University of Alberta

Anonymous and Pseudo-Anonymous Behaviors Online: Are Full Identities Truly Better?

by

Kristen Smirnov

A thesis submitted to the Faculty of Graduate Studies and Research
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

in

Marketing

School of Business

Fall 2013

Edmonton, Alberta

Abstract

In discussions of identity in online environments, the current momentum in both academic literature and industry practices pushes toward full identity revelation. Along with market-driven justifications, there is also a stated belief that revealing one's identity will lead to better behavior toward others, and in general. This has a long history of support in computing, psychology, and economics literature, but these investigations have examined full identities relative to completely anonymous behaviors. To do so ignores a third option: that of a persistent pseudonym which gains investment and reputation over time.

This dissertation examines the difference between real identities and pseudonyms (versus anonymous behavior) in two series of studies centered upon the production of product reviews. In them, I expect to find a short period of investment into pseudonymous identities after which they perform at quality levels equal to those people using their real names, as they now feel accountable to this secondary identity. Real-world product review data from Amazon.com was collected and two laboratory studies were run. It was found that in situations of voluntary identity disclosure, the investment period was much longer than anticipated. However, during mandatory identity disclosure, real name users suffered strong performance penalties that were generally avoided by those reviewers using a pseudonym. Potential explanations are offered and future research questions are identified.

Acknowledgment

I would like to first thank my parents, without whose support in every aspect of my life I would have found this a nearly impossible task. Next I would like to thank my advisor, John Pracejus, for his advice and support over the years, and who brought focus to my very broad interest in online identity. As well, deep thanks go to Paul Messinger, the first to invite me into a research project and a key contributor to the interest in identity that would later become this dissertation. I am also deeply indebted to the other members of my dissertation committee: Gerald Häubl, Sarah Moore, Sharla King, and Amna Kirmani. I also offer thanks to my friends Todd Harper, Karen Healey, and James Renken for being a constant source of empathy about the challenges of this work, and to Neil Brigden for helping me learn the best ways to conduct a study on MTurk. I hope I was able to offer some tips in return.

Table of Contents

List of Tables

List of Figures

Introduction

> "Facebook has always been based on a real-name culture," said Elliot Schrage, vice president of public policy at Facebook. "We fundamentally believe this leads to greater accountability and a safer and more trusted environment for people who use the service." (Sengupta 2011)

> "The real risk to the world is if information technology pivots to a completely authentic identity for everyone," said Joichi Ito, head of the Media Lab at the Massachusetts Institute of Technology. "In the U.S., maybe you don't mind. If every kid in Syria, every time they used the Internet, their identity was visible, they would be dead." (Sengupta 2011)

In 2011, Salman Rushdie discovered that his Facebook account had been deactivated. When he sent them a copy of his passport to prove his identity and ownership of the account, it was reactivated, but not under the name through which he became famous. Instead, he was listed under the identity on his passport: "Ahmed Rushdie" (Taylor 2011). The famous recluse, target of a fatwa on his head that forced him to live under police protection for years, had his public identity selected for him by a company insisting upon specific name usage by its customers (Madrigal 2011; Sengupta 2011). After pressure via Twitter, a social networking service that allows complete anonymity in contrast to Facebook's strict naming policies, Facebook retracted their earlier decision and allowed Rushdie to be listed under his chosen name. Even then, his naming options were limited; he was only allowed to list himself as Salman because it was his middle name, rather than a completely unrelated pseudonym (Madrigal 2011; Sengupta 2011; Taylor 2011).

Access to users' real identities and information can play out poorly in the marketplace, beyond publically notable figures such as Mr. Rushdie. Girls Around Me was an iPhone application that was released through Apple's App Store. This application pulled the user's location and compared it to recent activity in the Foursquare system. (Foursquare is a social network that keeps track of who has recently opted to share their arrival in certain locations; e.g., they may report that they are "checking in" to their local Panera when they stop there for a sandwich.) Upon providing these results, Girls Around Me also linked back to these women's Facebook accounts, so anyone considering approaching these women would have a full information sheet about them handy for use. In short, they would not be flirting with someone with whom they must chat ignorantly about their surroundings or the weather, nor with someone who would understand why this stranger knew exactly which bands to mention or school to claim as their alma mater. After a surge of complaints, Girls Around Me was removed from the Apple App Store (Bilton 2012). However, this application only streamlined and collected information that was already available to any public viewer: these women had chosen to start Facebook accounts and post personal information about themselves, they had chosen to not lock down this information, and they had chosen to share their current location in an ongoing stream of information to another social network. When these side effects of sharing social information were distilled into an application that made use of it, consumers were shocked.

Financial drivers are one obvious reason behind this push to make consumers' real world information available online. It's estimated that companies

currently spend US$2 billion per year for access to this personal data that has been shared online (Sengupta 2011). If we are increasingly tied into a persistent online identity that reveals our chosen shopping locales, entertainment preferences, political leanings, and so forth, these profiles become very valuable sources of data. However, as is apparent from the two quotes that opened this piece, there are conflicting opinions on matters beyond that. Those pushing for greater privacy measures warn against total identity revelation, while those in favor of using our real names online often speak not of money, but of a greater social good that will provide consumers with a safer, more enjoyable online environment.

This dissertation examines these assumptions of modern businesses, namely that more identity revelation is superior to maintaining an anonymous online identity, and questions certain behaviors currently playing out in the marketplace. It examines other options for improving users' behaviors while allowing them to maintain anonymity. It also seeks answers to a question that has gone unanswered in the drive toward this monetization of personal information: are there times when remaining anonymous, or rather *pseudononymous* (with a persistent *chosen* identity, as opposed to a persistent *real* identity), may actually be *more* beneficial than revealing one's entire real identity? To my knowledge, no examination exists of the basic behaviors customers engage in when allowed to be anonymous, possible benefits to anonymity in specific consumer behavior realms, nor how their behaviors may be manipulated through structural decisions by the hosting sites. In other words, an obvious hole exists in the literature: what

fundamental differences does anonymity make on basic consumer behaviors, and how do changes in the structure of online sites alter these consumer behaviors? By addressing this understudied area, I believe this work will make a significant contribution to the literature on anonymity in computer-mediated consumer behavior.

Beyond these obvious financial drivers, there are many reasons why companies would believe it beneficial to have their users identify themselves online. Visual contact is a key feature of human interaction. Essentially, we crave methods of assessing people upon meeting them. With only minor mental effort paid (Olson and Marshuetz 2005; Sui and Liu 2009), a viewer will see a person and believe him or her to be more talented (Landy and Sigall 1974), trustworthy (van 't Wout and Sanfey 2008; Zebrowitz, Voinescu, and Collins 1996), or generally well-regarded as a person (Dion, Berscheid, and Walster 1972), based on elements of physical appearance. Our brains take in large amounts of information about a person quickly, and just as quickly make judgments about that other member of a conversation. When an appearance is no longer available to us, we lose out on a significant portion of our standard method of interaction. If online communication exists in a space in which we've effectively lost one of our senses, it contains some inherent logic to not further limit the information we do have about others.

Furthermore, within psychology, it is well studied that remaining anonymous can lead to worsened behavior (Diener et al. 1976; Festinger, Pepitone, and Newcomb 1952; Singer, Brush, and Lublin 1965). As will be

discussed at length during the literature review, it is a common assumption among both the general population and specific businesses that forcing people to disclose their full identities will result in improved behavior. Given the level to which we expect to see information about a person and use such information to make judgments about them, this appears to have surface validity. Such motivations surely contribute to Facebook's insistence upon using one's real name on an account, and an increasingly common trend among comment-driven websites where users use their real names (for example, as associated with their Facebook account, or as required for elite status on the Yelp review site) to participate in discussions[1]. As this returns some sense of social normality to the discussion—we may once again "see" our conversation partner—it appears to be a positive. However, dangers lurk for both companies and consumers when such a broad brush is used.

Marketing literature largely examines a traditional setting in which both parties meet face to face, or when an unseen partner is represented by an impersonal figure such as an online retailer. However, the modern reality of consumer interactions includes computer-mediated communications that may or may not end with revealing personal identities. Thus, traditional studies on "anonymity" do not acknowledge many aspects of anonymity as experienced by modern consumers. There exists a hole in our understanding of both how people

[1] One exception to this trend is Twitter, which also allows log-ins for some sites in a manner similar to Facebook, but does not require real identities to be used. Indeed, due to how Twitter only verifies identities if Twitter has contacted someone associated with the account, e.g. a major celebrity's agent or corporation's PR division, there is no Twitter-based way of verifying the true identity of the account holder (Twitter.com 2012). These accounts do, however, exist as persistent pseudonymous identities, a concept which will receive further discussion in subsequent sections.

behave under online anonymity in general, and in particular, under the different levels of identity disclosure possible in a modern online environment.

This dissertation examines changes in online behavior that accompany using different levels of disclosed personal identities; specifically, in the realm of online product reviews. This is done first through real world data gathered from a major retailer, for correlational analysis across several sample populations. The findings are then taken into an experimental setting. Through this, I address questions left unanswered in the current literature on online consumer behavior, and on modern online identity in general. This literature is reviewed in depth in the following section, followed by my research questions and the studies themselves.

Literature

In this section, I first review literature across fields of marketing, social psychology, and computing science to identify what is known about anonymity in an online context, and what is known in a broader sense that would have applicability to online users. More importantly, the weaknesses in how these findings apply to modern online behaviors are identified, which was a motivation for this dissertation. Second, I develop a series of models that communicate the framework of this study, in an application of that literature.

Anonymity in Online Marketing Research

Anonymity is understudied within the field of online consumer behavior. Despite the increasing prevalence of online consumer interactions as we accustom

ourselves to online socialization and consumption, researchers have delved into only certain, limited aspects of consumer behavior. It is acknowledged that deception during communications is easier behind a veil of anonymity, even by corporations themselves who choose to interact with consumers without revealing their identities, when identifying themselves would be more ethically appropriate (Dellarocas 2006; Mayzlin 2006). Another of these studied areas is consumer response to the attempts they make to interact with companies. Here, it was found that anonymity has an impact on how much information they are willing to share during online transactions (Hoffman, Novak, and Peralta 1999). As well, consumers value their privacy differently, and recommendations are made to segment consumers on this personal difference, based on their willingness to share data with companies (Hann et al. 2007; Milberg et al. 1995). Attempting to predict consumer behavior by their preferred level of privacy is challenging, however, as there are indications that any personal differences in comfort may be unstable (John, Acquisti, and Loewenstein 2011). Put differently: for companies to measure consumers' willingness to expose their personal information is challenging for reasons of both complexity and limited consumer trust toward companies on this matter, and attempting to force any given level of privacy loss may go poorly. This second finding is particularly relevant, given the example about Salman Rushdie that opened this dissertation.

A relevant study for this work finds that consumers trust anonymous reviews of hedonic versus utilitarian products differently (Sen and Lerman 2007), but an explanation is not given for what drives consumers to make such

anonymous reviews in the first place. In other research potentially relevant to this work, however, there are studied impacts on the behavior of others based on whether they read negative or positive reviews. If anonymity promotes the posting of more negative reviews due to a sense of social freedom, this is worthy of consideration. Reading negative reviews can counter social pressure to buy something (Huang and Chen 2006), and has an outsized effect on the earliest stages of a product's introduction (Basuroy, Chatterjee, and Ravid 2003). Although retailers might be wary of promoting negative reviews, even indirectly, by allowing consumers to post them under a freeing identity, it is also worth noting that the volume of discussion about a product can improve its overall sales, regardless of that discussion's valence (Liu 2006).

Also relevant is a study done of auction site design. The persistent identities available on eBay—for one example—are an example of a structure that is examined in this work. There, a user may choose any identity they wish (such as VintageBaseballCards) that does not reveal their personal identity. As the user participate in auctions and gets positive feedback from other community members, this ID gradually gets a higher user rating next to itself on site pages. These auction settings are one of the areas in which the topic of moderate identity disclosure has received an increase in attention over marketing as a whole. There, users do recognize the potential ill effects of anonymity, and see the value of a persistent online identity in the form of reputation (Diekmann and Wyder 2002), rather than necessarily desiring full identity disclosure in the Facebook model.

The most thorough examination of online consumer privacy motivations provides a model that does not generally examine how consumers behave behind a veil of anonymity, but specifically what leads them to provide fabricated information online (Lwin and Williams 2003). Furthermore, it has a particular interest in the effects this has on corporate strategy, such as the increased market research costs that must be borne when customers deliberately provide false information, rather than on the behavior of the consumers themselves. Studies have examined anonymous reviews (Chatterjee 2001; Sen and Lerman 2007), with the acknowledgement that negative information is more heavily weighted than positive (Mizerski 1982). However, to my knowledge, there has been no overall comparison of anonymous reviews versus those with an attached identity. Consumers may believe that the use of anonymity removes some of the reporting bias, or an unwillingness to accurately portray reality (Eagly, Wood, and Chaiken 1978), meaning that other consumers are more willing to paint a more accurate picture of their experiences if they are allowed to participate anonymously. However, given the previous discussions about auction participants wishing to have some form of persistent identity during their transactions, they may put more value on those reviews with names attached. However, this is currently unclear, and any differences in the quality of output are not yet identified.

My own previous work with Messinger, et al. (2008; 2009) contributed to this dissertation's focus on a persistent online identity. In these papers, we examined the behavior of users of certain online environments and the visual avatars they chose to create. We found that great care was taken in the

construction and maintenance of many of these identities, and that they bore strong ties to the users' true, real world selves. These users had *invested* themselves quite strongly into these fictional identities that were used to socialize, entertain, and even purchase goods and consume advertising materials online. With that bridge between more formalized consumer behavior studies and our look at more general online environments, a broader look at anonymity in general online environments is now taken in full.

Anonymity in Online Environments

This literature looks at several levels of computer-mediated communication (CMC), including: purely anonymous text-based interactions, interactions with an associated, persistent fictional identity (pseudonyms), and interactions with an associated real identity, as is desired by Facebook and similar organizations. Through examining this spectrum, I hope to determine which factors of anonymity can alter certain aspects of consumer behavior. For, despite being largely overlooked in the consumer behavior literature as a potential driving force for users increasingly participating in a space with malleable identities, it may have a large impact on how people behave. At this point it seems appropriate to move beyond the marketing realm, and look at computing science literature that demonstrates similar findings and may apply to my consumer studies.

This idea has two major features identified in traditional psychological literature, both of which are relevant to my work on CMC as whole. The concepts of anonymity—not being tied to a social identity—and invisibility—not having their faces or behaviors seen—are two of the key factors for disinhibition in

online environments, or the removal of typical constraints on one's behavior (Suler 2004). (These concepts are explained in greater depth in the following section, which focuses upon social psychology literature.) Both clearly apply to many online interactions, in which we are never seen as more than words on a screen and in which we are allowed to represent ourselves with identities separate from our true names.

Suler labels anti-social behavior arising from disinhibition as "toxic disinhibition." Numerous examples exist. A recent and illuminating example involved a singer named Rebecca Black who posted a vanity music video to YouTube. After the video went viral and earned attention and mocking from major media outfits, she revealed that she received death threats from anonymous strangers because of her song. Black was thirteen years old at the time (Moses 2011). It is unlikely that these users would have threatened an adolescent's life to her face and yet they felt free to do so in the anonymous and invisible online environment. Older work found that anonymous subjects encouraged those expressing suicidal thoughts to follow through with their threats (Mann 1981), further demonstrating the possible severity of how people may feel free to behave when free of the constraints of identity and visibility. In a world where we are separated from any sort of identity, whether our own or that of a target's, it becomes all too easy to view another person as nothing more than a line of text on a screen, to whom no real respect needs to be paid, and cruel behavior is indeed more present in CMC relative to face-to-face interactions (Kiesler et al. 1985; Orenga et al. 2000; Siegel et al. 1986; Sproull and Kiesler 1986). This poor

behavior ranges from generally being less polite to other users to extremes of encouraging them to carry through with stated suicide threats. With this in mind, the drive toward full identity revelation online becomes an understandable pursuit, and even noble. Returning to the earlier discussion of Facebook's drive toward full identity revelation, it seems like a favorable confluence of motivations: they are able to both benefit their business' balance sheet and the overall behavior of their users, both within the Facebook system and in the greater internet ecosphere (Sengupta 2011).

However, generally speaking, people prefer not to share information about themselves (Cialdini 1993; Kelly and McKillop 1996; Lane and Wegner 1995). This holds true not only in specific online contexts, but also as a general trend in human behavior. Returning to CMC, subjects allowed to remain anonymous engaged in far more self-disclosure compared to those who revealed their identities (Christopherson 2007). This is useful in forming certain social bonds and addressing therapeutic needs. Teenagers can establish identities, and those concerned about revealing information (such as homosexuality) have a safe space to do so with similar parties, where they will not be found out by people they know in real life (Christopherson 2007). Platonic and romantic relationships can form with less pressure (Hinduja 2008), and people may be generally more willing to talk about themselves and their lives (Joinson 2001; Lee, Im, and Taylor 2008). This "benign disinhibition" even allows for people go out of their way to help others (Suler 2004). For example, people may be willing to share potentially sensitive or embarrassing information that they are aware could help

others, if they will not be personally compromised or shamed once that information is made public. As previously discussed, people prefer to retain control over their identities, and when we feel safer because of this control we may engage in more open and trust-based activities (Schul, Mayo, and Burnstein 2008). A man may forgo sharing his experience with brands of erectile dysfunction drugs under his true name, as sharing his product usage would reveal sensitive personal information; in contrast, he might be willing to make recommendations to other men in need if he is allowed to share his experiences anonymously, or pseudonymously.

