variety of ways, but some dimensions to consider include name, gender, embodiment, a physical (or virtual) appearance that may include age, ethnicity, and attractiveness, a personality, a voice with a certain tone or expression (if speech-based), and a conversational style that can range from open-ended to predefined

⁵ Honest or ethical anthropomorphism is the idea that “robot designers should not use anthropomorphism to deliberately mislead users as to privacy features” (Leong & Selinger, 2019, p. 300). For simplicity, from now on we use the words anthropomorphism, anthropomorphization, or anthropomorphic features to refer to ethical anthropomorphism.

answers (see Leong & Selinger, 2019). Anthropomorphization is especially effective when the anthropomorphic features of the CA (such as ethnicity (Qiu & Benbasat, 2009) or personality (Al-Natour, Benbasat, & Cenfetelli, 2006, 2011)) are designed to be similar to those of the consumer, a phenomenon we term *homophilous anthropomorphism*.

Anthropomorphization influences behaviors through a number of mediating mechanisms. First, anthropomorphised agents provide nonverbal cues that often generate “mindless” responses from people, to the extent that people apply *social scripts*—scripts for human-to-human interaction—to CAs, “essentially ignoring the cues that reveal the essential asocial nature of a computer” (Nass & Moon, 2000). These social responses occur as a result of conscious attention to a subset of contextual cues that trigger various scripts from the past (Langer, 1992; Moon, 2000, 2003; Nass & Moon, 2000; Nass, Steuer, & Tauber, 1994) that are applied mindlessly even when such behaviors seem irrational, inappropriate, or unnecessary (Nass & Moon, 2000).

Second, in addition to evoking social scripts that encourage social interaction, social responses to anthropomorphised CAs can nudge consumers towards self-disclosure through evoking *norms of reciprocity*. For example, Moon (2000, p. 328) shows that people share intimate data with computers “when computers initiate the disclosure process by sharing information first” and then follow a “socially appropriate sequence of disclosure by escalating gradually from superficial to intimate disclosures.” Thus, to nudge consumers to share information, anthropomorphic CAs can also incorporate design elements of reciprocity without violating patterns of escalation in disclosure (e.g., lie, share too much information too fast, or ask people to disclose data too early).

Third, anthropomorphism also increases *social presence*, defined as the degree to which a communication medium allows one to perceive the communicator as being psychologically present during an interaction (Short, Williams, & Christie, 1976). Social presence resulting from anthropomorphization has been associated with trust, engagement, and satisfaction (Kumar & Benbasat, 2006; Picard, 1997; Qiu & Benbasat, 2009; Turkle, 2017; Bleier, Harmeling, & Palmatier 2019). In examining specific anthropomorphic features, researchers have shown that certain personality characteristics such as friendliness and expertise (Verhagen et al., 2014) as well as embodiment and communication style (Qiu & Benbasat, 2009) influence perceptions of social presence. For example, recommendation agents with animated faces (rather than disembodied ones) and voice outputs providing rich social cues (rather than text) enhance social presence and generate higher trust, enjoyment, and perceived benefits (Qui and Benbasat 2009).

Since a CA's anthropomorphic features can affect both people's interactions and engagement with the CA, as well as their

perceptions of its trustworthiness and usefulness, they are also likely to influence their *privacy-calculus assessments*. Consumers weigh privacy concerns and related risks against the benefits of information disclosure (e.g., Dinev & Hart, 2006). The extent to which CAs are perceived as more engaging, enjoyable and useful will magnify their perceived benefits while the extent to which they are perceived as more trustworthy will reduce the perceived privacy risks. Both effects will shift the privacy-calculus towards greater information disclosure. Furthermore, emotions impact the privacy calculus of consumers, with positive affect leading to lowered perceptions of risk (Li, Sarathy, & Xu, 2011), higher intentions to disclose information (Anderson & Agarwal, 2011), and more self-disclosure (Kehr, Kowatsch, Wentzel, & Fleisch, 2015; Yu, Hu, & Cheng, 2015).

However, there is a non-linear relationship between anthropomorphization and outcomes. While anthropomorphization can generate positive marketing results (Aggarwal & McGill, 2007), too much of it can lead to negative effects. For example, some attempts to provide overly humanized agents have created unrealistic consumer expectations that turned into frustration (Knijnenburg & Willemsen, 2016) and abuse (Neff & Nagy, 2016). Excessive anthropomorphism can also trigger consumer discomfort (also known as the “uncanny valley” concept—see Mori, MacDorman, & Kageki, 2012; van Doorn et al., 2017) leading to decreased favorability toward the CA (Mende et al., 2019).

Additional research is thus required to better understand the right level and type of anthropomorphic design features for each context and different consumer groups (i.e., for Buff consumers vs. Ghost consumers). For example, while homophilous anthropomorphism across a number of features (age, gender, personality, race, dialect, etc.) may be possible with Buffs (since identifiable data and personal characteristics are captured and used) there are limits to the level of homophilous anthropomorphism possible with Ghosts (given that only de-identifiable data at the aggregate level is used). More work is also needed to investigate which anthropomorphic cues are consequential to organizational outcomes and under what conditions and which of these may encourage Ghost consumers towards greater information disclosure.

Fairness and transparency Alternatively, self-disclosure can be encouraged when firms provide assurances of algorithmic fairness and transparency (Garfinkel et al., 2017). In terms of fairness, audits and certifications provide assurances that the firm’s CAs are fair and unlikely to generate biased interactions towards particular subgroups, such as customers of certain race, gender, or socioeconomic status. This concern is likely to be more prevalent for the Buffs who disclose such information. In terms of transparency, developing CAs with predictive models that provide explainability, such as logistic regression (rather than “black box” predictive options like neural

networks) or using models such as LIME (Local Interpretable Model-Agnostic Explanations) that generate explanations to describe predictions made by machine learning algorithms, provide assurances that the firm is committed to accountability and transparency, and willing to share how the information delivered by CAs have been derived. Empirical evidence (e.g., Wang & Benbasat 2008) suggests that such transparency, i.e., explaining the logic behind a particular thought, decision, or recommendation, engenders trust.

Table 4 presents some research questions on how to encourage voluntary self-disclosure through conversational agents. Of great importance is identifying which of these are most effective for nudging Ghost consumers towards disclosing more personal information.

Personalizing interactions for Buff and Ghost consumers

Given that designing CAs to incentivize self-disclosure will nudge some consumers but not others, firms need to design CAs that personalize interactions for both types of consumers. While some CA design attributes are similar across the two groups, there are also differences and unique features for each that derive from the amount and type of data available for personalization.

To better understand which design elements impact the personalization of consumer interactions with CAs for Buff and Ghost consumers, we structure our discussion around two interdependent processes that are essential to understanding what the consumer needs and how to tailor the CA interaction to match the consumer’s needs.⁶ The first process focuses on *understanding consumers* and involves the collection of available data and construction of consumer profiles. The second process focuses on *generating responses* and includes matching products, services, or information to consumers’ needs and emotions, and communicating these to the consumers in conversations that are tailored to the individual.

While some dimensions of these processes are similar to current web personalization practices (e.g., session context modelling that uses click stream for data collection), others are unique to personalization with CAs (e.g., use of anthropomorphism for conversational presentation, emotion and sentiment tracking, etc.). One of the basic differences lies in the fact that CA interaction design needs to incorporate principles of both *intellectual quotient (IQ)* in being able to understand and respond with accuracy to consumer needs and *emotional*

⁶ These processes are based on the three-process framework for personalized recommendations by Adomavicius and Tuzhilin (2005). We excluded the last process of their framework which focuses on impact since our discussion is restricted to CA design and not downstream consequences. The processes are also adapted to the context of CAs based on specifics of the CA architecture and design guidelines for the chat module of the CA (e.g., see Shum et al. 2018).