Furthermore, the popular wisdom that deindividuation via anonymity is the driving force behind the criminal behavior of software piracy did not hold up in scientific study. Users were no more likely to pirate software if they valued being anonymous and pseudonymous online (Hinduja 2008). This challenges a recurring complaint that anonymity leads toward this broad area of anti-social behavior, although the authors note that outright criminal behavior will only be seen in a small segment of any population. As well, people may be more inclined to give anonymous others the benefit of the doubt when they create some given content, as too many specifics about a more richly detailed identity ends up distracting from the content itself (Klar and Giladi 1997). In particular, these two findings challenge specific *business* rationales for demanding more identity disclosure from consumers: the first one has obvious sales implications, and the second may impact the increasing level of user-generated content online. Thus, the first indications arise that it may not always be wise and helpful to remove

anonymity from users as a method of improving their behavior, and that perhaps not all of the assumptions being made hold true. As noted, this is an understudied area, and so it is difficult to make sweeping claims. However, I believe there to be strong indications that this is a valuable field of study in which to add to our understanding of consumer interactions with privacy online.

An additional layer of complexity arises from users' comfort levels with revealing their identities. It has been acknowledged that we make snap judgments of others based on their physical appearance. In the current integration happening with Facebook accounts, on sites such as USA Today's comment section, people *must* log in with a Facebook account if they wish to participate. Their Facebook profile picture is attached to every comment they make on those articles. In situations such as this, people may not wish to give up the anonymity of CMC as they perceive it to provide a more equal playing field (Haraway 1990; Poster 1990; Sproull and Kiesler 1991). Particularly for minority groups, this potential of reducing others' inherent social power is an attractive prospect as it grants a relative increase in their own social position (Dubrovsky, Kiesler, and Sethna 1991; Kiesler, Siegel, and McGuire 1984; Siegel et al. 1986). For example, an anonymous member of a lower social class who often feels that their opinions are viewed as uneducated, less important, or simply lesser may now feel capable of debating a political stance with someone with a diploma from Harvard Law School or the Kellogg School of Business. Demonstrating this, men are more willing to reduce anonymity in CMC than are women (Flanagin et al. 2002), as women feel that sexist behaviors limiting their ability to participate equally in

discussions no longer exist when they are allowed to participate under an anonymous handle. Famously, in the 1990s, Joanne Rowling became "JK Rowling" out of concern that she would not be broadly marketable as a female author (Duffy 1999). It cannot be said that she would *not* have succeeded in marketing *Harry Potter and the Philosopher's Stone* and its sequels as Joanne Rowling, but she is certainly an example of successful, sex-driven identity manipulation.

Furthermore, privacy, or maintaining control over the boundaries of others' access to one's self (Pedersen 1997), is generally positive for well-being. It allows us to gain the desired level of social contact at any given time (Altman 1975), delineate boundaries (Altman and Chemers 1980), adjust to new settings (Vinsel et al. 1980), and comfortably maintain relationships (Werner, Altman, and Brown 1992). This conflict means that, despite seeming like one obvious fix to improve behavior in CMC, the removal of anonymity may not be the ideal method to both balance consumer comfort and positive consumer behavior toward others. As well, when we have fewer concerns about losses of privacy about our true identity, we are more willing to disclose negative actions in which we've participated, as there is less feeling of accountability (John et al. 2011). This has strong applicability to many consumer contexts that might be encountered online, particularly the online product review context examined here.

Anonymity in Social Psychology and Consumer Behavior Research

Major papers in the field of social psychology have found strong effects for anonymity on worsened behavior. However, I find this literature to lack certain applicable features to the CMC researched in this work, due to assumptions made about the definition of anonymity in physical versus electronic contexts. Through a discussion of previous findings, and then an examination of the ways in which these findings do not apply to online environments, I identify certain elements of anonymity that differ between traditional social psychology research and modern consumer online behavior. In particular, the traditional understanding toward a sense of deindividuation is examined, with a focus on its contributing factors and the uniformly negative view much of the literature takes toward it. By this, and questioning whether an assumption of negative behaviors may actually prove false, I demonstrate my research's motivation and identify its contribution to the marketing literature.

Our behavior is altered when we no longer must show our own identities to the world. The concept of deindividuation becomes important when people no longer feel constrained by their typical identity or feel it become less salient. A key feature of this process is anonymity. Anonymity allows for anti-normative and anti-social behaviors to arise that would normally be constrained by feeling stronger ties to one's self (Diener et al. 1976; Festinger et al. 1952; Singer et al. 1965). Feeling this tie to our real selves reduces our likelihood to engage in

socially unacceptable behaviors, and feeling free of it increases the likelihood of pursuing them.

However, these findings—although strong and repeated by numerous researchers—reflect a model of anonymity that applies more to the physical world than electronic. There is a consistent assumption across numerous studies that being in a large group contributes to a feeling of deindividuation, and therefore anonymity. Diener et al.'s famous work (1976) on anonymity reported that children were significantly more likely to take more than their fair share of Halloween candy when they were in costumes that covered their faces and when they were in a larger group. Under this paradigm, the more people that surround us, the more anonymous we are. This is a very peculiar assumption to make in a CMC environment where people can interact remotely with no identity markers attached to them, and people seeking to be even more anonymous may look at options such as masking their IP address or using throw-away email accounts. The crowd paradigm for deindividuation makes little surface sense in such an environment. And, although one may argue that anonymous labels and masked IPs correspond to Halloween masks far better than they do to being in a physical crowd, the source of that poorer behavior by mask-wearing children is not clear. Do masks generate negative behavior? Or do children who are already prone to bad behavior choose "bad" costumes such as monsters, devils, and ghosts, which are more likely to cover their faces to achieve an inhuman look, while children prone to good behaviors choose heroic costumes such as Superman or Disney princesses that leave their faces on display? It cannot be determined from this

work, and so there exists a danger in overgeneralizing its results (Diener et al. 1976). Furthermore, subsequent meta-analysis of this research area found little support for consistent physically-driven deindividuation effects, and concluded that deindividuation was driven more by varying social and contextual cues at an individual level (Postmes and Spears 1998).

Despite this limitation, Diener et al.'s paper was highly influential in many early studies on CMC. However, there again exists a danger in generalizing these papers toward modern online behaviors, as the CMC described in these works was frequently that of in-office work groups who balanced their interactions between face-to-face and CMC (Connolly, Jessup, and Valacich 1990; Sproull and Kiesler 1986). This is, of course, a far cry from today's global Internet.

Looking at the history of CMC research in the context of technological developments, this focus makes sense. In the 1950s, when Festinger, et al.'s important early paper (1952) about deindividuation was published, CMC did not exist. Not until ARPANET of the 1970s and NSFNET of the 1980s were universities beginning to be connected via electronic communication networks (Glowniak 1998), and certainly, regular consumers did not get online until the spread of the Internet in the 1990s. This gave ample time for research streams to develop that made certain assumptions about the definition of anonymity. Diener et al. (1976) state that "[research] on anonymity suggests that this variable may have its strongest effect when combined with group presence." Mann (1981) further supported this finding. There, it was found that being in a large group contributed to a sense of deindividuation that led to a higher incidence of

prompting suicidal persons to jump off a building or bridge. To researchers working a non-Internet communications environment, this idea makes complete sense. We are not surprised when a person in a crowd throws a rock during a riot, even when we might be very surprised to see that same person perform the same act on a nearly empty street. When we are out in public, as these studies commonly assume, the deindividuation effect can best be achieved by a feeling of group membership overcoming our typical ties to our individual identity. However, modern, Internet-driven CMC provides the chance for individuals to interact with others across a broad spectrum of activities—social, informational, commercial, etc.—with *total* anonymity. There need not be any ties to a person's true self: no name, no history of previous behaviors, no face, and no potential of learning this information.

Under this understanding of online anonymity, much of the literature on anonymous behaviors applies poorly, if at all. There, they view being in the middle of a crowd as contributing to feelings of anonymity. However, as discussed in the Girls Around Me case, modern consumers felt distinctly exposed when their physical selves were related to their online information, even though this was, by definition of the Foursquare-dependent interface, happening in public spaces (Bilton 2012; Brownlee 2012b). By deliberately masking privacy options from the user (Brownlee 2012a), the programmers made it difficult to maintain a division between information shared online and interactions in the real world. These are reactions left unexplored by the traditional literature that treats identity as something passed on with spoken introductions or a business card, or with

faces that can be lost in a sea of people. However, in a world with facial recognition, online check-ins, and Google on everyone's smartphones, such control over identities in public spaces does not apply in the ways previously explored. Based on negative reactions to this application (Brownlee 2012b), it seems worth postulating that we view our online activities and identities as something we should be able to maintain control over, and that should not leak into the physical world without our express approval.

Through examining these three areas of literature, I believe I have found many holes or questions that my work addresses. First, from a marketing perspective, I address the lack of attention paid to how anonymity can affect consumer behavior. From a social psychology perspective, I seek to answer whether previous findings about anonymity, and in particular deindividuation, hold up in a modern online environment. Finally, from a computing science perspective, I research both the positives and negatives of anonymous representation and full identity disclosure (FID) along with the understudied pseudonymous option, rather than assuming any identity level is superior.

Varying Levels of Anonymity as a Research Framework

An alternate way of considering anonymity is to divide anonymity into multiple component elements, each of which can be independently manipulated. Rather than saying someone is using their *real identity* or is *anonymous*, we recognize that there are multiple levels of identity disclosure in which people might engage.

One way to view this was proposed by Hayne and Rice (1997). They propose a view where someone may have technical anonymity, or the removal of identifying facts about a person, and social anonymity, or the removal of clues with which an identity may be attributed. Although their work looked at subjects with pre-existing relationships, who could potentially identify a participant either by their full name *or* by clues to their identity (such as gender), much of the general concept holds true in an age where many of our lives are on Google. People online, if they share their real names, can quickly become "known" to anyone who enters them into a search engine or looks them up on Facebook; from educational history to employer to favorite television shows, a person's entire identity is easily associated with full disclosure of their name. However, unlike with existing real-world relationships, we may safely reduce some level of social anonymity without becoming "known." Choosing to disclose that one is the mother of three boys, lives in Little Rock, and works at a bank would give context to a user's opinions, but does not also reveal a full identity with the accompanying knowledge of Celiac Disease, a recent divorce, and any number of personal photographs.

Although the repeal of social anonymity (to continue the use of this generally applicable terminology) would not be viewed as a positive by those users who wish to mask their identity, technical anonymity seems certain to be the worse option to lose: fewer people would likely rather have their name and precise street address shared with strangers than their hobbies. Customers were uncomfortable with even giving that level of detailed information to online

companies for transactions (Hoffman et al. 1999), and so giving it to other users seems even more unlikely. However, when allowed to maintain control over their full identity, online users are free with limited self-disclosure (such as mentioning mental health challenges, opinions about professors, or some long-held personal secret) relative to face-to-face discussions (Joinson 2001; McKenna and Bargh 1998).

There exist systematic ways to *partially* reduce one's anonymity in online environments while maintaining control over awareness of one's full identity. This paper studies the use of a persistent online pseudonym, that remains with a user while he or she stays on a site (Hinduja 2008), such as the earlier eBay example or by using an user name like "Want2VisitParis" on a forum about world travel. These handles have the benefit of being easy to use and providing a precise amount of anonymity to match any given person's comfort level. While some audiences may perceive these chosen identities to be false (Rains and Scott 2007), there is strong trust formed between many online users in such settings (Parks and Floyd 1996). Parks and Floyd's work found that over 60% of Usenet (discussion groups that were commonly used in the early days of the Internet) users reported forming personal bonds with other users, in a platform where handles were common.

An alternate take is that CMC allows users to portray their identities as the identities they wish or believe themselves to be (Donath 1999). If these identities meet positive receptions, then this contributes to knowledge contribution within online communities (Jeppesen and Frederiksen 2006; Ma and Agarwal 2007).

Certain technical features of online communities that allow others to respond to these contributions, such as social feedback via comment/messaging systems and reputation systems, helps bolster relationships (Ma and Agarwal 2007), which then contribute to even better performance. As well, recognition from a company of some given identity that is producing good contributions on that site can further increase both the quantity and quality of future output (Jeppesen and Frederiksen 2006). The takeaway from these findings, relevant for this study, is that people can see these online identities—separate from their real identities—be rewarded and viewed positively by others, and invest them with pride, effort, and social connections similar (though not identical) to how they would behave while forming social bonds and producing output under their real names.

This perception reinforces the important of developing online identities over time, such as was researched in the auction literature, where a user's rating appeared in their listings and made them look like a reliable seller. Indeed, the possibility for seeing an identity's invested effort is vital for creating functional social platforms online, where the reputation of identities matters. Dellarocas (2010) identifies several key features that a web site must contain in order to maintain a successful reputation system, three of which apply to the data analyzed in this work: building trust, promoting quality, and sustaining loyalty. To achieve a goal of online identities becoming invested and valuable to users, and functional as intended in online interactions, Dellarocas recommends that users will benefit from such features as: being able to see their raw participations statistics and progressing through tiered membership levels; scoring mechanisms for user-

generated content and being ranked by quality; cumulative metrics that continue to accrue for as long as an individual maintains that single account on a site. Dellarocas' model strongly affected the website chosen for data collection.

As previously noted, privacy provides comfort to many individuals and anonymity grants some pro-social behaviors. However, complete anonymity can also promote hurtful behaviors. It is unlikely that the removal of technical anonymity, meaning the reveal of a person's specific identity, would both increase user comfort *and* pro-social behaviors given the findings previously discussed. This question is addressed in a laboratory study. However, the reduction of social anonymity by showing a partial representation of their identity may reduce disinhibition (and the subsequent anti-social behaviors) without reducing so much anonymity that comfort levels are affected. Dellarocas' recommendations support the idea of a persistent pseudonym as an ideal approach for online interactions. As people are rewarded for behaving well with some chosen identity and sticking with it, they are shepherded toward the benefits of total identity revelation without giving up the technical anonymity that they may wish to maintain. Examining these boundaries is central in this dissertation. Table 1 (later in this section) explores many technical features of site design that are present in these varying identity levels online, to help demonstrate various methods of partially reducing anonymity.

An alternate method of viewing these persistent pseudonyms is through a lens of impression management, for as previously stated, online identities can be idealized and carefully-managed versions of our selves (Donath 1999). To look at

it from a discussion of auction identities, to use the previous example of pseudonyms being studied in marketing, is to impose a commercialized assumption upon online behaviors. Although this dissertation does use data collected on a commercial site, there exist many non-commercial sites online for entertainment, socializing, and information seeking, and impression management more generally captures this breadth of behavior.

In the real world, we are tied to our identities. When dealing with strangers, we exert more effort into maintaining a favorable presentation of ourselves (Leary et al. 1994; Tice et al. 1995). As well, we also take care to present a more favorable image of ourselves when dealing with the opposite sex (Leary et al. 1994), and use impression management methods consistent with traditional gender roles (Guadagno and Cialdini 2007). Maintaining a socially desirable image of one's self takes time and effort (Holtgraves 2004), and at every step, we must customize our approach to our own identities and the social baggage attached to it. As we're aware of these societal assumptions potential working against us, we actively work to counter stereotypical impressions about our identities (von Hippel et al. 2005).

This echoes the earlier discussion about how some underprivileged users may prefer to stay anonymous online, to avoid stereotypical assumptions about their social group applying to them. The use of pseudonyms presents the potential to eliminate concerns over interactions with the opposite sex when participants' sexes are unclear, or to reduce concern over strangers judging us personally more harshly compared to the assessments a friend would make; after all, they can only

assess the pseudonym, not the real person. Although not all users prefer pseudonyms, as I will find and discuss during my study sections, many do. Perhaps these pseudonyms provide a respite from the effort of maintaining one's real identity. They may allow a level of comfort in which to defy or conform to stereotypes and freely interact with other users in a way that some might find impossible when tied to one's real identity, with the accompanying concerns for managing impressions of one's real self.

At this point, a more precise survey of the elements of online anonymity is warranted, before a model is proposed of how the manipulation of these elements affects online behavior. The following table identifies specific ways in which identity may be disclosed online, drawing from both the literature previously discussed and my own observations while conducting this research. As well, the assessments of the frequency of these occurrences also comes from my own observations during this research.

Table 1: Elements of Anonymity

	Anonymous	New Pseudonym	Invested Pseudonym	Full Identity
Real Name				✓
Persistent Identity		✓	✓	✓
Invested Identity			✓	✓
Picture of Self		Unlikely	Discretionary	Discretionary
Personal Location		Unlikely	Discretionary	Discretionary
Personal Description		Unlikely	Discretionary	Discretionary
Contact Method		Unlikely	Discretionary	Discretionary
Posting History		Beginning	✓	✓
Opinion Accountability		Beginning	✓	✓
Demonstrated Expertise			✓	✓
Metric-Based Reputation		Beginning	✓	✓
Social Reputation			✓	✓
Interactions with Others	✓	✓	✓	✓
Social Bonds Forming			✓	✓
Popularity Growth			✓	✓
Feel Free to Share Any Opinions	✓	Potentially	Potentially	Unlikely
Avoid Social Prejudice & Stereotypes	✓	✓	✓	
Avoid Excessive Information Disclosure	✓	✓	✓	
Avoid Repercussions to Sense of Identity	✓	✓	Potentially	
Avoid All Repercussions for Behavior	✓			

It is worth noting that not all proposed elements are analyzed in this work. An effort was made to be inclusive, for use in other research that may follow on this topic.

A goal in developing this structure was to identity the importance of the "persistent" part of "persistent pseudonym." In other words, whether we type in some new string of letters for a screen name or simply go by "Anonymous" when we make some post, we are likely to have little difference in our demonstrated behavior. Without any history with a pseudonym or investment into it, a fresh and unique pseudonym is very close to simply using the word "Anonymous." So many elements associated with a developed persistent pseudonym, such as a posting history to which people may refer, social ties formed with other users, or a profile customized to share chosen bits of information about one's self, are unlikely to be present when a user first starts posting under some given name. The identity has not yet become invested with any personal meaning or sense of self, and so I would not yet expect to see many (or any) changes in behavior come from using it, relative to anonymity.

The literature and these observations suggests that positive online consumer interactions can be sorted into three main categories:

1. Breadth of Participation: how willing are people to participate in total content production, and in the variety or type of this content? Although not specifically related to online review writing, many of the discussed papers identity a greater willingness to participate online when

allowed some level of veil of anonymity (Dubrovsky, Kiesler, and Sethna 1991; Haraway 1990; Kiesler, Siegel, and McGuire 1984; Poster 1990; Siegel et al. 1986; Sproull and Kiesler 1991). Whether this applies to the number of writings, the volume of these writings, the variety of these writings, or some combination of the three, is unknown.

2. Quality of Participation: how greatly do users value the quality of content associated with their identity, and make efforts to improve them? In our work on online identity (Messinger et al. 2008; Messinger et al. 2009) we saw great effort taken to improve the quality of one's online visual identity, regardless of the level of personal identity disclosure. Does this apply to the output made under a chosen identity, as well? Furthermore, from the perspective of impression management, we cultivate idealized images of ourselves online (Donath 1999). The question is then whether this applies both to real identities and pseudonyms (that we have invested with some level of identity), and whether the drive to create an idealized self image online also applies to the creation of a high-quality body of content associated with that identity.

3. Pro-Social Behaviors: how much effort do people put into making positive connections with other users, or how positive a response do they garner from others? The literature has mixed predictions for whether FID will lead to superior (Kiesler et al. 1985; Orenga et al.

2000; Siegel et al. 1986; Sproull and Kiesler 1986) or inferior (Christopherson 2007; Hinduja 2008; Joinson 2001; Lee, Im, and Taylor 2008; Schul, Meyer, and Burnstein 2008) performance in many areas.

With these categories in mind, I suggest the following models to demonstrate the different performances of users under different levels of anonymity:

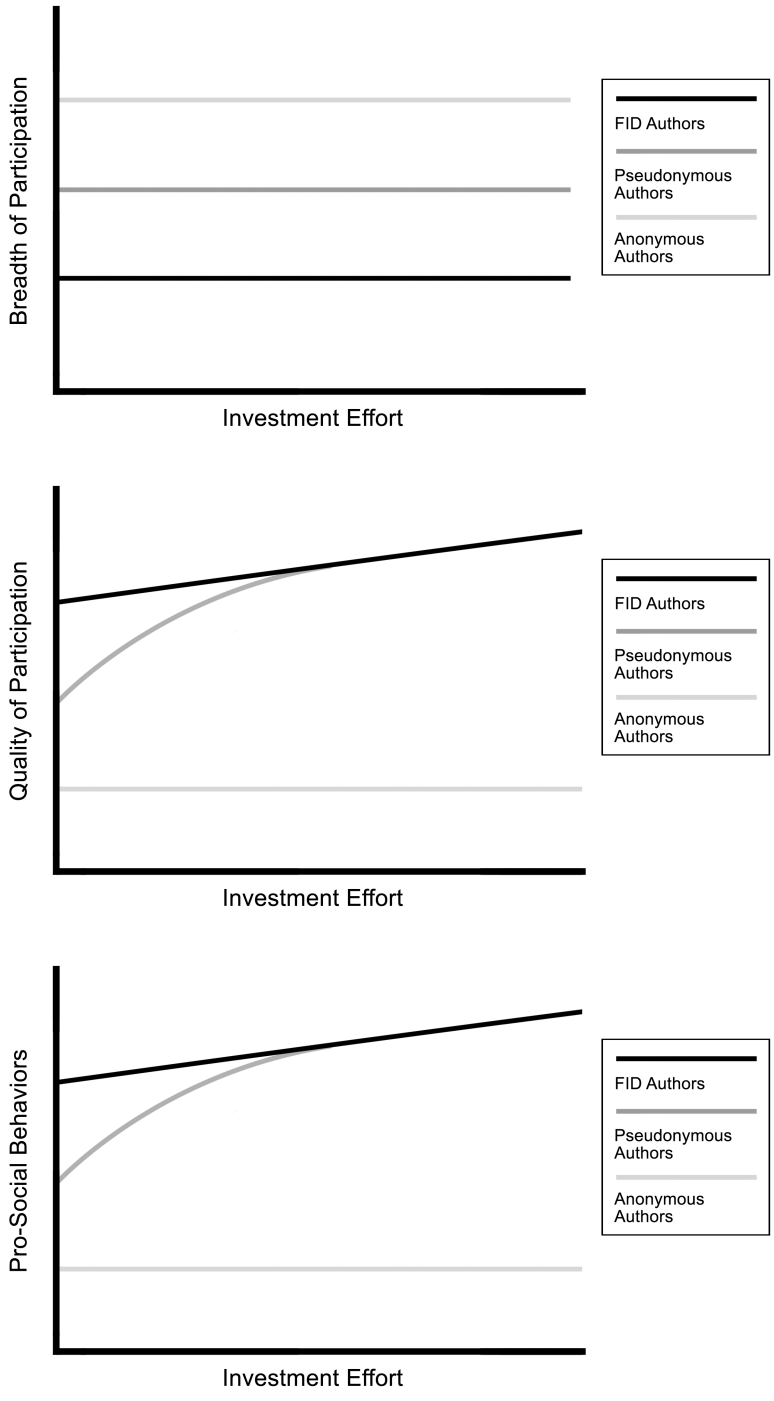

Figure 1: Proposed Models of Online Behavior Under Levels of Identity Revelation

As the reader can see, there are distinct differences between the expectation of behaviors related to the breadth of content production and to those related to pro-social behaviors and content quality. Given that people do prefer to maintain privacy, I expect certain baseline behavioral changes associated with different levels of identity disclosure. In other words, when we disclose more of our identity, we have a consistent negative reaction that impedes our willingness to write. However, for both content quality and pro-social behaviors and connections, anonymous users will perform the worst, while those using their real identities will begin with the best performances. There, the investment in one's identity will create a greater feeling of accountability, and therefore improved performance. Pseudonymous users using a new pseudonymous identity will not feel this strong sense of accountability to an invested identity, and so will not originally perform at a level approaching real name users. However, after a relatively limited time of investing themselves into this identity, through time spent with the identity and/or effort generated while using it, their performances in these areas will reach equity with real name users.

When one is anonymous, there is no posting history to which to refer, nor any pressure to behave positively toward others or easy way to form bonds with them. As well, under that anonymous identity, there is no reminder that, as a certain person who is known for being an expert on the history of the Red Sox franchise or Italian opera, all future postings on that topic should live up to the quality expected from such an individual. However, someone working under a

full identity will almost certainly feel less likely to participate in discussions with the freedom of someone operating anonymously; if we are using our real names, able to be Googled by people in our real lives, are we likely to discuss how we overcame the temptation to cheat on a spouse or the performance quality of a laxative tablet? In addition, although an individual may certainly become an expert in an area, it is effectively impossible for a whole identity to demonstrate a complete and total focus upon some topic area, as is possible for some chosen identity; in other words, even an expert's focused expertise may not be as great as if that expert was posting under an identity focused exclusively upon this topic.

This framework also suggests that invested pseudonyms may lead to improvements in performance over both anonymity and full identities, across measures of participation, quality and expertise, and social connection and reward. Given the steepness of the curves of the expected improvements by pseudonymous users, relative to real name users, it is possible that they will outperform real name users after this period of investment. Across multiple studies of both real world data and laboratory experiments, I explore these questions in depth through both laboratory and real world data analysis. Through them, I wish to demonstrate the two main features described by this model: (1) that pseudonyms allow for these performance improvements, and (2) that it is necessary to invest one's self into this chosen pseudonym before these performance improvements occur.

Pseudonymous Behavior in a Real-World Setting

Through a multi-stage process of analyzing real-world data, I identify behaviors occurring on a popular site (Amazon.com) as affected by levels of anonymity, and identify specific behaviors that will then be isolated and tested within controlled laboratory settings. In order to explore the magnitude and frequency of these variables in the real world, a highly externally valid correlational methodology is employed in several studies. Internal validity and direction of causality are further enhanced in subsequent controlled experiments.

Purpose

At this point I make use of consumer-generated content, later supplemented by laboratory testing. This initial focus is taken for several reasons. First, there is a wealth of such data available for analysis, and it remains inadequately examined in consumer research. Second, consumer-generated content is a way for online businesses to effectively outsource part of their business expenses to consumers, who willingly participate in generating content for them (Arakji and Lang 2007). Satisfying consumers' social and personal needs (Ren, Kraut, and Kiesier 2007) can be a strong economic driver for an online business (Armstrong and Hagel 1996). This is increasingly true when part of an online retailer's value offering is that it has many reviews for any given product, whose authorship by regular users can never be duplicated by company copywriters. User reviews of a product can have a marked impact on how that product performs (Chintagunta, Gopinath, and Venkataraman 2010).

For this dissertation, I chose product review data from Amazon.com. A discussion of the site's content is in the next section. Due to the specifics of Amazon's data structure, some technical descriptions of this data set are necessary. The key analyzed features of the users and reviews across the studies are the (1) levels of authorial anonymity, which have already been theoretically discussed, (2) the amount of effort expended upon their creation, (3) the amount of response they get from other users, and (4) the level of quality that they achieve.

Offering high-quality online content is an important goal for firms, as it boosts ties to a virtual space and improves opinions of the host site (Porter and Donthu 2008). However, as noted above, users must feel the right way about a site before they are willing to work for no compensation (Armstrong and Hagel 1996; Ren et al. 2007), and so understanding what might affect the quality of contributions is important. Effort in content generation matters as well. Across several studies, I will use review length as a proxy for effort. Longer length can contribute to higher review quality if more relevant features and benefits are addressed in greater depth. Also, the sheer amount of review information available can be more influential than its valence (Chintagunta et al. 2010). For these reasons, I include measures of review length as well as review quality across several studies.

This first study focuses on authors who, through assessment by other users on Amazon.com, have been identified as top performers for user-generated content. This is a short exploration, which serves to show the existence of top-performing authors who choose to write under a pseudonym across multiple

segments of the user population. As such users have gone unexplored in the literature, their identified presence then justifies the subsequent work both outside and inside of the lab.

Amazon.com Data

The United States version of Amazon was used for two primary reasons: one, it has a wealth of user reviews compared to other national versions of the site, and two, it has a far wider variety of product offerings compared to the Amazon sites in other nations. As well, Amazon offers huge numbers of products within each category, from top sellers to those which may only sell a handful of units in a year (Anderson 2006), effectively removing any seller-side constraint on which products and their reviews are encountered for analysis.

To leave a review on an account, it must first have at least one item purchased and delivered with it; this may be a physical delivery or an electronic download. In the author's experience, even with electronic downloads there is often a 48-hour activation period before an account with a delivered purchase can actually leave reviews. This both slows fresh account creation and attaches a monetary penalty to attempting to create many new accounts for the sake of leaving reviews. Although likely intended to prevent manufacturers or authors from flooding products with positive reviews, from the perspective of the individual consumer, this has the side effect of making people more accountable to their previous words, whether under their real identity or a chosen pseudonym. If consumers wish to leave behind their old words and start fresh, it will cost them both time and money, quite literally.

Upon beginning one's first review, Amazon asks the user how they would like to be represented. The default is the name attached to their credit card, verified through their existing purchases. This may appear as an abbreviated

version, such as "J. Tomilson," "Emily," or "Alex P.," or it may be a person's full first and last name, with or without a middle initial. They may also click a prominent link to use a chosen "pen name" instead of their real name, such as "Cascades Climber" or "jakesmom." The extent of the identity revelation is entirely up to the user's discretion, with the limitation that any chosen name may not be abusive or profane. If they use the default of the real name on that account's credit card, they receive a graphical Real Name badge on their profile page and are identified as using their actual identity upon making any review. Although this could result in selection bias, the external validity afforded by exploring a rich data set in this important substantive domain diminishes this shortcoming. Specifically, the possible selection bias may be outweighed by the benefits of Amazon's structure: (1) a large and rich body of data is available on the site, and (2) Amazon offers real name verification tied to a credit card, eliminating potential error associated with coders looking at names to decide whether they "look real." To further limit the impact of any potential selection bias, this correlational real world data is supplemented with experimental work where reviewers are randomly assigned to different identity conditions.

Studies

Study 1: Incidence of Pseudonymous Identification Among All-Time Elite Reviewers

Purpose

This study analyzes the types of users engaging in a particular, featured style of reviewing on Amazon.com's system, and identifies trends in their level of self-selected identity revelation.

Amazon's Hall of Fame system recognizes certain users who have achieved prominence within Amazon's user reviewers. Their site's algorithm generates an ongoing ranking of reviewers. Those who reach the top 10 ranking within these "Top Reviewers" are inducted into a permanent Hall of Fame. This Hall of Fame also recognizes users who were previously labeled as their top 100 reviewers under a previous sorting algorithm (Amazon.com 2011). A survey of the Hall of Fame reviewers reveals a list of users who are extremely prolific; one notable user currently has over 27,000 unique product reviews, although reviews in the thousands or hundreds are more typical of this group (Amazon.com 2012a).

It is important to understand the general purpose of this selection: the reviewers who reach the top of Amazon's sorting algorithm do so by demonstrating several behaviors valuable for analysis in this project. First, they have a large body of work available to analyze, and secondly, they are, by definition, currently writing high-quality reviews. This will show not only that

pseudonymous users can exist at this quality level, but will also serve as a contrast against other performance levels, to be discussed in subsequent sections.

For analysis, a key feature for obtaining the Amazon data is recording which "badges" people have on their profile. These small graphics note specific features of a profile's content, such as personal identity, or note recognitions granted to users by Amazon. There are a number of badges that may be earned by reviewers on the site. Along with generally being noted as a Hall of Fame reviewer, users are also granted entrance for specific years. For example, a user may be recently inducted in 2012, they may have been a Hall of Fame reviewer since several years back, or they may be an intermittent user who is there in some years and not others. Within their individual profile page, the years of their inductions are clearly listed. They may also be listed as a Top Reviewer, which, unlike the static Hall of Fame page, is a constantly shifting listing of the top reviewers of that moment, rather than the top historical reviewers chosen from Amazon's entire user population. For a comparison, the Hall of Fame page currently has fewer than two hundred featured reviewers, while the Top Reviewer listing displays the current top ten thousand reviewers. Top Reviewers may earn a general Top Reviewer badge, a Top 50 Reviewer badge, a Top 10 Reviewer badge, or a #1 Hall of Fame Reviewer badge.

As well, users may have a Real Name badge. As mentioned, the presence of this badge indicates that they have matched their public identity on Amazon with a credit card associated with their account. This is not an entirely inclusive method through which to determine who is using their real name online, as it does

require that optional step of verification and people may (for whatever reason) list their real name as their pen name, and so forgo the Real Name badge even while using their real identity. However, this step is not a difficult one, for one's credit card name is displayed as the default identity, and choosing to misrepresent one's real name as a pseudonym would require visiting an additional page instead of accepting it when first presented. As well, it is currently the only wholly reliable method of saying that, without question, someone is using his or her true identity online. A casual reader has no way of knowing whether a seemingly real name is truly that person's identity, or is simply a chosen pseudonym that happens to look like a typical North American name but is unrelated to their true self. Beyond this, the Real Name appears next to any review made from that person, and so it is a statement of verifiable identity that tracks the person throughout their participation in the site. It is not only verifiable, but it is also prominent, and it was used as the sole measure of real identity disclosure because of this. Put shortly, those with this badge truly own the disclosure of their real names (Amazon.com 2012d).

In this first study, I analyze a basic level of the Hall of Fame reviewers to determine how many of them have disclosed their real name. This provides the first steps to understanding how identity disclosure may lead to investment in a site and higher quality contributions to it.

Procedure

Within the Hall of Fame reviewers page (Amazon.com 2012a), data was

scraped to determine how many of the 123 honored reviewers (as of this writing) have verified their real names with the Amazon system. Then, they were analyzed further on three elements of their behavior on the site: how many reviews they have made, how many "useful votes" they have received from other users, and how many votes per review they receive, on average. Upon reading any review on the site, a user has the chance to label it as helpful to them, designating that they believe it to have some inherent quality and/or value. All such data was recorded.