Table 4 Research questions on how to encourage voluntary self-disclosure through conversational agents

Ethical Anthropomorphism	<i>Anthropomorphic Features</i>
	What anthropomorphic features are most effective towards encouraging self-disclosure? How do these features vary depending on whether the consumer is a Ghost consumer or a Buff consumer? How do these features vary based on the task (e.g., customer service vs. purchasing) and on the context (e.g., retail vs. healthcare vs. financial)? What homophilous anthropomorphic features are most effective, and do the features change based on outcomes of interest (e.g., trust vs. satisfaction)? What form of homophilous anthropomorphism is possible for Ghost consumers? What combination/level of anthropomorphic features leads to the “uncanny valley” and does this vary by Buff and Ghost consumer?
	<i>Mediating Mechanisms</i> What are the mediating mechanisms between anthropomorphic features and self-disclosure? Which anthropomorphic features operate through which mediating mechanisms? How do these mediating mechanisms vary by Buff and Ghost consumer?
	<i>Privacy Calculus and Trust</i> What anthropomorphic or other CA design features shift the privacy calculus towards self-disclosure? What anthropomorphic features influence the assessment of benefits of the CAs and what are these benefits (e.g., enjoyment, usefulness, etc.)? What anthropomorphic features influence the assessment of risks of the CAs (e.g., build trust that reduces assessment of risk)?
Transparency and Fairness	What institutional and structural assurances (e.g., certifications, audit, etc.) are effective in providing assurances of algorithmic fairness? How does their effectiveness in encouraging interaction and self-disclosure vary by Buff and Ghost consumer? How does transparency influence self-disclosure? What type and level of transparency? How does this vary by Buff and Ghost consumer? What is the appropriate level and type of transparency for CAs in different contexts (e.g., customer service vs. sales)? Can too much CA transparency backfire in different contexts?

quotient (EQ) in establishing an emotional connection and identifying the consumer’s emotions through the conversation and generating responses that are emotionally appropriate, social, and engaging (Shum et al. 2018). Figure 1 shows the design implications of each process for Buff and for Ghost consumers. We organize our discussion around these processes and follow the structure in Fig. 1.

Understanding Buff and Ghost consumers

To personalize the consumer interaction, one must understand the consumer (Adomavicius & Tuzhilin, 2005; Johar, Mookerjee, & Sarkar, 2014; Shum et al. 2018; Tam & Ho, 2006). In order to understand the two different types of consumers we discuss here, CAs must first elicit and collect data from the consumers (*data collection*), use this data to estimate their preferences and build profiles (*building profile*) that they will consequently use to tailor the interaction (Mobasher, 2007; Mobasher, Cooley, & Srivastava, 2000). There are unique differences in these steps between Buffs and Ghosts, since the type of data available to the CA a priori as well as some of the methods used to elicit data from them will be different for each type of consumer. Table 5 shows the differences between Buff and Ghost consumers on the information sharing spectrum.

Data collection Two different mechanisms currently inform the collection of data: explicit and implicit methods

(Adomavicius, Huang, & Tuzhilin, 2008; Li & Karahanna, 2015; Murthi & Sarkar, 2003). Explicit methods directly ask consumers for data, whereas implicit methods infer preferences by monitoring consumers’ behaviors (e.g., product views on a website). Since Buff consumers self-disclose private data to firms, CAs in this group can rely on stored demographic data and on implicit methods of data collection by tracking consumers’ online behaviors across devices and interactions. This implicit form of data collection is equivalent to how information is currently gathered and used for web, mobile, and other kinds of personalization services offered in the Surface Web (Chung, Wedel, & Rust, 2016). Such identifiable individual-level data can then be leveraged to build the profile of consumers.

The data collection method for Ghost consumers, however, will have to rely more heavily on explicit methods because such consumers are not identifiable a priori. As a result, CAs in this group must elicit stated needs and preferences by asking questions and engaging in two-way conversations about, for example, the purpose of buying a product, features or attributes of a desired item, etc. (Qiu & Benbasat, 2009). Research suggests that rich contextual information in long conversations with CAs may enable CAs to recognize consumers’ interests and intent even more accurately than having stored consumer profiles in which the data and information may be incomplete or ambiguous (Shum et al. 2018), making CAs’ potentially explicit methods of eliciting preferences of value to Buff consumers as well. Further, since people’s future actions are more

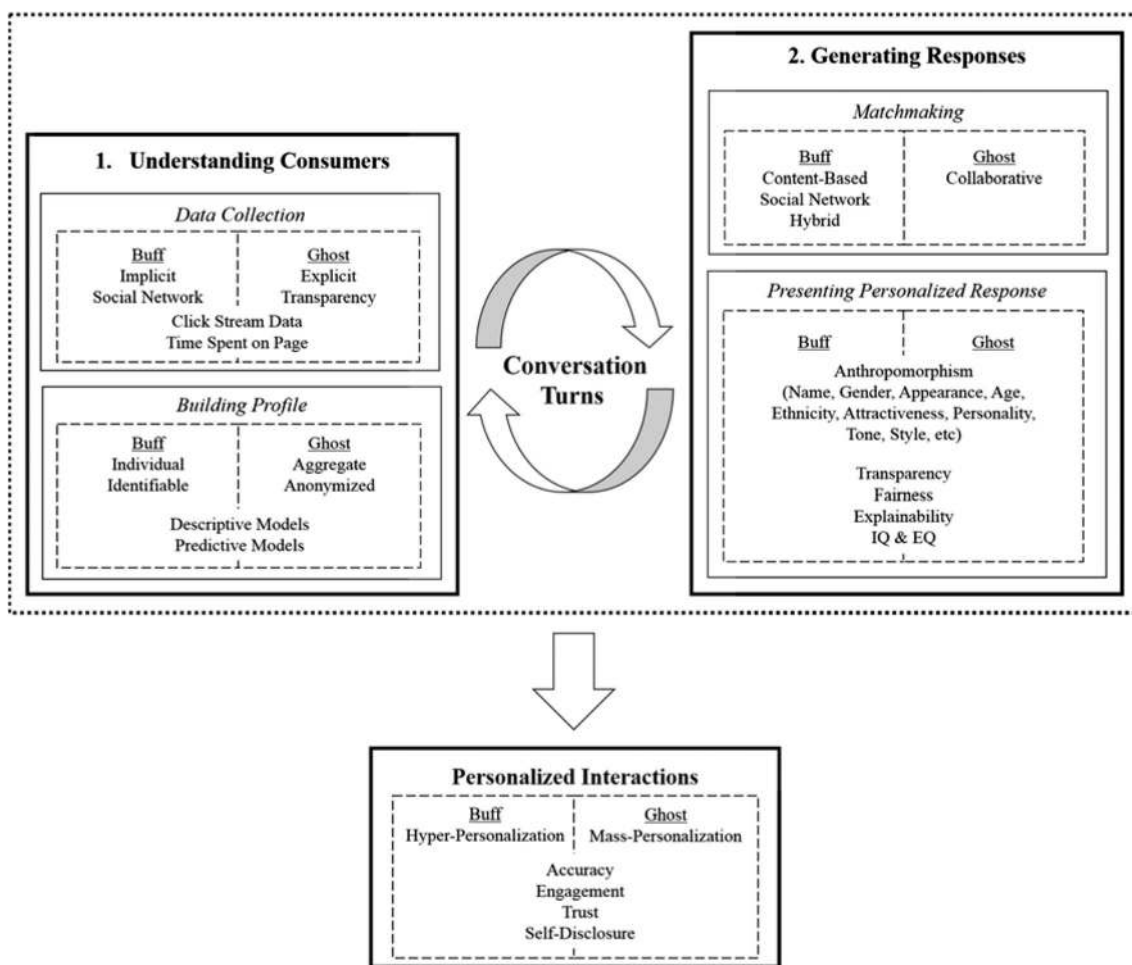


Fig. 1 Process of personalizing the CA conversation. Notes: Stages are adapted from the three-stage recommendation process model developed by Adomavicius and Tuzhilin (2005). The processes are adapted to the

context of CAs based on specifics of the CA architecture and design guidelines for the chat module of the CA (see Shum et al. 2018)

dependent on their past behaviors rather than their stated preferences (Hosanagar, 2019), CAs that also elicit data on prior behavioral activities (e.g., “what other products did you search for in the last week?”) will be able to build more accurate consumer profiles. As we have already discussed, anthropomorphization, norms of reciprocity in the conversation, and other conversational strategies are important in explicit methods of data collection to encourage data sharing, especially for Ghost consumers. For CAs to garner trust and enhance the provision of personal data from Ghosts, they may need to be able to be transparent about why a certain question was asked, how they will use the data provided, and what will happen to the data once the conversation is over. Such ask-but-explain-why transparency can mitigate feelings of vulnerability (Martin & Murphy, 2017) and incentivize Ghost consumers to share additional data with the CA. The ability to use CAs to collect personal data from people in one-on-one conversations is unique and key to understanding Ghosts better.