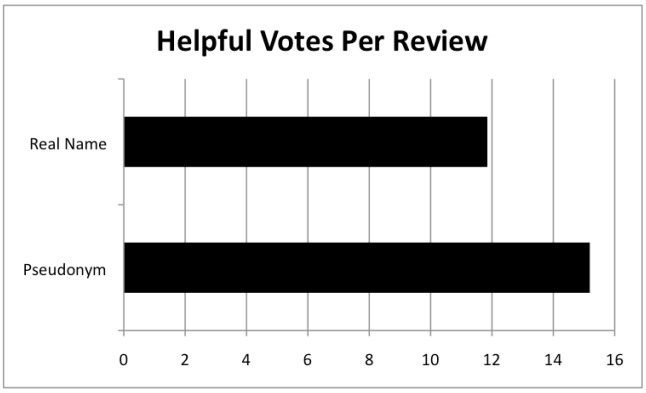

Figure 2: Helpful Votes for Hall of Fame Reviewers

Results

Of the 123 Hall of Fame reviewers, 46 (39%) are not using Amazon's Real Name system. While this does establish a majority of elite reviewers who have chosen to reveal their full names as part of their high-level community participation, there are still more than a third of these users who have achieved this level of presumed quality and community prominence not with their true identities, but with an invested pseudonym. As well, there exist both real name and pseudonymous users in the group of Hall of Fame reviewers with the longest record: 11 individual years of recognition.

Furthermore, the two groups have no differences in the scope or quality of their output. There is no difference in the number of reviews the two groups have

written ($M_{pseudonym}$ = 2282.93, $M_{realname}$ = 2499.49; F(1, 119) = 0.176, p = .675), in the total "helpful" votes from other users ($M_{pseudonym}$ = 21,847.28, $M_{realname}$ = 22,248.47; F(1, 119) = 0.019, p = .890), in the average number of helpful votes per review ($M_{pseudonym}$ = 15.19, $M_{realname}$ = 11.85; F(1, 119) = 2.772, p = .10), or in the average percentage of helpful votes ($M_{pseudonym}$ = 86.67, $M_{realname}$ = 87.19; F(1, 119) = .104, p = .748). Indeed, the only result that even approaches a significant difference between the groups is that of each review from the pseudonymous group, not the real name group, being judged as more individually helpful to other Amazon users.

Results

These results, although inconclusive, provide an initial challenge to the notion that users who reveal their true names behave "better" in both the scope and quality of their reviews. With the only near-significant difference being in favor of pseudonymous users, it is impossible to say that they outperform FID users; however, it is equally impossible to say that FID users are superior. Given the assumption in much of the literature for the latter to be so, this establishes the need for further analysis.

Study 2: Length of Review Content by All-Time Elite Reviewers

Purpose

With a deeper scraping of the data that collects every user's available review history, individual reviews can be recorded and analyzed, rather than using aggregate statistics. As individual reviews are treated as separate data points, this

also allows for the isolation of uniquely popular reviews that might be driving up the number of total votes for some users; if they had written a featured review for the last *The Hunger Games* book, for example. Through both greater statistical power and examination of individual review quality and content, a much deeper understanding of the users' behavior may be achieved.

Procedure

The full available body of reviews was taken for each user in the Hall of Fame by using customized software. Due to limits on Amazon's archival retrieval, this is not the full history of some users' reviews, but it does provide hundreds or thousands of output examples for each individual user (if they've written that much). Amazon's servers proved eventually unresponsive when many pages were scraped in an unbroken string. Some reviewers could not have their review history collected after three attempts, which was set as the maximum number of collection attempts. Given that the collection process lasted upward of eight hours for some individuals, server interruptions were expected. However, the vast majority of users were able to be collected, even with very large number of reviews occasionally needing to be retrieved. Between gathering the data for Study 1 and this study, four more users were inducted into the Hall of Fame. Of these 127 total users, 4 could not have their full review histories collected after three attempts, leaving 123 users' review histories available for analysis.

As well, there were approximately ten thousand reviews that were corrupted upon collection. The anticipated server interruptions for users themselves were instead mostly encountered on individual products' reviews.

There was no bias toward certain product classes, nor review ages, lengths, or valences being corrupted. Only one product (a certain model of Kindle e-reader) had more than two of its unusable reviews pulled in this population, or any subsequent studies. This resulted in 185,767 individual product reviews analyzed for this study.

Results

The size of the review base allowed for small but highly significant differences to be identified. Within this sample, pseudonymous users write more words per review than real name users ($M_{pseudonym}$ = 340.56, $M_{realname}$ = 335.64; F(1, 176,135) = 16.501, p < .001).

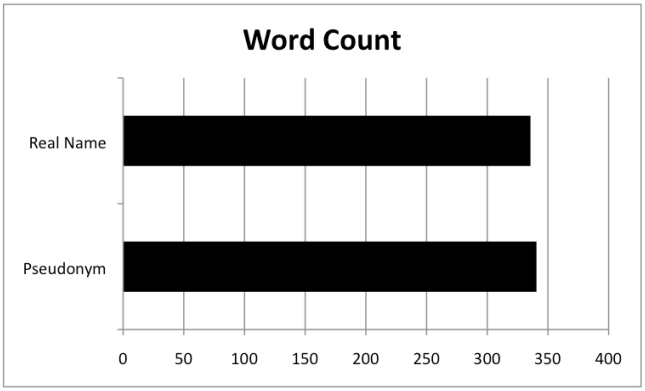

Figure 3: Hall of Fame Reviews Word Count

Discussion

Although the difference is small—five words across several hundred—it is also highly significant within this sample population. Here, we see a first indication that there may be a main effect of willingness to write more based on identity level. Greater length allows for more in-depth reactions and more details shared within a piece of review writing, though it does not guarantee it. This is

also, unlike Study 1, a clear indication of superior performance by pseudonymous users.

Study 3: User-Rated Content Quality by All-Time Elite Reviewers

Purpose

The previous study indicated the potential for higher quality reviews being produced by pseudonymous users through greater length. At this stage, numerous measures of individual review quality are used to explore the differences between pseudonymous and FID users in the most detail possible.

Procedure

The data collected for the Hall of Fame reviewers' full histories, as described in the second study, was analyzed here on quality measures. A large amount of information was available for each individual review beyond the previously examined word count. Just as each individual reviewer's profile contains the aggregate of all of the votes they have received, these votes are available for each individual product review. Further collected features unique to individual reviews include the number of comments, the product class and identifying information (such as name, ISBN, etc.), the creation date, the sales rank of the product, and the reviewer's scale assessment of the product's quality.

Results

Pseudonymous users received more votes per review than real name users ($M_{pseudonym}$ = 12.33, $M_{realname}$ = 10.86; F(1, 176,135) = 44.483, p < .001) and more helpful votes ($M_{pseudonym}$ = 10.89, $M_{realname}$ = 9.42; F(1, 176,135) = 50.983, p < .001). However, the number of negative votes resulted in a slightly higher positive percentage for real name users, not pseudonymous users ($M_{pseudonym}$ = 71.28, $M_{realname}$ = 72.23; F(1, 176,135) = 26.876, p < .001). Put differently, analyzed at this level, pseudonymous Hall of Fame users spark more engagement with other users, while other users judge reviews written under real names to be more useful.

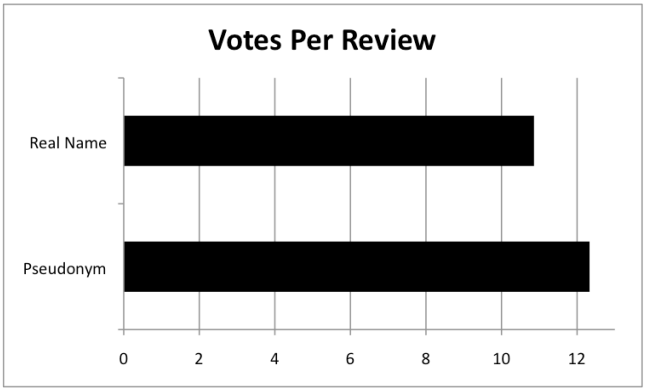

Figure 4: Total Votes per Hall of Fame Review

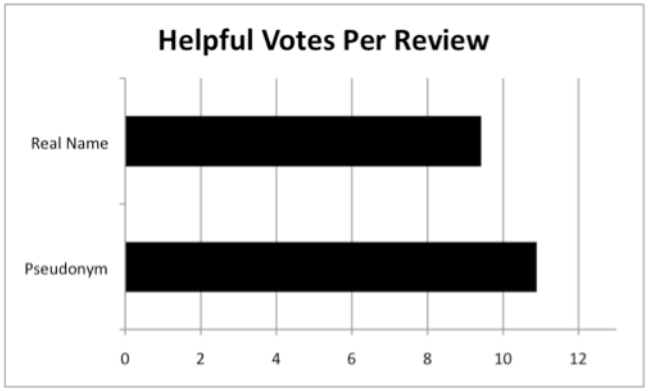

Figure 5: Helpful Votes per Hall of Fame Review

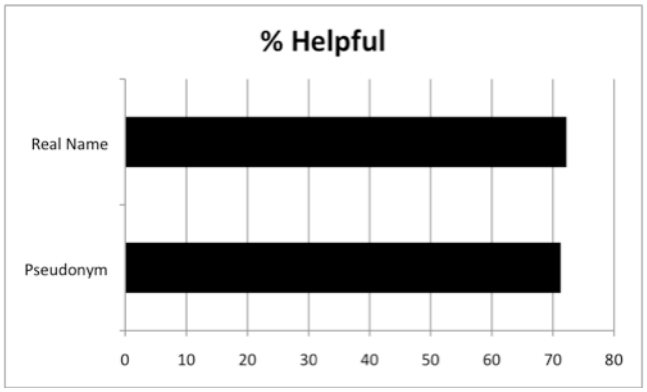

Figure 6: Helpfulness Ratio for Hall of Fame Reviews

Discussion

Within this all-time elite group, we see review-level findings that are more precise than the generally equitable performance seen in the original inspection of the Hall of Fame reviewers. Already, isolating unique reviews as their own data point, rather than rolling them blindly into a user's aggregate statistics, has allowed for newly identified behaviors.

The Hall of Fame users are an Amazon-desirable but highly unique group. Exploration of the current elite Top Reviewers and, particularly, the random population sample is an important next step.

Study 4: Incidence of Pseudonymous Identification Among Current Elite Reviewers

Purpose

Any site would likely welcome the Hall of Fame users as producers of content. However, we cannot expect that many customers are willing to write thousands of reviews for no compensation, even among high-quality authors. Here, a second, less established group of high performers has their reviews evaluated, to see if pseudonymous users are also well represented among a more accessible level of recognition.

Top Reviewers are the current top ten thousand highest-ranking reviewers within the Amazon system. Unlike the static Hall of Fame, the Top Reviewers list is a constantly changing snapshot of reviewer performance. As previously noted,

a reviewer must achieve top ten status within this group in order to be inducted into the permanent, static Hall of Fame. However, all users within this group are producing work that Amazon's algorithms recognize as being generally superior. (Amazon does not make its quality sorting algorithms transparent anywhere on its site, and did not respond to questions regarding these matters.) To expand the analysis of the prominence of non-real identity reviewers within high quality user content, in this study we conduct a similar analysis of the Top Reviewers.

Procedure

Study 1's procedure was duplicated here, but with the Top Reviewers (Amazon.com 2012b) segment of the user base rather than the Hall of Fame Reviewers (Amazon.com 2012a). One minor difference in the pages' presentation is the presence of "Fan Voters." Amazon explains this group thusly:

> Fan voters are people who consistently appreciate the author's reviews. These votes are not reflected in the total vote count to provide our customers with the most unbiased and accurate information possible (Amazon.com 2012c).

The scraped data thus does not include these users, who are seen by Amazon's system as insufficiently representative of their general population. These votes are relatively insignificant; a snapshot from the first page includes listings where 30 fan votes were removed to leave a pool of 37,873 votes, 574 to leave 46,385, and 14 to leave 10,493. In the first fifty reviewers, the highest percentage of removed Fan Votes was 2% of the remaining votes (271/12,958) (Amazon.com 2012c). I

am thus treating this group as negligible, and am ignoring them in the analysis, as does Amazon.

The problems previously encountered while collecting full histories for Hall of Fame users were encountered during basic profile collection here, given the number of pages visited. Given that it is impossible to break up Top Reviewer scraping across several days to overcome this technical challenge, as the list is constantly changing, the top 3,000 users of the 10,000 available had only their user IDs collected in a single scraping. With this small amount of identifying data that did not overwhelm the Amazon servers, their individual profile information was safely recorded across a broader timespan. At the time of the collection, 106 Hall of Fame users were represented in the top 3,000 reviewers. These were removed from the file to avoid duplicate analysis, leaving 2,894 users. As well, there were minor server interruptions that prevented some individual variables from being properly read for specific profiles during the time

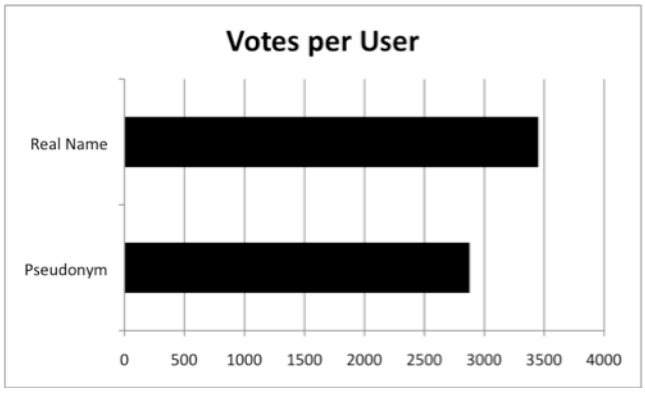

Figure 7: Votes for Top Reviewers

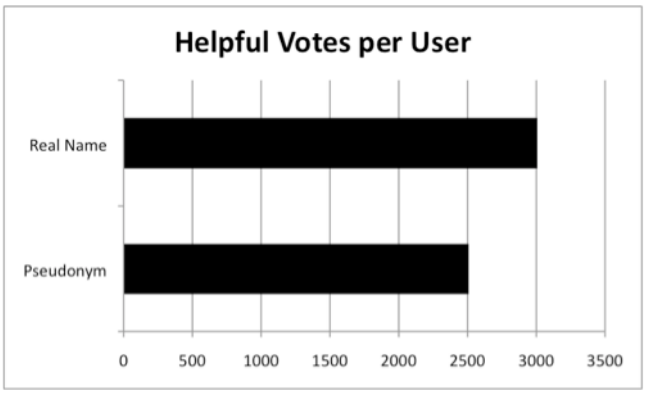

Figure 8: Helpful Votes for Top Reviewers

of access. These were inconsistent and very rare: less than five instances for any given variable.

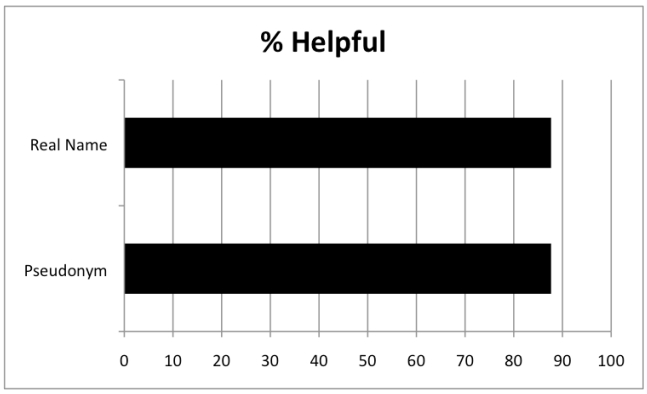

Figure 9: Helpfulness Ratio for Top Reviewers

Results

Unlike Hall of Fame users, this Top Reviewer subset has more pseudonymous users (1631) than real name users (1263). They have written far fewer reviews on average than Hall of Fame users (391), although this is of course still a sizeable investment of effort by consumers. No significant differences between real name and pseudonymous users exist on the number of reviews written.

Here, the earlier equitable situation seen with the Hall of Fame users in terms of gross quality measures sees an intriguing change. Although there are significant differences in the number of overall votes received ($M_{pseudonym}$ = 2,877.29, $M_{realname}$ = 3,451.41; $F(1, 2891) = 22.845$, $p < .001$) and the number of helpful votes ($M_{pseudonym}$ = 2,508.21, $M_{realname}$ = 3,006.81; $F(1, 2891) = 22.542$, $p < .001$), the helpfulness percentage between the two reviewer types is identical ($M_{pseudonym}$ = 87.67, $M_{realname}$ = 87.67; $F(1, 2891) = .001$, $p = .982$). Even though these pseudonymous reviewers are perhaps overlooked because of their lack of FID, they are writing identically helpful reviews as rated by this smaller population of readers.

Study 5: Length of Review Content by Current Elite Reviewers

Purpose

The earlier length analysis performed on the Hall of Fame reviewers is duplicated here. The Top Reviewer Group remains of interest as a more accessible population to which sites might hope their users aspire: producing high-quality content for the site, and a significant volume of it, even without a string of reviews spanning thousands of products and years of writing.

Procedure

The procedure for gathering review-level data was repeated for the Top Reviewer sample population. As with the Hall of Fame population, there were errors in gathering data at both the user and review levels, which were again centered more upon individual product review errors rather than user-level errors. Ultimately, 304,106 reviews were collected for analysis.

Results

Here, the clean results seen in the Hall of Fame reviewer population disappear. Although a directional difference remains, there is no

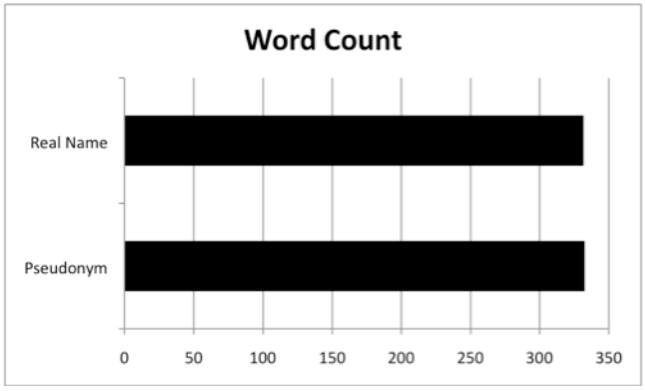

Figure 10: Top Reviewers Word Count

statistical difference in review length between pseudonymous and FID users ($M_{pseudonym}$ = 332.44, $M_{realname}$ = 331.66; F(1, 304,104) = .673, p < .412).