In addition to explicit methods of data collection, session context modelling by CAs (e.g., Shum et al. 2018) allows them to gather and use interactive click stream data *during the specific interaction* between the CA and the consumer (i.e., implicit data, see Johar et al., 2014; Padmanabhan, Zheng, & Kimbrough, 2001). This modelling approach can dynamically inform the CA’s understanding of the consumer and their intent and is an especially useful source of data to personalize interactions with Ghost consumers. For example, if a Ghost consumer clicks on a specific product and stays on the product’s page for some amount of time, CAs can utilize such behavior to gauge their interest in the product and provide helpful responses that can nudge the consumer towards purchase. The profile of Ghosts will, therefore, be based on their explicit answers to programmed questions, dynamic behavior during a specific session extracted through session context modelling, and also informed by anonymized aggregate data of other consumers behaving in similar ways.

Table 5 Differences between Buff and Ghost consumers

	Bufs	Ghosts
Data Available <i>Prior to</i> CA Interaction	Demographic data, prior transaction data, and tracking consumer online behaviour across devices and interactions	None
Data Collected <i>During</i> CA Interaction	Answers to explicit questions Session Information (clickstream data) Emotions and Sentiments	Answers to explicit questions Session Information (clickstream data) Emotions and Sentiments
Data Retention <i>After</i> CA Interaction	Full history of interaction and clickstream data	Anonymized data is retained and aggregated with the other anonymized data to update aggregate profiles for Ghost consumers
Profile Available <i>Prior to</i> CA Interaction	Identifiable-Individual of the Buff (Rich)	Anonymized-Aggregate of Other Ghosts (Poor)
Profile Available <i>After</i> CA Interaction	Identifiable-Individual of the Buff (Richer since more data were collected)	Anonymized-Aggregate of this Ghost and Others (Less poor since more data were collected)
Personalized Interaction	Hyper-Personalized	Mass-Personalized

Given that CAs rely on natural language understanding, two other activities (separate from consumer profiling and session context modelling) are also important to understanding the consumer (Shum et al. 2018): (1) understanding the message, and (2) emotion and sentiment tracking. Understanding the message involves semantic encoding and intent understanding, that is, understanding the purpose of the message (e.g., Tur & Deng 2011; Vinyals & Le, 2015). This task is easier when the CA has a consumer profile in place (as in the case of Buff consumers) where preferences and a history of interactions can facilitate intent inference. However, emotion and sentiment tracking are generally based on the current interaction (e.g., Yang et al. 2016, Chen et al. 2016) and can be used to personalize interaction with Buff and Ghost consumers alike.

Building profile As the consumer data is collected by the CA, generating personalized interactions requires integrating the data collected to iteratively build accurate and holistic consumer profiles (Adomavicius et al., 2008; Gao, Liu, & Wu, 2010). Personalization of CA interactions occurs *at each conversation turn* as part of the CA’s response generation process, and is influenced by the consumer profile that has been developed up to that point in the conversation.

The more data collected by the CA, the smaller the group segmentations for Ghosts (with the segments getting smaller in size as the information disclosed is increased) and the more complete individualization for Buffs (i.e., segments of one) resulting in more personalized CA interactions. Thus, the amount of data collected by CAs will determine the number of different profiles constructed along with the level and value of CA personalization. Many systems utilize a collection of individual or aggregate consumer facts (e.g., demographic information, favorite product, amount spent in an online store) to represent factual profiles in relational databases (Adomavicius & Tuzhilin, 2001). But, as Adomavicius et al. (2008, p. 65) note “factual profiles may not be sufficient in

certain more advanced personalization applications.” This observation is particularly true for Buff consumers, who will interact with CAs that utilize more advanced profiling techniques and leverage granular aspects of their behavior. These techniques may include descriptive models, such as rules, sequences, and signatures (Tuzhilin, 2008), or predictive models such as logistic regressions, neural networks, decision trees, support vector machines (SVM), and Bayesian networks (Adomavicius et al., 2008; Murthi & Sarkar, 2003). Though these models can be applied to both Buff and Ghost consumers, our discussion that follows highlights the differences in the types of data used for each group and the types of inferences made in terms of preferences.

Descriptive rules rely on CAs to examine the attributes of consumers and their respective activities (identifiable for Buffs and anonymized for Ghosts) to derive preferences using a variety of data mining techniques, such as association rules and classification rule discovery (Adomavicius & Tuzhilin, 2005). The sequence approach uses processual browsing activities to infer consumer preferences (Mannila, Toivonen, & Verkamo, 1997; Niu, Yan, Zhang, & Zhang, 2002). With this technique, CAs can leverage frequent episodes and other methods to learn sequential patterns of behavior, constructing profiles for both Buff (unique individual path) and Ghost (typical journey) consumers. Signatures are the data structures used to capture the evolving behavior learned from large streams of simple transactions (Cortes et al., 2000). An example signature is “top 5 most frequently browsed product categories over the last 30 days.” This signature can be stored in the profile of a specific person (Buff) or in the profile of a typical consumer (Ghost) (Adomavicius & Gupta, 2009). Finally, predictive models are based on various aspects of consumer behavior and can be built either for a specific person (Buffs) or a whole segment of similar individuals (Ghosts) (Adomavicius et al., 2008). While descriptive and predictive models represent advanced profiling techniques, more research is still needed to understand what models (descriptive

vs predictive) are more effective for CAs interacting with different types of consumers under different conditions. Similarly, more research is also needed to develop new models (e.g., prescriptive) of consumer profile building in this new hyper-private environment.

Generating responses for Buff and Ghost consumers

As the data are collected and the profiles of different consumers are iteratively built, CAs need to leverage the information they have in order to engage in personalized interactions with consumers. The first step here is *matchmaking*, which involves the identification of products, services, and information that accurately match the profile of consumers (Adomavicius & Tuzhilin, 2005; Johar et al., 2014; Mobasher, 2007). This matching would imply personalizing and guiding the conversation to align with what the CA has assessed as the consumer's behaviors, preferences, emotions, and needs while minimizing the consumers' effort (e.g., if a Buff customer interacts with a CA every week to ask about the status of their portfolio of investments, the CA can anticipate this, tailor the conversation, and provide this piece of information unprompted). After matching consumer preferences and selecting an appropriate response, CAs must engage in conversations with consumers to *present the personalized information*. (This dynamic two-way conversational channel is unique to CAs—web personalization, for instance, is not a dialogue but rather one-way where consumers receive personalized recommendations for related products or services.)

Matchmaking Existing approaches differ in that they use different sources of information to match consumers preferences for products, services, or information (Adomavicius et al., 2008). These approaches include (a) *content-based*, which typically uses the consumers' stored profile that includes historical ratings, viewing behavior, and purchases to match preferences; (b) *social networks*, which leverages the social connections of consumers to match their preferences assuming that people who are friends with one another tend to have similar characteristics and preferences; (c) *collaborative filtering*, which matches consumers' preferences based on the preferences of others who exhibit similar behaviors/preferences as the consumer; and (d) a *hybrid* approach combining the above methods (Adomavicius & Tuzhilin, 2005; Adomavicius & Gupta, 2009; Arazy, Kumar, & Shapira, 2010; Li & Karahanna, 2015). The type of approach leveraged by the CA will heavily depend on the type of data used to build consumer profiles (Li & Karahanna, 2015), and consequently on whether the consumer is Buff or Ghost. For example, while collaborative filtering can be used for both types of consumers, social network approaches would only be feasible for Buff consumers, and content-based approaches for Ghost consumers would be constrained to data extracted from the

session context modelling because of the lack of historical data on these consumers.