Discussion

Given the prominence and high number of reviews by both the Hall of Fame and Top Reviewer sample populations, it would seem logical to expect similar behaviors between them. However, this was not so. The less prolific (but still high quality) Top Reviewers see no statistical difference between groups even with tremendous statistical power. Although lacking the more intriguing result of superior performance by pseudonymous users, it nonetheless supports the argument that FID is not inherently superior.

Study 6: User-Rated Content Quality by Current Elite Reviewers

Procedure

The data collected for the Top Reviewers' full histories was analyzed here on quality measures, as performed with the Hall of Fame reviewers.

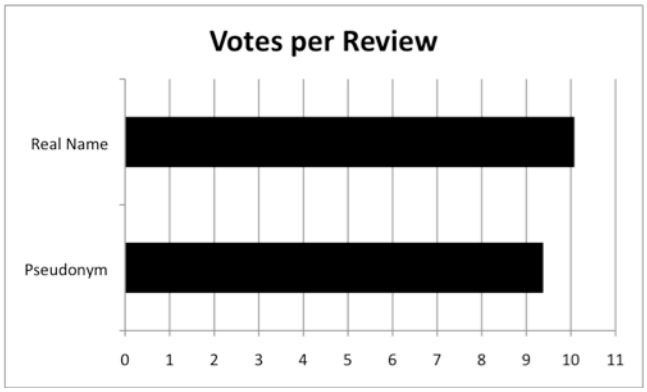

Figure 11: Total Votes per Top Reviewer Review

Results

Here, the earlier balanced situation seen with the Hall of Fame users in terms of gross quality measures suffers.

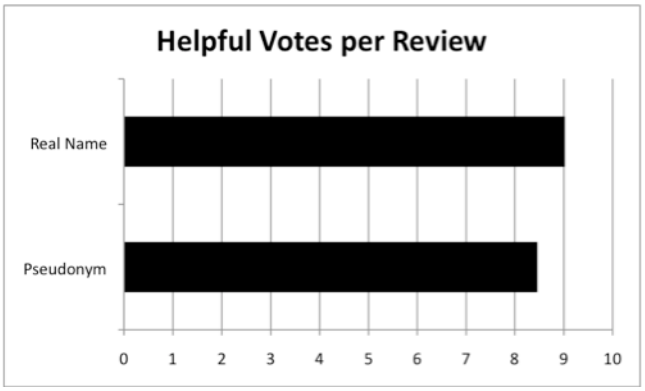

Figure 12: Helpful Votes per Top Reviewer Review

There are significant differences in the number of votes received per review ($M_{pseudonym} = 9.38$, $M_{realname} = 10.08$; $F(1, 304{,}104) = 13.201$, $p < .000$) and the number of helpful votes ($M_{pseudonym} = 8.46$, $M_{realname} = 9.02$; $F(1, 304{,}104) = 9.336$, $p < .002$), and the helpfulness percentage ($M_{pseudonym} = 64.85$, $M_{realname} = 68.94$; $F(1, 304{,}104) = 743.374$, $p < .000$).

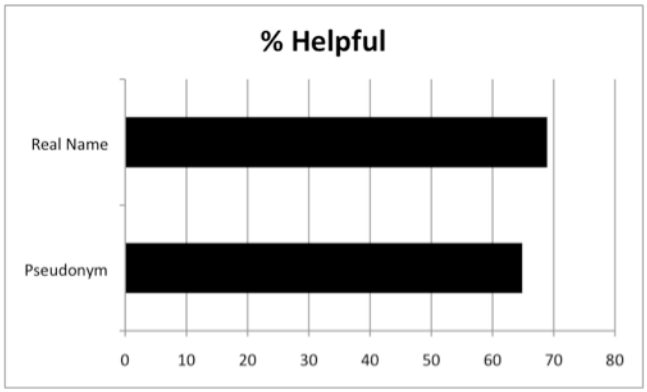

Figure 13: Helpfulness Ratio for Top Reviewer Reviews

Discussion

The Top Reviewer population is again less kind to pseudonymous users than was the Hall of Fame sample. At this stage, I move to the most representative sample possible, that of a random draw from the entire reviewer population. Afterward, I dig deeper into this sample to further explore review behaviors in what is the most accurate sample of the population, without any of the performance bias present with the high-quality population samples.

Study 7: Incidence of Pseudonymous Identification Among The General Population

Purpose

The previous studies demonstrate that pseudonymous users can match the output of real name users on several levels, although this distinction varies by length of time spent on the site. For the most highly invested pseudonymous users—members of the Hall of Fame—there are many areas in which they outperform real name users. For those with shorter histories on the site, however, pseudonymous users do not match the performers of real name users when review-level analysis is performed, although these differences vanish at the aggregate user level. For sites interested in high-quality user content, these results show that pseudonyms can be a useful method for generating content that is viewed quite positively by other users of their sites, although long-term use appears key. This finding applies particularly when the greater preference toward

pseudonymous representation in the Top Reviewer population is considered, as this is the source from which these elite users are recruited.

However, this gives little insight into the typical Amazon user. Even Top Reviewers have written nearly 400 reviews, on average. For both the time spent on writing and money spent on purchases, this seems like an unreasonable expectation to make of a representative user. The Hall of Fame and Top Reviewers listings are convenient from a data collection standpoint, as Amazon provides public lists of their user IDs. However, it is unknown how they compare to the general population of Amazon users, as there is no comparable listing of IDs for the general user base. To gain a better understanding of typical site users, a random sampling was taken of all Amazon reviewers and analyzed on the quality measures used for the two previous groups.

Procedure

All Amazon user profiles have the same URL structure: a leading address followed by an eleven-character alphanumeric identifier that is unique to each user. With this consistent structure in mind, a three-stage process was coded for generating a completely random sample of the Amazon reviewer population. A single script was written to handle all parts of this collection, and so the completed list was generated blind to researcher input. First, a random eleven-character string was created. Second, this string was appended to the leading URL and visited; if only an error page was found, it was rejected. Third, the presence of at least one review on the page was verified, as a secondary check that it was a

valid user profile. Character strings that passed both checks were added to the output file. As there was no static page to easily draw IDs from and each successfully identified profile involved many more rejections, this required a high amount of Amazon server use relative to Top Reviewer collection. 1,500 valid IDs were collected for the random population. 38 users had written reviews but not yet received any votes for them, and so were analyzed with zeroes in the fields for total votes and helpful votes. No data was recorded for the helpfulness ratio, as a zero percent entry there would have skewed results.

Results

Despite drawing from a representative sample of the Amazon population, the findings from the Top Reviewers have been duplicated. Pseudonymous users were again more common (936) than real name users (564). And, as before, no significant difference exists between the number of reviews written by pseudonymous and real name users. As to be expected, there was a smaller average number of reviews written compared to the other sample populations (139). The median was also examined to identify if there were outliers driving this high number, as unlike the other two groups, there is no required history of strong performance to be counted here. Indeed, this was the case. The median number of reviews written for the random sample was only 11, which has face validity as the number of reviews one might find for the typical Amazon user who cares to leave any reviews at all.

There were again significant differences in the number of overall votes received ($M_{pseudonym}$ = 1,156.30, $M_{realname}$ = 1,838.76; $F(1, 1498) = 5.481$, $p < .019$) and the number of helpful votes ($M_{pseudonym}$ = 982.01, $M_{realname}$ = 1,574.34; $F(1, 1498) = 5.407$, $p < .02$). However, once again, the helpfulness percentage between the two reviewer types has no significant difference, and is in fact directionally favorable toward pseudonymous users ($M_{pseudonym}$ = 85.06, $M_{realname}$ = 84.60; $F(1, 1460) = .310$, $p = .578$).

Study 8: Length of Review Content by the General Population

Purpose

Do pseudonymous users write more or not? Is the hypothesis detailed in the full dissertation introduction, that supposes that increased performance comes only after some certain and significant investment of time, accurate? The current results seem to support this: Hall of Fame users (with a longer review history) wrote slightly more than Top Reviewers, and there it was users with those highly invested pseudonyms who out-performed FID users. At a lower level of involvement in the site, pseudonymous users maintain a directional advantage in terms of writing length, but the statistical difference has disappeared. At a cursory level, the idea of greater writing volume for highly invested pseudonyms appears to be the story told. However, it is not a terribly applicable one. Across an unknown total number of reviewers on Amazon, only a handful make it to the Hall of Fame. Do pseudonymous users outperform FID users only when they have written thousands of reviews?

Here, I look to the random reviewer data. Of all three sample populations, this is the most representative and easy to generalize to online users in general. These are not users who have any minimum presence on the site; indeed, there are many who have only ever written a single review, and that one was not necessarily received well. The results from the random reviewers' posts are therefore most likely to tell the true story of how typical online users generally behave under varying identity levels.

Procedure

Review-level data was collected for all random user IDs, replicating the collection method for the two previous studies. After removing corrupted entries, 192,693 reviews remained for analysis.

Results

Within the randomly sampled population, pseudonymous users write significantly more than those using their real names ($M_{pseudonym}$ = 321.81, $M_{realname}$ = 304.81; $F(1, 192{,}691) = 178.954$, $p < .000$).

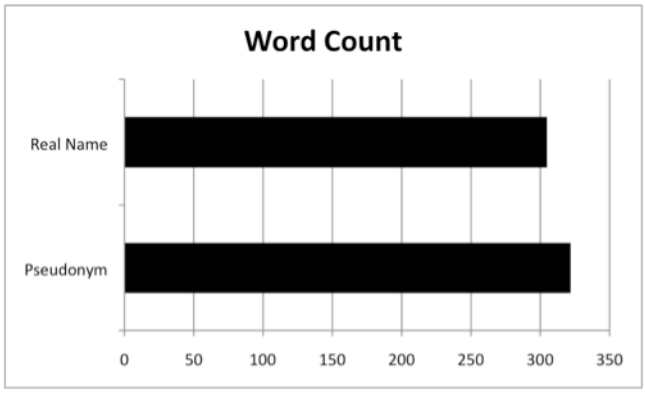

Figure 14: Random Reviewer Word Count

Discussion

Pseudonymous users again demonstrate a greater willingness to write more content in their reviews. This randomized, representative sample of users and their reviews shows a larger difference than either of the other two populations. The results from the scraped Amazon user data suggest one thing very clearly: across two out of three sample populations, including the most representative one, there is a main effect of increased willingness to write more for pseudonymous users. Even the population that did not repeat these results with

statistical significance repeated it directionally. A claim can be made that, with this real world data, identity affects willingness to produce content.

Returning to the proposed model of online behavior discussed in the literature section of this dissertation, there appears to be general, if not uniform, support for the hypothesis that using a non-FID identity leads to greater input. This is a new finding in the literature, which had previously considered only total anonymity.

Study 9: User-Rated Content Quality by The General Population

Purpose

This random population sample is now analyzed on more specific quality measures. Previously, it has been shown that Hall of Fame pseudonymous users outperform FID users in many quality measures, and also write more than them. In a less invested group of users, pseudonymous behavior suffers relative to the Hall of Fame group. However, in this most representative sample with a very small median number of reviews, pseudonymous users saw the largest word count advantage over FID users in any of the populations. Quality measures now follow up this word count finding to further explore the behavior of the most typical, least invested users examined in this collected data.

Procedure

The data collected for the random reviewers' full histories was analyzed here on quality measures, in the same steps seen the other populations.

Results

Despite the aggregate statistics being kind to pseudonymous users before, the review-level data also shows superior performance by real name users. There are significant differences in the number of votes received per review ($M_{pseudonym} = 8.82$, $M_{realname} = 11.63$; $F(1, 192{,}691) = 195.571$, $p < .000$) and the number of helpful votes ($M_{pseudonym} = 7.54$, $M_{realname} = 9.97$; $F(1, 192{,}691) = 165.407$, $p < .000$), and the helpfulness percentage ($M_{pseudonym} = 63.16$, $M_{realname} = 67.70$; $F(1, 192{,}691) = 585.202$, $p < .000$).

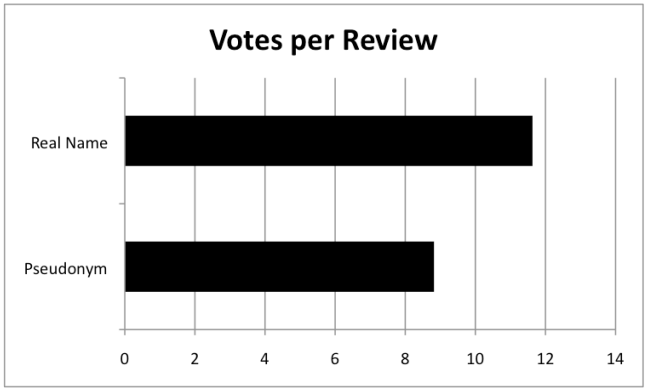

Figure 15: Votes per Random Reviewer Review

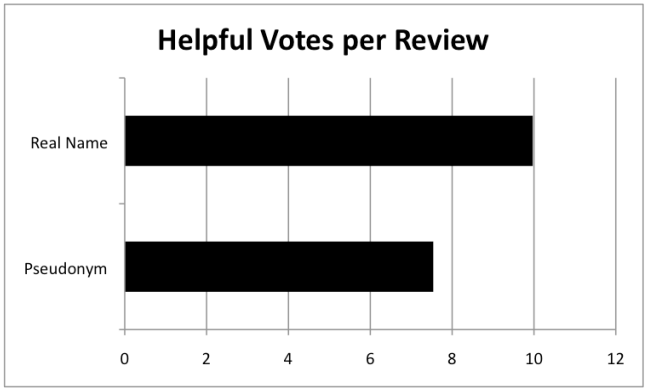

Figure 16: Helpful Votes per Random Reviewer Review

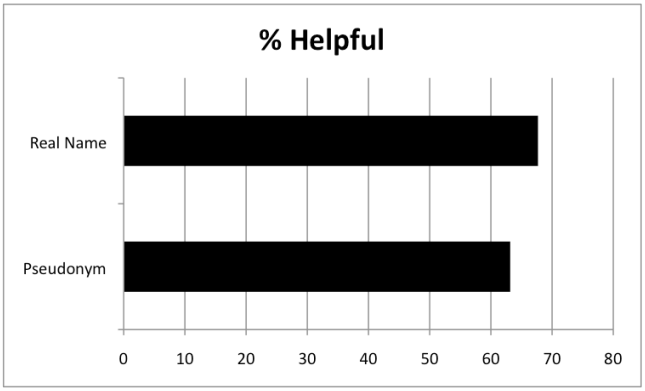

Figure 17: Helpfulness Ratio for Random Reviewer Reviews

Discussion

Here, the random reviewer data has duplicated the Top Reviewers. With greater statistical power, there appears to be support for real name users outperforming pseudonymous users at all but the top levels of identity investment on quality measures, although volume remains in pseudonymous users' favor. This supports the proposed model where there are no incentives for FID users to improve the amount of their participation, and pseudonymous users only improve their performance on quality measures when investment occurs. However, the investment needed for this effect, going by current findings, is quite high.

However, only a few measures of review quality have been taken until this point. There are many other aspects of the reviews that have, as yet, gone unmeasured. Subsequent studies examine these measures in detail, focusing upon the random sample, to identify any other differences in behavior between the groups that might impact other users' ratings.

Study 10: Quality Assessments as Affected by Review Content

Purpose

Up to this point, quality has been measured by straightforward, other-measured votes of a review's quality with an eye only towards that review's level of user identification. However, there are other review features to consider: valence and product category.

Do users react more or less positively to reviews that are harsh or kind toward the product? We may assume that critical reviews are seen as more useful, but is that so? Furthermore, Amazon sells an enormous range of products, and

these products are presented on the site in very different ways. Originally, Amazon sold only media products, such as CDs, tapes and DVDs, and books. They have since added other product categories, but the differences between these categories' sizes are stark. Within a smaller category, it is possible to take an optimizing approach toward product selection; therefore, a more critical review that removes items from consideration may be seen as more valuable. On the other hand, no one could ever hope to try all products within some of Amazon's larger categories; they are likely the closest example that exists anywhere to a complete set of product offerings. In such a circumstance, optimizing is impossible, and a consumer must take a simple satisficing approach. In this case, negative reviews should be seen as less useful, as the goal is to find a product that satisfies an acceptability threshold rather than removing those that do not.

Procedure

The star rating (1-5, poor to excellent) given to each product was already collected in the scraped data. For the product categories, the decision was made to split the dataset into two broad areas: "media" and "non-media" products. Media products are the aforementioned core of Amazon's historical product offerings: music, books, and movies. This was done for two reasons. One, the scraped product category data was not always consistent in specificity. For one scraped product, Amazon returned the very large category of "DVD." For the product directly under it in the database, it returned "Wireless Phone Accessory," a much more specific identifier. This is likely due to Amazon's organizational structure,

where any product can only be narrowed down by four categories (for example: Books > Business & Investing > Economics > Econometrics) before no further category links appear. (To demonstrate how large these media categories can be, the book category used as an example returned over 570,000 results on econometrics.) For smaller categories, narrowing it down one or two levels returns quite specific results; for larger categories, the results are still very broad. Therefore, broader, fewer classifications were used for analysis across all products.

As for the second reason, the sales ranks for the respective groups were tested after they were coded. The higher the average sales rank in a category, the more items are contained within it. Under Amazon's structure, the top-selling item in any category is "1," the second is "2," and so on through the total number of items in that category. Therefore, a category with an average sales rank of 200,000 would have many more items within it than one with an average sales rank of 20,000. The difference in average rank between the media and non-media categories was stark: a mean of 124,940 for non-media products and 709,249 for media products. This indicates far larger category size in the media products, as anticipated, and strongly suggests that different aspects of reviews might be valued if satisficing vs. optimizing methods are being used. However, there also exists the possibility that more negative reviews are seen as more expert, and would be more valued on smaller, less frequently encountered products relative to familiar purchases such as movies and books.

With this done, reviews were analyzed on several measures. As discussed, the ratings given to products by reviewers were considered. As well, previous measures of performance—word count and the percentage of helpful votes—were also used. Finally, there were some products whose product category could not be retrieved at all. These were removed from the analysis, leaving 185,904 reviews.

Results

Table 2: Product Rating by Identity and Product Class

Product Type	Author Label	Star Rating
Media	Pseudonymous	4.24
Media	Real	4.27
Non-Media	Pseudonymous	4.18
Non-Media	Real	4.14

Table 3: Review Length by Identity and Product Class

Product Type	Author Label	Word Count
Media	Pseudonymous	349.67
Media	Real	336.15
Non-Media	Pseudonymous	266.36
Non-Media	Real	212.36

Table 4: Number of Helpful Votes by Identity and Product Class

Product Type	Author Label	Helpful Votes
Media	Pseudonymous	7.07
Media	Real	10.13
Non-Media	Pseudonymous	9.33
Non-Media	Real	9.37

First, product ratings were modeled on identity and product category. Although there is no main effect of identity on rating ($M_{pseudonym} = 4.23$, $M_{realname} = 4.25$; $F(1, 185{,}900) = .214$, $p < .643$), there is a main effect where media products are rated more highly than non-media products ($M_{media} = 4.25$, $M_{non\text{-}media} = 4.17$; $F(1, 185{,}900) = 224.695$, $p < .000$). A significant interaction ($F(1, 185{,}900) = 34.329$, $p < .000$) exists between identity and product type, as well, where pseudonymous users are more favorable toward non-media items but more critical toward media items. See Table 2 for these means.

A return to word count analysis also finds significant differences. As previously seen, pseudonymous users write significantly more than real name users in this random sample of the reviewer population. However, there is also a significant main effect of product category ($M_{media} = 343.62$, $M_{non\text{-}media} = 245.74$; $F(1, 185{,}900) = 3891.989$, $p < .000$) as well as a significant interaction between identity level and product category ($F(1, 185{,}900) = 148.675$, $p < .000$). Real name users, relative to pseudonymous users, have a larger reduction in review length for non-media items. See Table 3 for these means.

Table 5: Helpfulness Ratio by Identity and Product Class

Product Type	Author Label	% Helpful
Media	Pseudonymous	65.58
Media	Real	70.16
Non-Media	Pseudonymous	54.97
Non-Media	Real	56.70

Intriguingly, this occurs alongside a significant (($F(1, 185{,}900) = 33.391$, $p < .000$) interaction between identity and product type on the number of helpful

votes received; see Table 4 for these means. Along with a main effect of media reviews being judged as far more helpful than non-media reviews (M_{media} = 55.63, $M_{non\text{-}media}$ = 67.63; F(1, 185,900) = 2379.057, p < .000), real name users receive a greater boost from their identity when they are writing for media items than for non-media items. See Table 5 for these means.

The final check remains the incidence of review quantities between conditions. The difference is significant, and a higher percentage of FID users' reviews are about the more generously assessed media items ($M_{pseudonym}$ = .80, $M_{realname}$ = .84; F(1, 185,900) = 500.439, p < .000).

Discussion

Although there exist clear differences in assessed quality, as seen in previous studies, these differences assumed that all product reviews were treated equally. These results suggest that they may not be. It could be that pseudonymous and real name users have preferences toward the products they choose to review and some product classes are seen as benefitting less from accompanying product reviews, regardless of quality. On the other hand, these results could come from genuinely worse writing quality by pseudonymous users. We cannot know whether reviews of less commonly encountered non-media items are truly less helpful in any objective sense, or if users are simply more generous toward those users who choose to review the familiar media products that make up the bulk of Amazon's sales. Consider the sharp reduction in word count for real name users relative to pseudonymous users on non-media items.

Are the authors of these reviews truly writing superior content, or are they benefitting from other review features, such as the product categories chosen by each group or the simple revelation of one's identity?

At this time, correlational studies are set aside for laboratory experiments in which these elements may be individually manipulated and isolated. Through this, I seek to identify whether non-elite reviewers have their reviews assessed as superior and produce superior reviews based on objective elements of review quality, or whether they are benefitting only from the use of their real names.

Pseudonymous Behavior in an Experimental Setting

Study 11: Assessment of Review Quality as Affected by Manipulated Identity

Purpose

As previously discussed, we make many judgments about a person upon first seeing them, from talent to trustworthiness to general assessments of their likability (Landy and Sigall 1974; van 't Wout and Sanfey 2008; Zebrowitz et al. 1996). In particular, Landy and Sigall's famous study found that associating a highly attractive person with a piece of writing led to higher assessments of that writing's quality. In an environment where the only personal identifier is the level of identity disclosure, does the presence of a real name duplicate the presence of an attractive face? To twist their famous question of "is beauty talent," this study asks "is identity talent?"

Multiple correlational studies have shown that, except in the most elite and invested Hall of Fame condition, pseudonymous users are assessed as worse than FID users, even though they write lengthier reviews. There may also be an explanation reliant on the type of reviews being written, or it could be that reviewers are simply more critical about some products than others, and as shown, pseudonymous users have a tendency to write more in this area. This experiment seeks to identify the answer to this question through manipulation of the presented authorial identity to subject reviewers.

Procedure

A wide variety of products sold on Amazon.com were pretested for familiarity, usefulness of online product reviews, and the perceived embarrassment level associated with using the products. (The latter was included to account for findings that anonymous users may be perceived to be more forthcoming in some discussions (Eagly, Wood, and Chaiken 1978), although to my knowledge this has not been found to relate to this specific domain.) All pretested products had at least three unique products within that category with thirty or more user reviews. This was done to assure that there would be both high and low quality reviews for whatever product was selected, as real Amazon reviews were presented to subjects.

From this, nine products were selected that were rated as benefitting from online review reading, that pretesters had not commonly tried themselves and so would not be relying on their own memories instead of assessing the review's quality (this resulted in the loss of most entertainment media), and that spanned a range of embarrassment levels. For each selected product, one individual product listing with at least 30 user reviews was randomly chosen as the source of the reviews used in this study. These reviews were rank ordered by helpfulness votes, and the most and least helpful reviews were collected. One exception to this procedure occurred when the least helpful review for a certain product was not about the product itself in any way, but was instead a complaint about slow package delivery. Then, the least helpful review that was exclusively about the product was collected. The list of selected products may be seen in Table 6.

Table 6: Products Selected

Adult diapers
Contact lens saline
Fiber supplement wafers
Hummingbird feeder
Hunting video game
Jock itch powder
Laxative tablets
Scented hand lotion
Weed killer spray

Originally, the design was intended to measure three levels of identity: FID, pseudonymous users, and complete anonymity, which is not available on Amazon but we would present as an experimental condition here. However, concerns arose during pretesting. From the list of random Amazon reviewers gathered during the data scraping studies, 30 pseudonyms were randomly chosen and rated for likability and plausibility. (I.e., that this was a real reviewer's pseudonym and not one I had made up.) A database of name frequencies was collected from Social Security sources and used to generate 30 random "real" names of middling frequency in the American population, which were also rated on likability and plausibility. Selecting a real name was a simple procedure, as there were many names that were of medium likability and were seen as highly plausible; furthermore, there was little variation in responses. However, the randomly-selected pseudonyms had more variation in the pretesters' responses. Furthermore, many of the names (see Table 7) had some unique features related to interest or focus (such as "VideoCritic") that raised concern over whether they would be believable in a random assignment to many different product categories. Thus, only two identity levels were used in this experiment: "real" name (using the pretested name of "Michael McKelvie") and complete anonymity. Although this removes the chance to compare pseudonymous versus FID performance, as seen in previous studies, it does explore one question raised

during several stages of analysis: do other users rate reviews more highly exclusively based upon the use of a real name?

Table 7: Pretested Pseudonyms

PT Cruiser	Trapshooter
Cubist	Halicon72
Pumpkin Man	atmj
Penny Pinching	SkyeNoir
Polly	Musicfan
Wix	DukeOfEarl
Relytia	Aravyndra
Dad of Divas	Mountain Woman
Z Hayes	SystemStructure
reg	prisrob
Parka	b-gat
Sparky Jones	VideoCritic
novchyld	MotherLodeBeth
Compusurge	tvtv3
The Mad Scientist	BrieBrie
Grimmy	

900 subjects participated in this study on MTurk. 17 failed the attention check, leaving 883 for analysis. They were randomly assigned into condition by product, review quality (high/low), and author identity (anonymous/real name). Upon beginning the study, subjects were given brief instructions and then shown an image and basic information (via a screencapture of the Amazon.com page) about the product to which they had been assigned. They then saw a real review of that product, of either high or low quality. Above the review was a line identifying the author: either "an anonymous author" or "Michael McKelvie," with Amazon's Real Name badge next to his name. (For Michael, there was also an explanation in very small font directly under his name that the presence of that graphic meant that he was using his real name as verified by the credit card on his account. This was done to replicate the pop-up explanation of the badge that can be read any time the badge is encountered on Amazon.)

Upon reading the review, subjects made several assessments of its quality. They first rated it on the binary helpfulness scale used by Amazon, then rated its

helpfulness more precisely on a slider scale from 1 to 100. Next, they completed a commonly used and adapted scale of writing quality (Diederich 1974). Subjects here completed the original version (less an inapplicable question about handwriting quality), in which they filled out subscales related to broader compositional issues (ideas and organization) and to more specific grammatical issues (wording, punctuation usage, spelling, grammar usage, and general "flavor"). An attention check asking for a description of the product they had just seen reviewed completed the task.

Results

The binary yes/no helpfulness rating proved to have insignificant difference between any groups, due to an overwhelming subject tendency to rate a review as helpful even when they subsequently rated the review as well below 50 on the 1 to 100 helpfulness slider. Only 19% of reviews were rated as unhelpful. Due to both this and the greater granularity, the helpfulness slider was used for analysis. The Diederich scale was collapsed into its broader categories (a 2-point composition scale ($\alpha = .868$) and 5-point grammatical scale ($\alpha = .903$)), as intended, and those subscales were used.

These three measures were modeled on author identity, review quality, and product in a GLM. Unsurprisingly, higher-quality reviews had a strong positive effect on evaluations ($F(3, 845) = 151.719$, $p < .000$), but a review being presented under a real name also had a strongly significant positive effect ($F(3, 845) = 3.910$, $p < .009$). Review quality had a significant effect on all individual

quality measures: helpfulness ($F(1, 847) = 263.112$, $p < .000$), composition ($F(1, 847) = 390.591$, $p < .000$), and grammar ($F(1, 847) = 328.505$, $p < .000$). The same was also true for the identity manipulation: helpfulness ($F(1, 847) = 5.767$, $p < .017$), composition ($F(1, 847) = 11.493$, $p < .001$), and grammar ($F(1, 847) = 6.058$, $p < .014$). The mere presence or absence of a real name associated with a review affected all quality measures. The interaction between identity and quality level is also significant ($F(3, 845) = 2.643$, $p < .048$), and the significant movement between anonymous and real identities can be localized to the bad review condition. These means may be seen in Table 8.

Table 8: Means by Review Quality and Author Identity

Review Quality	Author Label	Helpfulness	Composition	Grammar
Bad	Anonymous	49.41	4.72	2.94
Bad	Real	56.24	5.51	3.15
Good	Anonymous	75.93	7.57	3.95
Good	Real	76.30	7.65	4.05

Discussion

The most basic finding here is the key one: simply by using a real name on a review, that review is rated as being of better quality across multiple measures. Identity is indeed talent, to paraphrase Landy and Sigall. Associating a review with a real name improved subject assessment of the review, particularly when the review was of low quality. Although this cannot conclusively address the increase in assessed helpfulness of real name reviewers in the scraped data, it does indicate that some of that increase may have been due to perceptions that real name reviews are better. The second finding, of the interaction, supports the idea that identity is talent in the eyes of online consumers in a different way. For excellent reviews, those ones rated as the single most helpful for a product, there exists no significant difference in assessed quality between the identity conditions. However, poor reviews' assessments have a significant amount of leeway given to FID authors. Though there appears to be a ceiling effect on how much of a boost one may receive from disclosed identity when the review is already of high quality, it has a major impact on response toward sloppily written

reviews. Consider that most FID quality improvements in the scraped data were found not between identity conditions in the most elite population, but in the other groups with lower helpfulness ratings, which had room for this effect to occur.

Study 12: Creation of Reviews as Affected by Manipulated Identity

Purpose

All data up until this point has examined identity levels that are opted into voluntarily by the reviews' authors. Amazon review names are chosen at the time of review creation, and although the identities were manipulated in the previous study, the original authors wrote them under identities that they selected. However, as established in the introduction to this dissertation, voluntarily choosing one's displayed identity level is not the trend among many major industry players. Under the Facebook model, as opposed to Amazon's, every review would be left under one's real name, and if one were found not to be using that real name, a penalty would be incurred or the real name would be revealed.

Sites increasingly rely upon user-generated content to add value (Lee 2008). However, without an intrinsic reward system (such as cash rewards), there exists a complex interplay of factors that compel people to contribute. Personal satisfaction, creative demonstrations, attention from others, and a variety of other socially and personally-driven factors give people the motivation to exert time and effort in order to create content for which they won't be paid (Cha et al. 2007). Because of this, it is unsurprisingly both easy and common for consumers

to be put off participating in these sites if they encounter some negative personal or social feelings.

An industry example demonstrates the risk of forced identity disclosure on participation. Blizzard Entertainment, the creators of *World of Warcraft* (a popular online role-playing game), runs web forums that serve as both social gathering points and consumer interaction platforms. There are community areas in which to chat and seek help, as well as business-oriented areas, such as receiving technical support and asking account questions. For many issues, this web forum is the only reliable place to receive assistance from the company. In 2010, they announced the introduction of their Real ID system. With it, everyone in the forums would participate using the real name associated with the credit card on their account. Blizzard offered this as an introduction to the system:

> "Removing the veil of anonymity typical to online dialogue will contribute to a more positive forum environment, promote constructive conversations, and connect the Blizzard community in ways they haven't been connected before (Pegoraro 2010)."

This justification is much the same that was given by Facebook: by sharing everyone's identities, it will be possible to avoid some of the worst behavior associated with anonymity.

Reaction was immediate and strongly negative. One forum thread gathered more than 2,000 pages of discussion, largely unsupportive of the decision (Pegoraro 2010). A center for free online speech decried the move, warning that it could remove the benefit of escapism provided by entertainment and raising the

specter of negative real-world consequences for simple disagreements over a game (Pegoraro 2010). For underprivileged groups, there was concern over stalking and harassment if their visibly female or non-white name was made public without their consent (Anonymous 2010; Brooks 2010).

A writer at *The Washington Post* came to the same conclusion made in this work: that any pro-social value comes not in forcing true identity revelation, but in being held accountable to *some* persistent identity:

> "To me, Blizzard seems to be making two core mistakes. First, in most dysfunctional online forums (anybody remember Usenet?) the problem isn't anonymity but unaccountability: If anybody can easily look up everything you've written, and if a site's administrators limit duplicate "sock puppet" accounts, you can't hide from your past words, even if you didn't post them under your name. (Anonymous 2010)"

Real ID, as originally implemented, was shut down by Blizzard only three days after release. Ultimately, the company released what they labeled a new "BattleTag" system to address the Real ID issues. Under it, customers could interact with one single persistent pseudonymous identity, rather than being forced into using their real names (Brooks 2010; Parrish 2011). If people had been forced to use their real names to participate in this environment, then—for a variety of reasons—some would have chosen *not* to participate. Others would have reduced their participation levels.

Here, I examine whether forced identity disclosure affects effort (length) and writing quality. While the real world data shows that pseudonyms are associated with longer reviews, it is not clear whether this is due to a sense of freedom afforded by not revealing one's full identity or whether people who write longer reviews do not wish to protect their identities more than people who write shorter reviews. Also, given that the (very small) increase in helpfulness ratings of real name reviews may be due to an "identity is talent" effect, it is important to see whether assigning people to identity condition impacts review quality.

Participants

Undergraduate participants were selected only if they indicated that they had a persistent online identity that they believe represents them and which they would be unhappy if they were forced to change (determined during "mass testing" at the start of the term). By random session block, participants were randomly assigned into one of three username conditions upon entering the study: no anonymity [their real name], invested pseudonym [an important persistent online identity], or high anonymity [only the word "anonymous"]. Ultimately, ninety-nine participants completed the study ($N_{real} = 35$, $N_{pseudonymous} = 34$, $N_{anonymous} = 30$).