Presenting personalized response The natural language dialog interaction style of CAs offer the possibility not only to personalize the response to match the consumer's product or service preferences, but also to personalize the conversation in an anthropomorphic way (e.g., tone, style, accent, humor, sociability) to further match the consumer's personality and emotions. Therefore, in addition to identifying *what* to respond to the consumer through the various matchmaking approaches, it is important to identify *how* to respond to the consumer.⁷ According to Shum et al. (2018, p. 6), a CA "may generate responses in attractive styles (e.g., having a sense of humor) that improve user engagement. It needs to guide conversation topics and manage an amicable relationship in which the user feels he or she is well understood and is inspired to continue to converse with the bot," which is important for both types of consumers but more so for engaging Ghost consumers and understanding their needs. Generating responses that reflect a consistent CA personality makes the conversation easier and more predictable for the consumer and generates trust (Shum et al. 2018). As such, CA personality information (e.g., age, gender, etc.) is often incorporated into the process of generating responses (e.g., see Li et al. 2006; Mathews et al. 2015, Shum et al. 2018).

In addition, a conversation style that embodies both IQ in terms of the accuracy of the responses provided as well as EQ in terms of emotion appropriateness, can facilitate the generation of trust that what is being presented in the conversation is accurate, fair, explainable, and made benevolently in the interest of the consumer. While the personalization of "what" is delivered to the user may be hampered by the limited profile data of Ghost consumers, the personalization of "how" the conversation is conducted is likely less hampered, since it relies heavily on session information and personality settings of the CA.

Personalized interactions

As illustrated in Fig. 1, a *personalized interaction* is the resulting outcome of an iterative process of *understanding consumers* and *generating responses* that takes place as CAs chat with Buffs and Ghosts. Given that Buff consumers allow firms to implicitly collect their personal data and create identifiable profiles based on their historical interactions with the firm, the interactions they have with CAs will be hyper-personalized, more accurate, and generated with content-based or

⁷ Clearly, there are many different things impacting how CAs respond to consumers (e.g., the strategic framing of a message and its linguistic features). Due to space limitations, however, we focus our discussion of the "how" on anthropomorphism.

social network approaches, or a hybrid of the two. Such hyper-personalization makes it easier and faster for customers to interface with CAs since their existing patterns can be used to anticipate future requests, identify information relevant to their needs, and recommend products or services that match their preferences, and thus reduce search costs. In contrast, Ghost consumers will enjoy interactions that are tailored based on mass-personalization. Since the personal preferences of Ghost consumers are not stored in their consumer profiles, CAs will have to engage in conversations and elicit their stated preferences, needs, and prior behaviors during each conversation in order to provide them with more personalized experiences. Then, personalizing the interaction will be based on the data elicited through the conversation (the more data collected by the CA the more personalized the interaction will be) and collaborative filtering, where the consumer’s profile (dynamically built during the interaction) will be matched with

aggregate profiles of other similar consumers. The interaction will therefore be personalized based on patterns that emerge from these other aggregate profiles.

The discussion above is meant to be illustrative of differences and similarities in CA design across the two consumer groups and is neither meant to be exhaustive nor comprehensive. In general, more research is needed to inform which design elements result in the creation of engaging and personalized, but also ethical and unobtrusive CAs. Table 6 presents some illustrative research questions along these lines organized around our framework. Deriving personalization benefits of using CAs without being “creepy” while morally attending to different needs of the two groups is important because, despite successes, there have been many failures and many challenges still remain in designing CAs that provide high quality interactions (Ben Mimoun, Poncin, & Garnier, 2012; Chakrabarti & Luger, 2015; Gnewuch et al., 2017;

Table 6 Research agenda for personalizing consumer interactions through conversational agents in the era of hyper-privacy

Process of Personalizing the CA Conversation	Research Questions
Understanding Consumers	
<i>Data Collection</i>	
Explicit / Implicit	What data collection techniques, or combinations, are most effective in collecting the type of data needed to build consumer profiles and personalize interactions for Ghost consumers?
Click Stream	What is the right combination of explicit and implicit methods that minimizes effort for Ghost consumers without giving rise to privacy concerns (e.g., explicit methods require more effort but generate fewer privacy concerns whereas session modelling (implicit) methods does not involve effort but can give rise to privacy concerns)?
Time Spent on Page	
<i>Building Profile</i>	
Individual / Aggregate Identifiable / Anonymized	What models (descriptive vs predictive) are more effective for CAs interacting with Buff and with Ghost consumers and how do these vary depending on the task (e.g., whether it is transactional or informational) and context (e.g., healthcare vs. shopping vs. financial)?
Session Context Modelling	Development of new models (e.g., prescriptive) of consumer profile building that rely heavily on explicit methods of questioning, session context modelling, and emotion and sentiment tracking
Emotion & Sentiment Models	
Generating Responses	
<i>Matchmaking</i>	
Content-Based	What matching methods are most effective for Ghost consumers? And does this vary by the nature of the task and context?
Collaborative	Under what conditions is it more effective to match the intent and preferences of Ghost consumers based on collaborative, content-based, and social networks rather than only using a typical hybrid approach (content-based and social networks)?
Social Network	
Hybrid	
<i>Presenting Personalized Response*</i>	
Anthropomorphism	What linguistic and dialog design practices can CAs use to engage Buff and Ghost consumers? How does the effectiveness of these vary between the two types of consumers?
Transparency	What is the right level of EQ for Buff and Ghost consumers? Does too much lead to privacy concerns?
Fairness	
Explainability	
IQ & EQ	
Personalized Interactions	
Mass- / Hyper- Personalization	Which models and combination of models (data collection, profile building, matching, response generation) provide more accurate and engaging interactions and a higher level of perceived personalization?
Accuracy	Do these vary for Buff and Ghost consumers?
Engagement	To what extent is personalization for Buff consumers more accurate than for Ghost consumers? How does this vary by task and by context?
Trust	How does the level of personalization (or perceived personalization) influence (a) trust, (b) the privacy-calculus, (c) self-disclosure? How does this vary by Buff and Ghost consumers?
Self-Disclosure	How does the level of personalization (or perceived personalization) influence consumer outcomes (e.g., customer satisfaction, loyalty) and firm outcomes (e.g., sales)? Does this vary by Buff and Ghost consumer?

Note: * See Table 4 for Ethical Anthropomorphism, Transparency, and Fairness Research Questions

Jenkins et al., 2007; McTear, Callejas, & Griol, 2016; Schuetzler et al., 2014; Shechtman & Horowitz, 2003).

Conclusion

The Web is an evolving complex system that impacts both firms and consumers. We review the evolution of the Web, and note that Web changes are driven by dynamics of information control embedded in market power disputes between firms and consumers. Based on this review, we suggest that the Web will significantly change in the next five to ten years. Stakeholder privacy violations will lead government agencies to enforce new legislations, resulting in an information sharing culture of “opt-in” where consumers will by default be ghost consumers and no longer allow firms to collect, use, and share their personal information with other organizations. As a

result, firms will lose the ability to create in-depth profile of consumers, and personalized practices like micro-targeting may be at risk.

To survive this coming change, marketers will have to incentivize consumers to remain in the buff, as many are today, while also serving the needs of consumers who deny access and become ghost consumers. We suggest that CAs will play an increasingly important role in helping firms market to both ghost and buff consumers. In particular, we argue that CAs can be used to understand and engage both Ghost and Buff consumers by developing personalized interactions for both groups of consumers (see Fig. 1), albeit in different ways. We also suggest that CA design may be instrumental in nudging consumers to self-disclose private information. CAs that do so well can become a source of differentiation and competitive advantage for the firm.