Procedure

All participants were told that they would review products live on Amazon.com and would be using (their real name/their favored online

identity/complete anonymity). It was emphasized that these reviews would be submitted to Amazon's server in the course of the experiment and would be available for anyone in the world to read, as well as being indexed by Google for later association with the authoring identity. Real name participants had it emphasized that their real identity would be associated; pseudonymous participants had it emphasized that their reviews would be associated with their cherished persistent on-line identity; anonymous participants had it emphasized that no one would be able to associate their reviews with any identity of theirs.

Four products actually sold on Amazon were tried and then reviewed. All products were chosen due to the expectation that they have some product feature that might make discussion of them a bit awkward for some users. Three of these products were similar to products in the previous study, while one was entirely new. A major constraint on product selection for the study was which products could be adequately sampled and assessed within a short time inside a full lab environment. These four products were sampled in a set order but with a random starting point for each participant.

Participants were told that they could write as much or as little as they liked, including no review at all. (This saved them effort and potential exposure but not time, as the study was conducted in a series of uniform rounds through which participants could not rush.) Two products were chosen with the expectation that they would be somewhat embarrassing for all participants. Two were chosen with the expectation that they would not be inherently embarrassing, but would be perceived as strongly gendered, and thus participants of the other

gender might feel slightly reluctant to have their identity associated with the product. This did result in one product per subject that was not expected to be embarrassing at all (the non-inherently embarrassing product which was aligned with that subject's gender). The two inherently embarrassing products were Metamucil cinnamon fiber wafers (a dietary supplement and constipation aid) and a novelty pencil sharpener shaped like a cat where the pencil opening corresponded to the cat's anus (the cat meows loudly as the pencil is sharpened). The two gendered products were an African safari hunting video game (male) and a cucumber-scented lotion (female).

At the start of each round, participants were asked to sample as much (if any) of their assigned product for that round as was desired, and then write the review (if any) that they wished to leave. Although Amazon requires a minimum word count to publish, subjects were allowed to write as little as they wished. It was emphasized that reviews did not need to meet some pre-set standard, but they were to write the review they would naturally choose to leave for these products, whether negative or positive, humorous or serious, long or short. After again emphasizing that opting out was acceptable, subjects were asked to raise their hand if they chose not to write a review for a round. Upon completing four rounds of product testing and reviews, they completed general product questions and demographics.

Results

Opt-Outs

Due to Amazon server unreliability I discovered during the design stage, my records of the review content were taken when I asked participants to copy and paste their reviews into a page at the end of the study for back-up. I saw here that there were more opt-outs than subjects had formally informed us, but the number was small overall (97% of expected reviews were completed) and could also be due to improperly pasting the content before moving on. Of the skipped reviews that were reported to the researcher, all four were about the fiber wafers. One participant in the real name condition stated that he did not wish to have his name associated with a constipation aid. Two participants in the pseudonymous condition said that they had allergies or gluten intolerance. One participant in the anonymous condition said that she didn't want to bother thinking of anything to say about the product. As well, one participant in the real name condition came to the researcher after a session and privately asked for his public reviews to be deleted. Although he had posted all four reviews and entered them into our research server for collection, he had become increasingly stressed over the course of the hour about the idea of his real name being on these product reviews online.

Length

Under some time constraint to try and review 4 products in one hour, experimental reviews were shorter than the scraped Amazon reviews. Due to this, the number of characters per review was analyzed rather than the number of words. A MANOVA analysis revealed significant differences in the amount written across the three identity conditions ($F(8, 188) = 1.995$, $p < .049$). Following up by individual product, the differences across identity conditions did not meet accepted significance levels for the wafers ($F(2, 96) = 2.392$, $p < .097$). The game ($F(2, 96) = 4.036$, $p < .097$), lotion ($F(2, 96) = 3.594$, $p < .021$), and pencil sharpener ($F(2, 96) = 4.869$, $p < .01$) all demonstrated significant differences across identity conditions. The mean number of characters per review is in Table 9. (Mean standard deviations are in parentheses.)

Table 9: Review Length by Condition

	Real Name	Pseudonym	Anonymous
Wafers	211 (111)	199 (120)	278 (220)
Video Game	219 (110)	203 (100)	303 (215)
Lotion	195 (102)	238 (119)	292 (205)
Pencil Sharpener	245 (125)	259 (154)	372 (236)

Follow-up contrasts were next conducted for the three products which differed across the three identity conditions. Real name participants wrote significantly less than anonymous subjects for all 3 products, game ($F(1, 63) = 4.100$, $p < .047$) pencil sharpener ($F(1, 63) = 7.595$, $p < .008$) and lotion ($F(1, 63) = 6.136$, $p < .016$). Anonymous respondents wrote significantly more than pseudonym respondents in their reviews of the game ($F(1, 62) = 5.627$, $p < .021$) and pencil sharpener ($F(1, 62) = 4.973$, $p < .029$) but not the lotion ($F(1, 62) = 1.554$, $p < .217$). No significant

differences were detected in the amount written between pseudonym and real name participants for any product.

Content

Computerized measures of review readability were also taken. Ultimately, reviews posted to a mass retailer like Amazon should be useful to other consumers as they make purchase decisions; "useful" is, after all, the phrasing used for the negative and positive voting that was recorded in the scraped data. To this end, I analyzed how easy the reviews were to read, by condition, using the Flesch Reading Ease measure, in which a higher score denotes an easier passage to comprehend (see Table 10). For this analysis, responses in each condition were removed because the review ended in mid-sentence (almost certainly an artifact of participants not cutting and pasting their review properly), after the analysis was initially run and returned an error in completing the assessment. For two reasons, this group was kept separate from the analysis above. First, the deleted reviews were obviously, and near-uniformly, missing only part of their last word; for example, they ended similarly to this real entry of "the moisture is long lasting and feels grea" *(sic)*. However, such a review would not be comprehensible to the readability analysis. Second, anonymous users were far more careless than the other groups in pasting their reviews. Even after removing all reviews with which the readability assessment had any issue, there were never any identity*product cells with fewer than 25 reviews for pseudonymous and real name users. However, anonymous users were so careless in their cutting and pasting that all

products only had 15-18 reviews for readability analysis. This damaged the potential for valid statistical tests, which were (unsurprisingly) more difficult to achieve in this smaller set. More importantly, this shows that the results above were obtained even when anonymous users were leaving out part of their review's full length with greater frequency than the other two groups. Had they pasted their reviews properly at a rate to match the other two conditions, their word count advantage over the other conditions would have only increased.

Table 10: Review Readability

Flesch Reading Ease			
	Real Name	Pseudonymous	Anonymous
Wafers	75.98	84.62	79.91
Video Game	77.56	82.18	74.41
Lotion	79.49	82.39	74.29
Pencil Sharpener	81.12	79.84	71.95

An omnibus MANOVA determined that writing quality differed across the three identity conditions for the 4 products ($F(8, 138) = 2.910, p < .005$).

Follow up contrasts show that anonymous reviews were marginally less readable than pseudonymous reviews ($F(4, 40) = 2.596, p < .051$). Real name reviews were significantly less readable than pseudonymous reviews ($F(4, 49) = 3.591, p < .012$). Although there is a significant difference between anonymous reviews and real name reviews ($F(4, 44) = 2.809, p < .037$), this is driven entirely by improved real name performance on the pencil sharpener reviews $F(1, 47) = 5.721, p < .021$); there are no significant differences on the other products. The three individually significant products between all conditions were the wafers ($F(2, 71) = 3.880, p < .025$), game ($F(2, 71) = 3.651, p < .031$), and pencil sharpener ($F(2, 71) = 3.272, p < .044$).

Discussion

Taken as a whole, these results suggest, surprisingly, that real name reviewers performed the worst out of the three conditions, given that they earned top scores on neither measure: they wrote less than anonymous reviewers and were less comprehensible or helpful than pseudonymous reviewers. Although pseudonymous reviewers wrote less than anonymous reviewers, their results were more comprehensible. Whether companies place more value either on the sheer quantity of output or the usefulness of their product reviews to a general audience, this study clearly does not support the view that FID leads to optimal performance when that disclosure is forced upon the user.

Within the studies that examined voluntary identity disclosure, support was found in correlational data for superior performance assessments of FID users as driven not by actual review quality, but by bias toward their real names in all but the most elite conditions. Users displaying their real names were shown to write less, particularly on items less central to Amazon's business model, and yet were judged as more helpful. When this was moved into an experimental manipulation, the indicated findings of the correlational, real world data were replicated in the laboratory.

When forced identity disclosure occurs, pseudonymous users find themselves as clearly stronger performers than real name users, who suffer even compared to those anonymous users that heads of industry have spoken out against. The first set of findings, about voluntary disclosure, prompts future

research into what, precisely, is driving a bias toward real name reviews in any level of assessment, even as pseudonymous users give indication that they may offer more useful reviews in various ways: longer, more critical, and less focused upon simple media items. The last study's findings about forced disclosure directly challenge industry trends and supports the Amazon model of voluntary identity disclosure, even if bias then occurs.

Conclusion

Across both real world data and laboratory experiments, identity has an effect on behavior. In this concluding section, I first summarize the findings of the dissertation's findings. Next, I apply these findings to the originally proposed model of Figure 1, create updated versions, and offer potential explanations for the behaviors therein. Lastly, I discuss potential future work, both in terms of specific experiments and broader theories to test.

In the following table, recall that the sample of the Hall of Fame population changed between aggregate collection and review-level collection, due to new members being inducted.

Table 11: Summary of Amazon.com Results

Amazon.com Data			
Measured Behavior	Sample Population	Identity Condition	Results
Incidence of Pseudonymous Identification	Hall of Fame	Pseudonymous	46 (39%)
		Real Name	77 (61%)
	Top Reviewers	Pseudonymous	1,631 (56%)
		Real Name	1,263 (44%)
	Random Reviewers	Pseudonymous	936 (62%)
		Real Name	564 (38%)
Average Aggregate Number of "Helpful" Votes	Hall of Fame	Pseudonymous	21,487.28
		Real Name	22,248.47
	Top Reviewers	Pseudonymous	2,508.21
		Real Name	3,006.81
	Random Reviewers	Pseudonymous	982.01
		Real Name	1,574.34

Measured Behavior	Sample Population	Identity Condition	Results
Average Aggregate Percentage of "Helpful" Votes	Hall of Fame	Pseudonymous	86.67
		Real Name	87.19
	Top Reviewers	Pseudonymous	87.67
		Real Name	87.67
	Random Reviewers	Pseudonymous	85.06
		Real Name	84.60
Average Word Count of Individual Reviews	Hall of Fame	Pseudonymous	340.56
		Real Name	335.64
	Top Reviewers	Pseudonymous	332.44
		Real Name	331.66
	Random Reviewers	Pseudonymous	321.81
		Real Name	304.81
Average Number of Votes per Review	Hall of Fame	Pseudonymous	12.33
		Real Name	10.86
	Top Reviewers	Pseudonymous	9.38
		Real Name	10.08
	Random Reviewers	Pseudonymous	8.82
		Real Name	11.63
Average Number of "Helpful" Votes per Review	Hall of Fame	Pseudonymous	10.89
		Real Name	9.42
	Top Reviewers	Pseudonymous	8.46
		Real Name	9.02
	Random Reviewers	Pseudonymous	7.54
		Real Name	9.97
Average Percentage of "Helpful" Votes per Review	Hall of Fame	Pseudonymous	71.28
		Real Name	72.23
	Top Reviewers	Pseudonymous	64.85
		Real Name	68.67
	Random Reviewers	Pseudonymous	63.16
		Real Name	67.70

Measured Behavior	Product Type	Identity Condition	Results
Average Product Ratings	Media	Pseudonymous	4.24
		Real Name	4.27
	Non-Media	Pseudonymous	4.18
		Real Name	4.14
Average Word Count	Media	Pseudonymous	349.67
		Real Name	336.15
	Non-Media	Pseudonymous	266.36
		Real Name	212.36
Average Number of "Helpful" Votes per Review	Media	Pseudonymous	7.07
		Real Name	10.13
	Non-Media	Pseudonymous	9.33
		Real Name	9.37
Average Percentage of "Helpful" Votes per Review	Media	Pseudonymous	65.58
		Real Name	70.16
	Non-Media	Pseudonymous	54.97
		Real Name	56.70
Incidence of Reviews	Media	Pseudonymous	80%
		Real Name	84%
	Non-Media	Pseudonymous	20%
		Real Name	16%

Table 12: Summary of Laboratory Results

Study 11: Manipulation of Presented Identity				
Review Quality	Identity Condition	Helpfulness	Composition	Grammar
Bad	Anonymous	49.41	4.72	2.94
	Real Name	56.24	5.51	3.15
Good	Anonymous	75.93	7.57	3.95
	Real Name	76.30	7.65	4.05
Study 12: Manipulation of Authorial Identity				
Identity Condition	Product	Review Length by Characters		
Anonymous	Wafers	278		
	Video Game	303		
	Lotion	292		
	Pencil Sharpener	372		
Pseudonymous	Wafers	199		
	Video Game	203		
	Lotion	238		
	Pencil Sharpener	259		
Real Name	Wafers	211		
	Video Game	219		
	Lotion	195		
	Pencil Sharpener	245		
Identity Condition	Product	Flesch Reading Ease		
Anonymous	Wafers	79.91		
	Video Game	74.41		
	Lotion	74.29		
	Pencil Sharpener	71.95		
Pseudonymous	Wafers	84.62		
	Video Game	82.18		
	Lotion	82.39		
	Pencil Sharpener	79.84		
Real Name	Wafers	75.98		
	Video Game	77.56		
	Lotion	79.49		
	Pencil Sharpener	81.12		

Given these results, a look back at the originally proposed behavioral models (Figure 1) is warranted. At this point, they are revised to reflect my findings. One of the most significant changes is splitting up the models into two groups: one for those users operating under *voluntary* identity disclosure, and the other for users operating under *mandatory* identity disclosure. As this was a key finding of this dissertation, it is important to reflect it in the models.

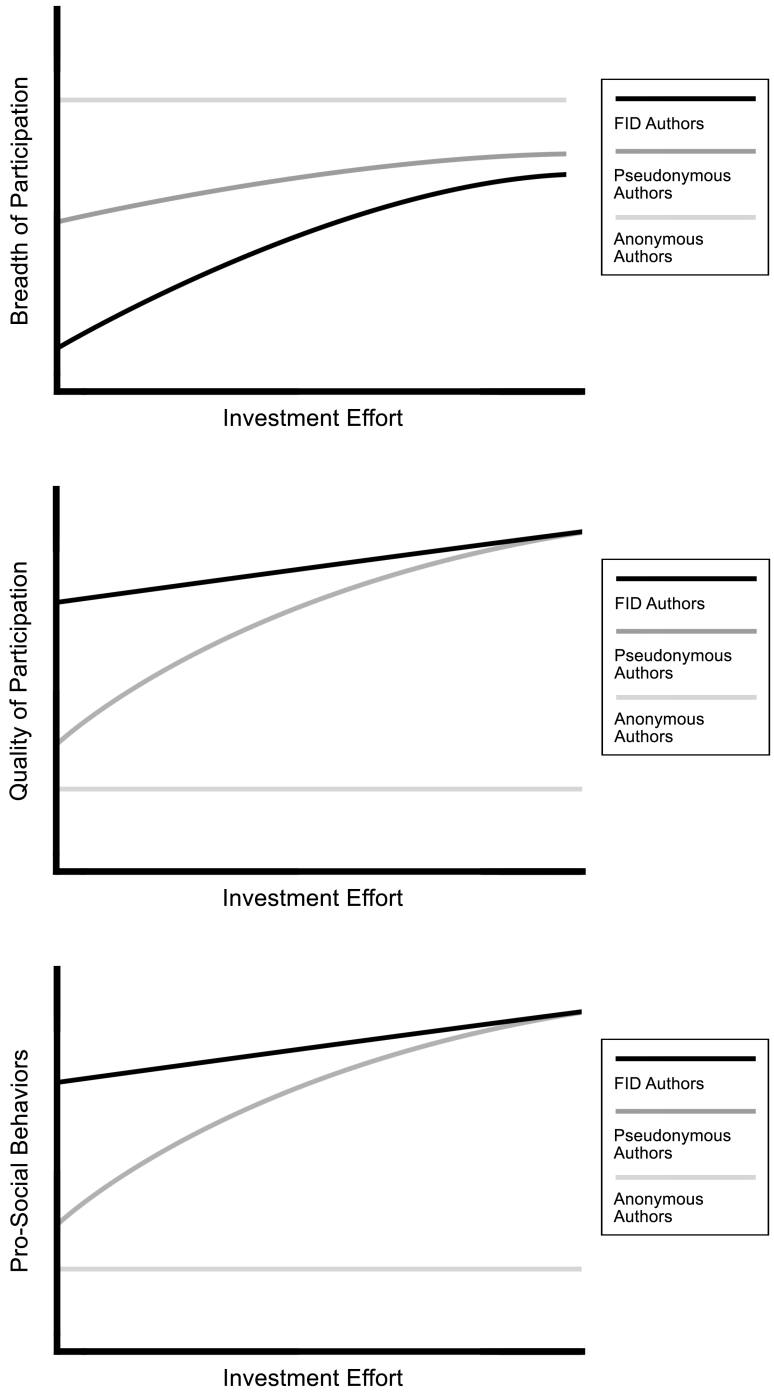

Figure 18: Revised Models of Online Behavior Under Voluntary Levels of Identity Revelation