Appendix

Table 7 Existing studies on anthropomorphic features of conversational agents

Article	Features	Theory	Variables	Results	Context	Method
Aggarwal & McGill (2007)	Smile	Schema Congruity	IV: Feature Type Mediator: Anthropomorphic Product Perceptions Moderator: Affective Tag DV: Product Evaluation	The ability of consumers to anthropomorphize a product and their consequent evaluation of that product depend on the extent to which that product is endowed with characteristics congruent with the proposed human schema. Consumers' perception of the product as human mediates the influence of feature type on product evaluation. The affective tag attached to the specific human schema moderates the evaluation but not the successful anthropomorphizing of the product	Marketing	Experiment
Al-Natour et al. (2006)	Appearance; Gender; Personality; Decision Strategy; Online	Decisional Guidance & Speech Act Theory	DVs: Perceived Similarity	Anthropomorphic features manifest desired personalities & behaviors	E-commerce	Experiment
Al-Natour et al. (2011)	Appearance; Gender; Personality; Decision Strategy; Online	Social Psychology & HCI Research	IV: Perceived Similarity DVs: Usefulness; Ease of Use; Trust; Social Presence; Enjoyment;	Perceived similarity influences enjoyment, social presence, trust, ease of use, & usefulness	E-commerce	Experiment
Araujo (2018)	Name; Language Style; Framing; Online	Computers as Social Actors	Mediator: Social Presence DVs: Satisfaction; Mindless / Mindful Anthropomorphism; Attitude; Emotional Connection	The usage of human-like language or name were sufficient to increase perception of the agent as being human-like. Adopting an intelligent frame does reduce perceptions of mindless anthropomorphism for machine-like agents, yet no such difference was	E-commerce	Experiment

Table 7 (continued)

Article	Features	Theory	Variables	Results	Context	Method
Mimoun & Poncin (2015)	Gender; Appearance; Ethnicity; Online	Selling Research	IVs: Playfulness; Social Presence DV: Satisfaction; Purchase Intention; Decision Quality Mediator: Shopping Value	found for the human-like agent. The usage of human-like cues had significant positive influence on emotional connection Use of a CA increases shopping value; shopping value mediates the effects of playfulness & social presence on satisfaction & intentions (hedonic) & decision quality (utilitarian)	E-commerce	Survey
Ben Mimoun et al. (2017)	Gender; Appearance; Ethnicity; Online		IVs: Involvement; Product Familiarity; Internet Skills; Need for Interaction DV: E-Consumer Productivity	Objective e-consumer productivity depends on interaction with CA. Individual characteristics affect perceived productivity, either independently from CA use (for involvement or product familiarity) or in interaction with CA (for Internet skills and need for interaction)	E-commerce	Experiment
Chattaraman, Kwon, & Gilbert (2012)	Appearance; Online	Computers as Social Actors & Social Response Theory	Mediators: Trust; Social Support; Perceived Risk DV: Trust; Intention	Embedding a CA that serves search and navigational support functions leads to increased perceived social support, trust, and intention. The effect of CAs on intentions is mediated by and perceived risks and the effect on trust is mediated by perceived social support	E-commerce	Experiment
Ciechanowski et al. (2019)	Appearance; Language Style; Online	The idea of the uncanny valley	DVs: Affect, Electromyography, Respirometer, Electrocardiography, & Electrodermal Activity DV: Trust	Participants experienced less uncanny effects and less negative affect in cooperation with a text chatbot than with an animated avatar chatbot. The text chatbot also induced less intense psychophysiological reactions	Academy Enrolment Process & Casual Conversation	Experiment
Elkins & Derrick (2013)	Tone; Online	Interpersonal Adaptation Theory	DV: Trust	Vocal pitch was inversely related to perceived trust, but temporally variant; vocal pitch early in the interview reflected trust. The CA was perceived as more trustworthy when smiling	Mock-Screening Interview	Experiment
Knijnenburg & Willemsen (2016)	Appearance; Language Style; Online	Human-Computer Interaction & Agent-Based Interfaces	Mediator: Anthropomorphic Beliefs DVs: Exploitation; Use	Users of more humanlike agents try to exploit capabilities that were not signaled by the system. This severely reduces the usability of systems that look human but lack humanlike capabilities. Users of humanlike agents also form anthropomorphic beliefs about the system: They act humanlike towards the system and try to exploit typical humanlike capabilities they believe the system possesses	Tourism	Experiment
Li et al. (2019)	Voice; Online	Computers as Social Actors	IV: Space Type Mediator: Self-Awareness DV: Self-disclosure	Users were most willing to disclose information about their tastes and interests and least willing to disclose money information. Users in the living space were willing to disclose more information than those in the workspace, which was mediated by users' expectations for the reciprocal services of CAs rather than the awareness of other persons or external factors	Control & Question & Answer Task	Experiment
Louwerse et al. (2005)	Appearance; Voice; Gender; Online	Social Agency Theory & Social Cue Hypothesis	DV: Comprehension & Impression Scores	Participants preferred natural agents with natural voices. Although female agents with male voices formed an exception, this was explained by a stereotype effect	Education Services	Experiment Experiment

Table 7 (continued)

Article	Features	Theory	Variables	Results	Context	Method
Mende et al. (2019)	Appearance; Ethnicity; Tone; Name; Physical	The idea of the uncanny valley	Mediator: Eeriness Level, Perceived Threat to Human Identity; Machinization of CA Moderator: Social Belonging & Food Type DV: Compensatory Response; Customer Food Consumption	Consumers display compensatory responses when they interact with CAs rather than a human. CAs elicit greater consumer discomfort (i.e., eeriness and a threat to human identity), which in turn results in the enhancement of compensatory consumption. Compensatory responses are mitigated when social belongingness is high, food is perceived as more healthful, and the CA is machinized (rather than anthropomorphized)		
Qiu & Benbasat (2009)	Appearance; Tone; Online	Social Relationship Perspective	Mediator: Social Presence DVs: Trust; Enjoyment; Intention	Anthropomorphism influences users' perceptions of social presence, which in turn enhances users' trusting beliefs, perceptions of enjoyment, and ultimately, their intentions to use the agent as a decision aid	E-commerce	Experiment
van Doorn et al. (2017)	Tendency to Anthropomorphize	Social cognition & Psychological Ownership	Mediator: Social Cognition; Psychological Ownership Moderator: Customer's Relationship Orientation; Technology Readiness DVs: Social Presence	We have proposed that social cognition and psychological ownership mediate the relationship between automated social presence (ASP) and service outcomes and that key factors (relationship orientation, anthropomorphism, and technology readiness) moderate the effects of ASP.	Services	Conceptual
Verhagen et al. (2014)	Appearance; Personality; Communication Style; Online	Implicit Personality, Social Response, Emotional Contagion, Social Interaction	IVs: Friendliness; Expertise; Smile DVs: Social Presence; Personalization; Satisfaction	An empirical study confirms the cross channel applicability of friendliness and expertise as determinants of social presence and personalization	E-commerce	Experiment

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Acquisti, A., Brandimarte, L., & Loewenstein, G. (2015). Privacy and human behavior in the age of information. *Science*, *347*(6221), 509–514.
- Adjerid, I., Acquisti, A., & Loewenstein, G. (2018). Choice architecture, framing, and cascaded privacy choices. *Management Science*, *65*(5), 2267–2290.
- Adomavicius, D., & Tuzhilin, A. (2005). Personalization technologies: A process-oriented perspective. *Communications of the ACM*, *48*(10), 83–90. <https://doi.org/10.1007/s11576-006-0098-7>.
- Adomavicius, G., & Tuzhilin, A. (2005). Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge & Data Engineering*, *6*(6), 734–749.
- Adomavicius, G., & Gupta, A. (2009). *Business Computing*. Emerald Group Publishing.
- Adomavicius, G., Huang, Z., & Tuzhilin, A. (2008). Personalization and Recommender Systems. In Z.-L. Chen, S. Raghavan, P. Gray, & H. J. Greenberg (Eds.), *State-of-the-Art Decision-Making Tools in the Information-Intensive Age* (pp. 55–107). <https://doi.org/10.1287/educ.1080.0044>
- Adomavicius, G., & Tuzhilin, A. (2001). Using data mining methods to build customer profiles. *Computer*, *34*(2), 74–82. <https://doi.org/10.1109/2.901170>.
- Aggarwal, P., & McGill, A. L. (2007). Is that Car smiling at me? Schema congruity as a basis for evaluating anthropomorphized products. *Journal of Consumer Research*, *34*(4), 468–479. <https://doi.org/10.1086/518544>.
- Aguirre, E., Roggeveen, A. L., Grewal, D., & Wetzels, M. (2016). The personalization-privacy paradox: Implications for new media. *Journal of Consumer Marketing*, *33*(2), 98–110.
- Al-Natour, S., Benbasat, I., & Cenfetelli, R. T. (2006). The role of design characteristics in shaping perceptions of similarity: The case of online shopping assistants. *Journal of the Association for Information Systems*, *7*(12), 821–861.
- Al-Natour, S., Benbasat, I., & Cenfetelli, R. T. (2011). The adoption of online shopping assistants: Perceived similarity as an antecedent to evaluative beliefs. *Journal of the Association for Information Systems*, *12*(5), 347–374. <https://doi.org/10.17705/1jais.00267>.
- Altman, I. (1975). *The Environment and Social Behavior: Privacy, Personal Space, Territory, and Crowding*.
- Anderson, C. L., & Agarwal, R. (2011). The digitization of healthcare: boundary risks, emotion, and consumer willingness to disclose

- personal health information. *Information Systems Research*, 22(3), 469–490.
- Araujo, T. (2018). Living up to the chatbot hype: The influence of anthropomorphic design cues and communicative agency framing on conversational agent and company perceptions. *Computers in Human Behavior*, 85, 183–189.
- Arazy, O., Kumar, N., & Shapira, B. (2010). A Theory-Driven Design Framework for Social Recommender Systems. *Journal of the Association for Information Systems*, 11(9). Retrieved from <https://aisel.aisnet.org/jais/vol11/iss9/2>
- Arora, N., Dreze, X., Ghose, A., Hess, J. D., Iyengar, R., Jing, B., et al. (2008). Putting one-to-one marketing to work: Personalization, customization, and choice. *Marketing Letters*, 19(3–4), 305–321. <https://doi.org/10.1007/s11002-008-9056-z>.
- Bakken, S. A., Moeller, K., & Sandberg, S. (2018). Coordination problems in cryptomarkets: Changes in cooperation, competition and valuation. *European Journal of Criminology*, 15(4), 442–460.
- Barratt, M. J., Ferris, J. A., & Winstock, A. R. (2016). Safer scoring? Cryptomarkets, social supply and drug market violence. *International Journal of Drug Policy*, 35, 24–31.
- Benjamin, V., Valacich, J. S., & Chen, H. (2019). DICE-E: A framework for conducting Darknet identification, collection, evaluation with ethics. *MIS Quarterly*, 43(1), 1–22.
- Ben Mimoun, M. S., Poncin, I., & Garnier, M. (2012). Case study: Embodied virtual agents: An analysis on reasons for failure. *Journal of Retailing and Consumer Services*, 19, 605.
- Mimoun, M. S. B., Poncin, I., & Garnier, M. (2017). Animated conversational agents and e-consumer productivity: The roles of agents and individual characteristics. *Information & Management*, 54(5), 545–559.
- Bennett, C. J., & Raab, C. (2006). *The governance of privacy: Policy instruments in global perspective*. Cambridge: The MIT Press.
- Bleier, A., & Eisenbeiss, M. (2015). The importance of trust for personalized online advertising. *Journal of Retailing*, 91(3), 390–409.
- Bleier, A., Harmeling, C. M., & Palmatier, R. W. (2019). Creating effective online customer experiences. *Journal of Marketing*, 83(2), 98–119.
- Bojei, J., Julian, C. C., Wel, C. A. B. C., & Ahmed, Z. U. (2013). The empirical link between relationship marketing tools and consumer retention in retail marketing. *Journal of Consumer Behaviour*, 12(3), 171–181. <https://doi.org/10.1002/cb.1408>.
- Brandtzaeg, P. B., & Følstad, A. (2017). Why People Use Chatbots. In I. Kompatsiaris, J. Cave, A. Satsiou, G. Carle, A. Passani, E. Kontopoulos, ... D. McMillan (Eds.), *Internet Science* (Vol. 10673, pp. 377–392). https://doi.org/10.1007/978-3-319-70284-1_30
- Brandtzaeg, P. B., & Følstad, A. (2018). Chatbots: Changing user needs and motivations. *Interactions*, 25(5), 38–43. <https://doi.org/10.1145/3236669>.
- Bucklin, R. E., & Sismeiro, C. (2003). A model of web site browsing behavior estimated on clickstream data. *Journal of Marketing Research*, 40(3), 249–267.
- Burke, R. R. (1997). Do you see what I see? The future of virtual shopping. *Journal of the Academy of Marketing Science*, 25(4), 352–360.
- Bursztein, E. (2017). Understanding how people use private browsing. Retrieved February 6, 2019, from Elie Bursztein's site website: <https://www.elie.net/blog/privacy/understanding-how-people-use-private-browsing>
- Caudevilla, F., Ventura, M., Fornis, I., Barratt, M. J., Vidal, C., Quintana, P., et al. (2016). Results of an international drug testing service for cryptomarket users. *International Journal of Drug Policy*, 35, 38–41.
- Caudill, E. M., & Murphy, P. E. (2000). Consumer online privacy: Legal and ethical issues. *Journal of Public Policy & Marketing*, 19(1), 7–19.
- Chakrabarti, C., & Luger, G. F. (2015). Artificial conversations for customer service chatter bots. *Expert Systems with Applications*, 42(20), 6878–6897. <https://doi.org/10.1016/j.eswa.2015.04.067>.
- Chattaraman, V., Kwon, W. S., & Gilbert, J. E. (2012). Virtual agents in retail web sites: Benefits of simulated social interaction for older users. *Computers in Human Behavior*, 28(6), 2055–2066.
- Chatterjee, P., Hoffman, D. L., & Novak, T. P. (2003). Modeling the clickstream: Implications for web-based advertising efforts. *Marketing Science*, 22(4), 520–541.
- Chen, H., Sun, M., Tu, C., Lin, Y., & Liu, Z. (2016, November). Neural sentiment classification with user and product attention. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing* (pp. 1650–1659).
- Chung, T. S., Wedel, M., & Rust, R. T. (2016). Adaptive personalization using social networks. *Journal of the Academy of Marketing Science*, 44(1), 66–87. <https://doi.org/10.1007/s11747-015-0441-x>.
- Claycomb, C., & Martin, C. L. (2001). Building customer relationships: An inventory of service providers' objectives and practices. *Marketing Intelligence & Planning*, 19(6), 385–399. <https://doi.org/10.1108/EUM000000006109>.
- Cortes, C., Fisher, K., Pregibon, D., Rogers, A., & Smith, F. (2000). Hancock: A language for extracting signatures from data streams. In *Proc. of the 2000 ACM SIGKDD Intl. Conf. on Knowledge Discovery and Data Mining*, 9–17.
- Danaher, B., Dhanasobhon, S., Smith, M. D., & Telang, R. (2010). Converting pirates without cannibalizing purchasers: The impact of digital distribution on physical sales and internet piracy. *Marketing Science*, 29(6), 1138–1151.
- Daugherty, P. R., & Wilson, H. J. (2018). Human + machine: Reimagining work in the age of AI. Retrieved from <http://public.eblib.com/choice/publicfullrecord.aspx?p=5180063>
- Deighton, J. (1997). Commentary on "exploring the implications of the internet for consumer marketing". *Journal of the Academy of Marketing Science*, 25(4), 347–351.
- Dholakia, N., & Zwick, D. (2001). Privacy and consumer agency in the information age: between prying profilers and preening webcams. *Journal of Research for Consumers*, 1(1).
- Dinev, T., & Hart, P. (2006). An extended privacy Calculus model for E-commerce transactions. *Information Systems Research*, 17(1), 61–80. <https://doi.org/10.1287/isre.1060.0080>.
- Duxbury, S. W., & Haynie, D. L. (2019). Criminal network security: An agent-based approach to evaluating network resilience. *Criminology*, 57(2), 314–342.
- Edvardsson, B., Tronvoll, B., & Gruber, T. (2011). Expanding understanding of service exchange and value co-creation: A social construction approach. *Journal of the Academy of Marketing Science*, 39(2), 327–339.
- Elkins, A. C., & Derrick, D. C. (2013). The sound of trust: voice as a measurement of trust during interactions with embodied conversational agents. *Group decision and negotiation*, 22(5), 897–913.
- Gao, M., Liu, K., & Wu, Z. (2010). Personalisation in web computing and informatics: Theories, techniques, applications, and future research. *Information Systems Frontiers*, 12(5), 607–629. <https://doi.org/10.1007/s10796-009-9199-3>.
- Garfinkel, S., Matthews, J., Shapiro, S. S., & Smith, J. M. (2017). Toward algorithmic transparency and accountability. *Communications of the ACM*, 60(9), 5–5. <https://doi.org/10.1145/3125780>.
- Ghose, A., Ipeirotis, P. G., & Li, B. (2012). Designing ranking systems for hotels on travel search engines by mining user-generated and crowdsourced content. *Marketing Science*, 31(3), 493–520.
- Givon, M., Mahajan, V., & Muller, E. (1995). Software piracy: Estimation of lost sales and the impact on software diffusion. *Journal of Marketing*, 59(1), 29–37.
- Gnewuch, U., Morana, S., & Maedche, A. (2017). Towards Designing Cooperative and Social Conversational Agents for Customer Service. 15.

- Hann, I. H., Hui, K. L., Lee, S. Y. T., & Png, I. P. (2008). Consumer privacy and marketing avoidance: A static model. *Management Science*, 54(6), 1094–1103.
- Holt, T. J., Smirnova, O., & Chua, Y. T. (2016). Exploring and estimating the revenues and profits of participants in stolen data markets. *Deviant Behavior*, 37(4), 353–367.
- Hosanagar, K. (2019). *A Human's guide to machine intelligence: How algorithms are shaping our lives and how we can stay in control*. New York: Viking.
- Huang, M.-H., & Rust, R. T. (2017). Technology-driven service strategy. *Journal of the Academy of Marketing Science*, 45(6), 906–924. <https://doi.org/10.1007/s11747-017-0545-6>.
- Jardine, E. (2018). Tor, what is it good for? Political repression and the use of online anonymity-granting technologies. *New Media & Society*, 20(2), 435–452.
- Jain, S. (2008). Digital piracy: A competitive analysis. *Marketing Science*, 27(4), 610–626.
- Jenkins, M.-C., Churchill, R., Cox, S., & Smith, D. (2007). Analysis of User Interaction with Service Oriented Chatbot Systems. Proceedings of the 12th International Conference on Human-Computer Interaction: Intelligent Multimodal Interaction Environments, 76–83. Retrieved from <http://dl.acm.org/citation.cfm?id=1769590.1769600>
- Johar, M., Mookerjee, V., & Sarkar, S. (2014). Selling vs. profiling: Optimizing the offer set in web-based personalization. *Information Systems Research*, 25(2), 285–306. <https://doi.org/10.1287/isre.2014.0518>.
- Kehr, F., Kowatsch, T., Wentzel, D., & Fleisch, E. (2015). Blissfully ignorant: the effects of general privacy concerns, general institutional trust, and affect in the privacy calculus. *Information Systems Journal*, 25(6), 607–635.
- Knijnenburg, B. P., & Willemsen, M. C. (2016). Inferring capabilities of intelligent agents from their external traits. *ACM trans. Interact. Intell. Syst.*, 6(4), 28:1–28:25. <https://doi.org/10.1145/2963106>.
- Kumar, N., & Benbasat, I. (2006). Research note: The influence of recommendations and consumer reviews on evaluations of websites. *Information Systems Research*, 17(4), 425–439. <https://doi.org/10.1287/isre.1060.0107>.
- Ladegaard, I. (2019). Crime displacement in digital drug markets. *International Journal of Drug Policy*, 63, 113–121.
- Lamberton, C., & Stephen, A. T. (2016). A thematic exploration of digital, social media, and mobile marketing: Research evolution from 2000 to 2015 and an agenda for future inquiry. *Journal of Marketing*, 80(6), 146–172.
- Lambrecht, A., & Tucker, C. (2013). When does retargeting work? Information specificity in online advertising. *Journal of Marketing Research*, 50(5), 561–576.
- Langer, E. J. (1992). Matters of mind: Mindfulness/mindlessness in perspective. *Consciousness and Cognition*, 1(3), 289–305. [https://doi.org/10.1016/1053-8100\(92\)90066-J](https://doi.org/10.1016/1053-8100(92)90066-J).
- Larivière, B., Bowen, D., Andreassen, T. W., Kunz, W., Sirianni, N. J., Voss, C., et al. (2017). “Service encounter 2.0”: An investigation into the roles of technology, employees and customers. *Journal of Business Research*, 79, 238–246. <https://doi.org/10.1016/j.jbusres.2017.03.008>.
- Leong, B., & Selinger, E. (2019). Robot Eyes Wide Shut: Understanding Dishonest Anthropomorphism. Proceedings of the Conference on Fairness, Accountability, and Transparency - FAT* '19, 299–308. <https://doi.org/10.1145/3287560.3287591>
- Li, S., & Karahanna, E. (2015). Online recommendation systems in a B2C E-commerce context: A review and future directions. *Journal of the Association for Information Systems*, 16(2), 72–107. <https://doi.org/10.17705/1jais.00389>.
- Li, H., Sarathy, R., & Xu, H. (2011). The role of affect and cognition on online consumers' decision to disclose personal information to unfamiliar online vendors. *Decision Support Systems*, 51(3), 434–445.
- Li, W., Chen, H., & Nunamaker, J. F. (2016). Identifying and profiling key sellers in cyber carding community: AZSecure text mining system. *Journal of Management Information Systems*, 33(4), 1059–1086.
- Louwerse, M. M., Graesser, A. C., Lu, S., & Mitchell, H. H. (2005). Applied Cognitive Psychology: The Official Journal of the Society for Applied Research in Memory and Cognition. *Social cues in animated conversational agents*, 19(6), 693–704.
- Mannila, H., Toivonen, H., & Verkamo, A. I. (1997). Discovery of Frequent Episodes in Event Sequences. 31.
- Martin, K. D., & Murphy, P. E. (2017). The role of data privacy in marketing. *Journal of the Academy of Marketing Science*, 45(2), 135–155. <https://doi.org/10.1007/s11747-016-0495-4>.
- Mathews, A. Xie, L., & He, X. (2015). SentiCap: generating image descriptions with sentiments, Conference of the Association for the Advancement of Artificial Intelligence (AAAI), 2015.
- McAfee, A., & Brynjolfsson, E. (2012). Big Data: The Management Revolution. Harvard Business Review, (October 2012). Retrieved from <https://hbr.org/2012/10/big-data-the-management-revolution>
- McTear, M., Callejas, Z., & Griol, D. (2016). *The Conversational Interface: Talking to Smart Devices* (1st ed.). Springer Publishing Company, Incorporated.
- Mende, M., Scott, M. L., van Doorn, J., Grewal, D., & Shanks, I. (2019). Service robots rising: How humanoid robots influence service experiences and elicit compensatory consumer responses. *Journal of Marketing Research*, 56(4), 535–556. <https://doi.org/10.1177/0022243718822827>.
- Mimoun, M. S. B., & Poncin, I. (2015). A valued agent: How ECAs affect website customers' satisfaction and behaviors. *Journal of Retailing and Consumer Services*, 26, 70–82.
- Mobasher, B. (2007). *The Adaptive Web* (P. Brusilovsky, A. Kobsa, & W. Nejdl, Eds.). Retrieved from <http://dl.acm.org/citation.cfm?id=1768197.1768201>
- Mobasher, B., Cooley, R., & Srivastava, J. (2000). Automatic personalization based on web usage mining. *Communications of the ACM*, 43(8), 142–151. <https://doi.org/10.1145/345124.345169>.
- Montgomery, A. L., Li, S., Srinivasan, K., & Liechty, J. C. (2004). Modeling online browsing and path analysis using clickstream data. *Marketing Science*, 23(4), 579–595.
- Moon, Y. (2000). Intimate exchanges: Using computers to elicit self-disclosure from consumers. *Journal of Consumer Research*, 26(4), 323–339. <https://doi.org/10.1086/209566>.
- Moon, Y. (2003). Don't blame the computer: When self-disclosure moderates the self-serving bias. *Journal of Consumer Psychology*, 13(1/2), 14.
- Mori, M., MacDorman, K. F., & Kageki, N. (2012). The uncanny valley [from the field]. *IEEE Robotics & Automation Magazine*, 19(2), 98–100.
- Murthi, B. P. S., & Sarkar, S. (2003). The role of the management sciences in research on personalization. *Management Science*, 49(10), 19.
- Nass, C., & Moon, Y. (2000). Machines and mindlessness: Social responses to computers. *Journal of Social Issues*, 56(1), 81–103. <https://doi.org/10.1111/0022-4537.00153>.
- Nass, C., Steuer, J., & Tauber, E. R. (1994). Computers Are Social Actors. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, 72–78. <https://doi.org/10.1145/191666.191703>
- Neff, G., & Nagy, P. (2016). Talking to bots: Symbiotic agency and the case of Tay. *International Journal of Communication*, 10, 17.
- Niu, L., Yan, X.-W., Zhang, C.-Q., & Zhang, S.-C. (2002). Product hierarchy-based customer profiles for electronic commerce recommendation. Proceedings. International Conference on Machine Learning and Cybernetics, 2, 1075–1080 vol.2. <https://doi.org/10.1109/ICMLC.2002.1174549>

- Padmanabhan, B., Zheng, Z., & Kimbrough, S. O. (2001). Personalization from Incomplete Data: What You Don't Know Can Hurt. In Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD-2001, 154–163).
- Peterson, R. A., Balasubramanian, S., & Bronnenberg, B. J. (1997). Exploring the implications of the internet for consumer marketing. *Journal of the Academy of Marketing Science*, 25(4), 329.
- Phelps, J., Nowak, G., & Ferrell, E. (2000). Privacy concerns and consumer willingness to provide personal information. *Journal of Public Marketing*, 19, 27–44.
- Picard, R. W. (1997). MIT Media Laboratory; Perceptual Computing; 20 Ames St., Cambridge, MA 02139 picard@media.mit.edu, <http://www.media.mit.edu/~picard/>. 16.
- Qiu, L., & Benbasat, I. (2009). Evaluating anthropomorphic product recommendation agents: A social relationship perspective to designing information systems. *Journal of Management Information Systems*, 25(4), 145–182. <https://doi.org/10.2753/MIS0742-1222250405>.
- Sahni, N. S., Wheeler, S. C., & Chintagunta, P. (2018). Personalization in email marketing: The role of noninformative advertising content. *Marketing Science*, 37(2), 236–258. <https://doi.org/10.1287/mksc.2017.1066>.
- Sankar, R., & Balakrishnan, K. (2016). Implementation of an inquisitive chatbot for database supported knowledge bases. *Sadhana*, 41(10), 6.
- Schuetzler, R. M., Grimes, M., Giboney, J. S., & Buckman, J. (2014). Facilitating Natural Conversational Agent Interactions: Lessons from a Deception Experiment. *Human Computer Interaction*, 17.
- Schumann, J. H., von Wangenheim, F., & Groene, N. (2014). Targeted online advertising: Using reciprocity appeals to increase acceptance among users of free web services. *Journal of Marketing*, 78(1), 59–75. <https://doi.org/10.1509/jm.11.0316>.
- Shechtman, N., & Horowitz, L. M. (2003). Media inequality in conversation: How people behave differently when interacting with computers and people. Proceedings of the Conference on Human Factors in Computing Systems - CHI '03, 281. <https://doi.org/10.1145/642611.642661>
- Short, J., Williams, E., & Christie, B. (1976). The social psychology of telecommunications. London; New York: Wiley.
- Shum, H., He, X., & Li, D. (2018). From Eliza to XiaoIce: Challenges and opportunities with social chatbots. *Frontiers of Information Technology & Electronic Engineering*, 19(1), 10–26. <https://doi.org/10.1631/FITEE.1700826>.
- Sinha, R. K., Machado, F. S., & Sellman, C. (2010). Don't think twice, it's all right: Music piracy and pricing in a DRM-free environment. *Journal of Marketing*, 74(2), 40–54.
- Sinha, R. K., & Mandel, N. (2008). Preventing digital music piracy: The carrot or the stick? *Journal of Marketing*, 72(1), 1–15.
- Smith, H. J., Dinev, T., & Xu, H. (2011). Information privacy research: An interdisciplinary review. *MIS Quarterly*, 35(4), 989–1016.
- Smith, H. J. (2001). Information privacy and marketing: What the US should (and shouldn't) learn from Europe. *California Management Review*, 43(2), 8–33.
- Spangler, W. E., Hartzel, K. S., & Gal-Or, M. (2006). Exploring the privacy implications of addressable advertising and viewer profiling. *Communications of the ACM*, 49(5), 119–123. <https://doi.org/10.1145/1125944.1125951>.
- Staff, F. M. (2016). Try Hello Hipmunk™ to Chat Before You Pack. Retrieved February 6, 2019, from Tailwind by Hipmunk website: /tailwind/lets-talk-travel-try-hello-hipmunk-to-chat-before-you-pack/.
- Stewart, D. W. (2017). A comment on privacy. *Journal of the Academy of Marketing Science*, 45(2), 156–159.
- Summers, C. A., Smith, R. W., & Reczek, R. W. (2016). An audience of one: Behaviorally targeted ads as implied social labels. *Journal of Consumer Research*, 43(1), 156–178.
- Tam, & Ho. (2006). Understanding the impact of web personalization on user information processing and decision outcomes. *MIS Quarterly*, 30(4), 865. <https://doi.org/10.2307/25148757>.
- Trusov, M., Ma, L., & Jamal, Z. (2016). Crumbs of the cookie: User profiling in customer-base analysis and behavioral targeting. *Marketing Science*, 35(3), 405–426.
- Tur, G., & Deng, L. (2011). Intent determination and spoken utterance classification. Spoken language understanding: systems for extracting semantic information from speech. Wiley, Chichester, 93–118.
- Turkle, S. (2017). Alone together: Why we expect more from technology and less from each other. /z-wcorg/.
- Tuzhilin, A. (2008). Personalization: The state of the art and future directions. In Business Computing. Retrieved from https://www.researchgate.net/publication/302958247_Personalization_The_state_of_the_art_and_future_directions
- Urban, G. L., Liberali, G., MacDonald, E., Bordley, R., & Hauser, J. R. (2013). Morphing banner advertising. *Marketing Science*, 33(1), 27–46.
- Van Buskirk, J., Roxburgh, A., Bruno, R., Naicker, S., Lenton, S., Sutherland, R., et al. (2016). Characterising dark net marketplace purchasers in a sample of regular psychostimulant users. *International Journal of Drug Policy*, 35, 32–37.
- van Doorn, J., Mende, M., Noble, S. M., Hulland, J., Ostrom, A. L., Grewal, D., & Petersen, J. A. (2017). Domo arigato Mr. Roboto: Emergence of automated social presence in organizational frontlines and customers' service experiences. *Journal of Service Research*, 20(1), 43–58. <https://doi.org/10.1177/1094670516679272>.
- Verhagen, T., van Nes, J., Feldberg, F., & van Dolen, W. (2014). Virtual customer service agents: Using social presence and personalization to shape online service encounters. *Journal of Computer-Mediated Communication*, 19(3), 529–545. <https://doi.org/10.1111/jcc4.12066>.
- Vinyals, O., & Le, Q. (2015). A neural conversational model. arXiv preprint arXiv:1506.05869.
- Vlahos, J. (2018). Barbie Wants to Get to Know Your Child. The New York Times. Retrieved from <https://www.nytimes.com/2015/09/20/magazine/barbie-wants-to-get-to-know-your-child.html>
- Wang, W., & Benbasat, I. (2008). Attributions of trust in decision support technologies: A study of recommendation agents for e-commerce. *Journal of Management Information Systems*, 24(4), 249–273.
- Warren, S. D., & Brandeis, L. D. (1890). Right to privacy. *Harv. L. Rev.*, 4, 193.
- Westin, A. F. (1967). *Privacy and freedom*. New York: Atheneum.
- Yang, Z., Yang, D., Dyer, C., He, X., Smola, A., & Hovy, E. (2016). Hierarchical attention networks for document classification. In Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies.
- Yu, J., Hu, P. J. H., & Cheng, T. H. (2015). Role of affect in self-disclosure on social network websites: A test of two competing models. *Journal of Management Information Systems*, 32(2), 239–277.
- Yue, T. W., Wang, Q. H., & Hui, K. L. (2019). See no evil, hear no evil? Dissecting the impact of online hacker forums. *MIS Quarterly*, 43(1), 73.
- Xu, H., Teo, H.-H., Tan, B. C. Y., & Agarwal, R. (2010). The role of push-pull technology in privacy calculus: The case of location-based services. *Journal of Management Information Systems*, 26(3), 135–174.