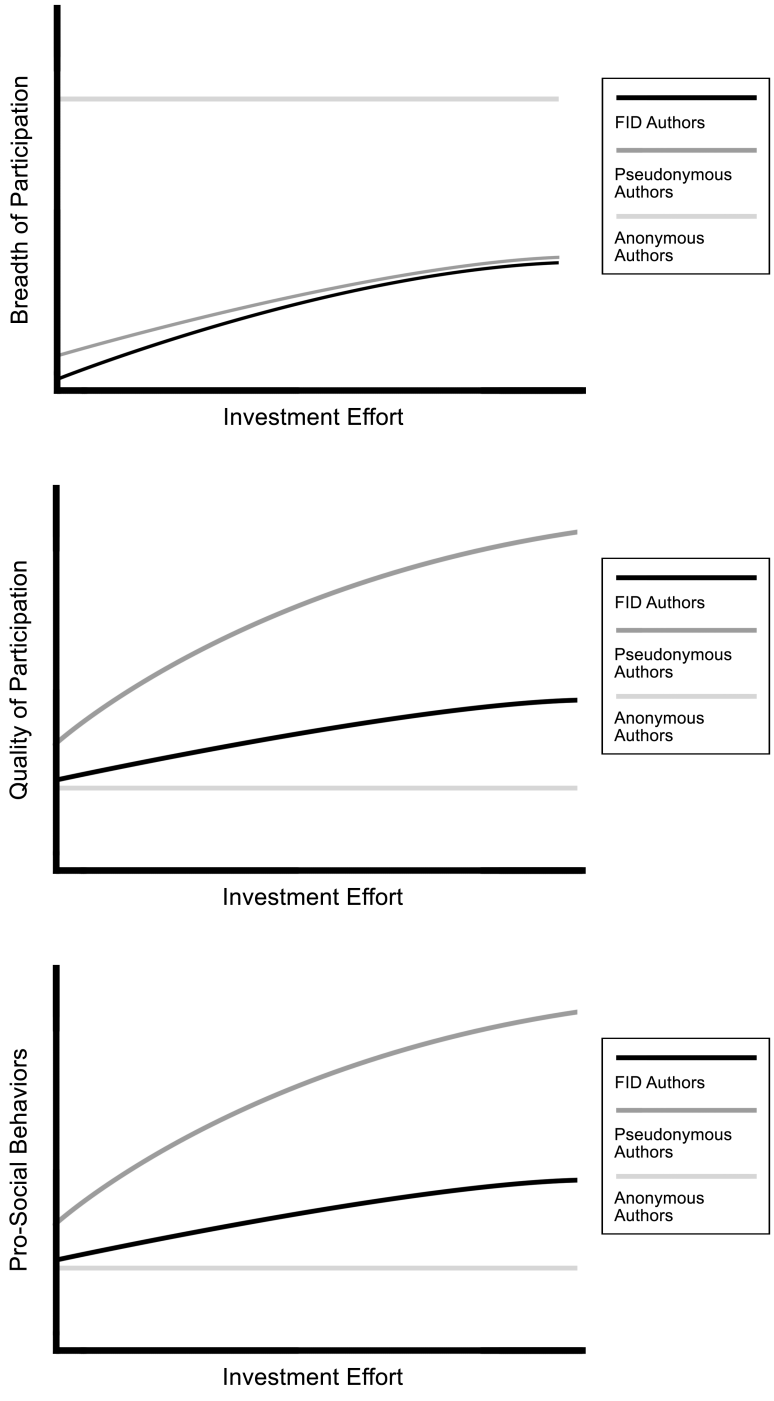

Figure 194: Revised Models of Online Behavior Under Mandatory Levels of Identity Revelation

Three changes were made to the original model to reflect findings in voluntary identity disclosure conditions. First, the breadth of participation increases over time, rather than having a flat effect based on identity level as I expected. Pseudonymous users do always maintain an advantage here over those users using their real names, as driven by two findings: one, that they write consistently longer reviews, and two, that their reviews are on more varied product types. Anonymous users are still hypothesized to be the strongest performers here, although they cannot be directly measured from Amazon.com data, and given that they cannot invest themselves into an anonymous identity, their curve remains flat. In the other two areas, the change comes from the altered slope of the pseudonymous improvement. I anticipated a relatively short investment in time in my original model, after which pseudonymous users performed on par with real name users in matters of quality and social connection. This did not prove to be the case. By the time of Hall of Fame data, it could be said that the groups had comparable (if not identical) levels of performance, but this took a significant amount of time and effort on those users' parts.

Within the mandatory disclosure models, the results look quite different. With questions of breadth, both pseudonymous and real name users are much more reluctant to participate, and there is little (but existent) hypothesized improvement over time. Anonymous users always maintain a strong advantage over them both. However, the social and quality-related curves remain the same for pseudonymous and anonymous users, but now real name users see their performances strongly depressed at all times. It should be noted that this model

series is more speculative about changes over time than are the Amazon.com-based ones, given that they are based on a single laboratory experiment and previously touched-upon theories about privacy.

Indeed, it is the literature on privacy that offers the most likely explanation for why there is such a difference between the two sets of models. As previously noted, it is possible to segment consumers based on their valuing of (or dismissal of) privacy during interactions with companies (Hann et al. 2007; Milberg et al. 1995). This idea of distinct groups of privacy-minded consumers does find support in this work, where consumers not only have different levels of output based on the identity they choose for themselves, but even focus on different types of products to review. When people are allowed to choose the level of privacy with which they are comfortable, the previous findings regarding greater anonymity leading to worsened behaviors (Diener et al. 1976; Festinger et al. 1952; Singer et al. 1965) appear to be dominant. However, when identity disclosure is mandatory, the research on privacy preferences comes to the forefront. As these consumers have their abilities to delineate boundaries and maintain desired levels of social contact taken away from them (Altman 1975; Altman and Chemers 1980), they behave poorly in return. A desire to maintain control over identity is important to us (Pedersen 1997), and that *control* appears to be the key piece in predicting how an online identity affects one's behavior. For both pseudonymous and real name users, their willingness to write suffered with mandatory identity disclosure. However, when it comes to social and quality matters, the pseudonymous identities did not receive the same penalties as real

name users, who began performing at levels roughly comparable to anonymous users. Mandatory identity disclosure hindered pseudonymous users' willingness to write, but not their quality or (as is speculated from this single experiment) their social performance.

Obvious future research questions arise while presenting these altered models. The most pressing is to conduct research on mandatory research disclosure with multiple pieces of output from the same users, to compare more directly to the Amazon.com data. Although laboratory experiments are possible, they seem unlikely to match the level of investment seen in Amazon data. A mandatory identity disclosure website with a significant volume of review content (or similar content) should be identified and collected in a similar fashion.

Large and powerful organizations with largely opaque agendas are making a concerted effort to convince the world of the general and perhaps universal value of FID in on-line commerce. The premise of sites like Facebook arguing for FID is centered upon an assumption that various kinds of positive behaviors are associated only with using one's real name online. The evidence provided in this paper demonstrates that for online product reviews, requiring FID is not entirely justified. This is important because many sites increasingly rely upon user-generated content to add value (Lee 2008; Parrish 2011). Without being paid for reviews, a complex interplay of factors lead people to contribute time and effort to such activities (Cha et al. 2007; Lee 2008). Because of this, it is both easy and common for consumers to be put off participating in these sites if they encounter

some negative personal or social feelings. The required disclosure of personal identity appears to be such a negative force.

There are other questions to immediately answer regarding some of the findings in this dissertation. At the end of the correlational studies, there was indication that differences in ratings could be driven not by true differences in output quality, but by the idea that "identity is talent." In other words, a rater bias toward real identities may exist and be driving these results. The obvious next step is to conduct a large-scale random sample of the reviews analyzed in the first part of this work and rate them for quality in a method similar to the review ratings done during the laboratory experiments. By analyzing their writing quality by hand (using MTurk workers who remain ignorant as to the authors' identities) and with computerized methods for both measuring quality and identifying sources of variation, any doubt about a bias toward real identities will vanish. As well, Amazon.com does make identities and status known with each review written, but as these identity labels are quite small, it is difficult to know whether readers are processing their meanings or simply noting their absence or presence at some subconscious level. A person using their real name always has one more graphic next to their name than someone who is not, and furthermore, these graphics look similar at a glance (being made of small blue text) to the badges given as a reward for strong performance. Having asked this, is it possible to manipulate the basic presentation of information on a website to eliminate variation in users' assessments of review quality by identity level? An online

experiment with MTurk users would be both efficient and illuminating on this question.

Given this work's findings in a largely unexplored field, there are other questions that arise for future study, as well. These are rooted in the potential theoretical explanations for the behaviors discovered here. One of the primary findings of this dissertation was the freeing effect of anonymity to write more consumer-generated content. An impression management perspective, touched upon during the initial literature review, may explain this. If we are indeed aware of stereotypical assumptions being made about our identities given some behavior that we present to the world (von Hippel et al. 2005), having that fear reduced could explain an increased willingness to discuss products that might strengthen any such assumptions. When we are anonymous, no assumption can be made about our identity; when we are pseudonymous, the assumption can be made only about this secondary identity that we have constructed. Given that we do exert more effort into presenting a favorable presentation of our identity to strangers (Leary et al. 1994; Tice et al. 1995), and online selves are a recognized form of managed identity (Donath 1999), this explanation makes sense for consumers who are, in effect, speaking through a megaphone about their personal experiences to a planet full of complete strangers.

Recall that in the laboratory study, with a smaller subject population, the word count effect was found only in total anonymity, while the large volume of scraped Amazon data consistently found the effects (whether with statistically significant differences or directionally) for a freeing effect on pseudonymous

users. This supports a perspective of impression management, given that effort is needed to maintain a socially desirable self (Holtgraves 2004): the further away from our permanent, real-world identities we are, the more we feel freed to write about any product before us. If we are released from the effort of maintaining a desirable identity, we can instead put that effort toward producing consumer content. Given these early findings, a reasonable future path of inquiry is to pursue these findings in the lab with an impression management perspective. If users with invested identities are once again recruited, and asked to perform content generation tasks related or unrelated to these constructed identities, is there a change in the freeing effect? While collecting the Amazon data, I saw a number of users whose profiles were centered around some specific topic; for example, music. Would a pseudonymous music reviewer feel free to write more about Windex or granola bars than he would about a Radiohead album, where he would always be conscious of the effect this review might have on his established identity as an expert reviewer? From the scraped Amazon data where real name users have written about thousands of different products, in nearly every product category imaginable, it can be said that there is a reliable stifling effect that happens to them across the board. If pseudonyms could be found to only write less in an area that relates strongly to that pseudonym's identity, while feeling freed to write more in unrelated areas, that would enhance our understanding of how people use and construct pseudonymous identities online. Some support has been found for this idea in studies of visual avatars, where users assigned to some

appearance online modified their behavior to align with this particular online identity (Yee and Bailenson 2007).

There must be a balance, however. It is one thing to simply note that an increased volume of discussion drives higher sales (Liu 2006); companies also wish for this content to be of high quality (Lee 2008; Parrish 2011). The discussion here relates to *invested* identities, and going by the scraped data findings, these identities must be very strongly invested indeed for their quality performance to match, or even outperform, those users who write with their real names. Here, two obvious questions have been identified for future study: one, are these differences in quality assessment driven not by the quality of the reviews themselves, but by more generous assessments of these reviews by others? The first lab study suggests that this might be so, and this would take previous findings about physical identity and show that it can, in some way, be mapped onto an online self (Landy and Sigall 1974). Other, more basic social behaviors have been found to translate to online environments (Yee et al. 2007), and so this is a reasonable assumption. If "identity is talent," and that, rather than an actual difference in quality, is driving the improvement in assessments, then how can this perceived gap be reduced? Dellarocas' (2010) examination of online reputation systems could be taken into this realm and manipulated to see how it affects not the production side of user content and participation, but the assessor side. Are there ways in which sites can display user identity alongside user content, such that others can recognize and reward the investment into a

pseudonym prior to that pseudonym having written thousands of reviews like the Hall of Fame members?

Two, is the effect of forced identity disclosure as pronounced as it appears to be from these initial dissertation studies? In the laboratory study, the quality benefits granted during voluntary identity disclosure (the Amazon data) vanished when users were forced to disclose their real names. This aligns with previous findings (outside the area of user-generated content) that consumers are reluctant to share their information with companies (Hoffman et al. 1999), prefer to maintain control over their identities (Cialdini 1993; Kelly and McKillop 1996; Lane and Wegner 1995), and engage in more self-disclosure when anonymous (Christopherson 2007; Joinson 2001; Lee, Im, and Taylor 2008). If these previous findings apply here to content generation, as this dissertation's findings support, then this is a challenge to current industry practices where such identity disclosure is commonly forced upon users, who can easily be put off producing valuable content for these sites if they feel uncomfortable (Cha et al. 2007; Lee 2008).

This research does have weaknesses, which should be addressed in future research. The current quality assessments do not take into account reviews' ability to affect product decisions, which is ultimately the value they can bring to a host site. In addition, the chosen products in the final study were chosen to be embarrassing to some degree. Future research may benefit from presenting a broader variety of product types from which to sample, as well as reducing the time pressure on subjects so that they may, if desired, write longer reviews. Finally, the readability measures used are preferable for a broad audience, but

may suggest reviews of lesser quality on other measures, particularly for complex products that may deserve a more complex (and therefore challenging) review.

In the scraped Amazon data, I saw elite reviewers who were willing to associate their real identities with vast quantities of user-generated content; more of them at the top level, in fact, than those who chose to use a pseudonym. Furthermore, although some populations identified more highly rated content for real name users, further analysis suggests that this may be an artifact of product focus, review valence, or simple identity disclosure, rather than actual review quality.

However, the key difference between the scraped Amazon data and the laboratory experiment is that all identity association was consensual for the scraped data content and forced upon our student subjects. Just as Salman Rushdie's identity was chosen for him by Facebook, some of my participants in the final study were forced into revealing their identities if they wished to write a review. When required to operate under a specific identity, as some large Internet companies wish to make mandatory, we saw both the quality and quantity of output fall. I believe, therefore, that it is worth revisiting the assumptions behind the push toward full identity disclosure in all areas of Internet commerce, and considering the role that invested pseudonyms could play in consumers' online experiences.

References

Altman, I. and M. Chemers (1980), *Culture and Environment*, Monterey, CA: Brook/Cole.

Amazon.com (2011), "Honoring our top reviewers over the last 11 years," http://www.amazon.com/forum/top%20reviewers?_encoding=UTF8&cdForum=Fx2Z5LRXMSUDQH2&cdThread=TxH7WGUB9C60P4.

--- (2012a), "Amazon's Top Customer Reviewers (Hall of Fame Tab)," http://www.amazon.com/review/hall-of-fame.

--- (2012b), "Amazon's Top Customer Reviewers (Top Reviewers Tab)," http://www.amazon.com/review/top-reviewers.

--- (2012c), "Fan Voters," http://www.amazon.com/review/guidelines/fan-voters.html/ref=cm_cr_tr_txt_popup.

--- (2012d), "Pen Names and Real Names™ ," http://www.amazon.com/gp/help/customer/display.html?ie=UTF8&nodeId=14279641.

Anderson, Chris (2006), *The Long Tail: Why the Future of Business is Selling Less of More*, New York: Hyperion.

Anonymous (2010), "Another Round of "Real Names Will Solve Everything": Blizzard Edition," http://geekfeminism.org/2010/07/07/another-round-of-real-names-will-solve-everything-blizzard-edition/.

Arakji, Reina Y. and Karl R. Lang (2007), "Digital Consumer Networks and Producer-Consumer Collaboration: Innovation and Product Development in the Video Game Industry," *Journal of Management Information Systems*, 24 (2), 195-219.

Armstrong, A. and J. Hagel (1996), "The Real Value of On-Line Communities," *Harvard Business Review*, 74 (3), 134-41.

Basuroy, S., S. Chatterjee, and S. A. Ravid (2003), "How Critical are Critical Reviews? The Box Office Effects of Film Critics, Star Power, and Budgets," *Journal of Marketing*, 67 (4), 103-17.

Bilton, Nick (2012), "Girls Around Me: An App Takes Creepy to a New Level," *New York Times*, March 30, 2012.

Brooks, Sean (2010), "Blizzard Looks To Chill Forum Speech with Real ID," https://www.cdt.org/blogs/sean-brooks/blizzard-looks-chill-forum-speech-real-id.

Brownlee, John (2012a), "'Girls Around Me' Dev: Our App's Not For Stalking Women, It's For Avoiding The Ugly Ones [Exclusive Interview]," http://www.cultofmac.com/158764/developers-behind-girls-around-me-stalking-app-explain-themselves-exclusive-interview/.

--- (2012b), "This Creepy App Isn't Just Stalking Women Without Their Knowledge, It's A Wake-Up Call About Facebook Privacy [Update]," http://www.cultofmac.com/157641/this-creepy-app-isnt-just-stalking-women-without-their-knowledge-its-a-wake-up-call-about-facebook-privacy/.

Cha, Meeyoung, Haewoon Kwak, Pablo Rodriguez, Young-Yeol Ahn, and Sue Moon (2007), "I Tube, You Tube, Everybody Tubes: Analyzing the World's Largest User Generated Content Video System," in *7th ACM SIGCOMM Conference on Internet Measurement*, New York: ACM, 1-14.

Chatterjee, P (2001), "Online Reviews: Do Consumers Use Them?," in *Advances in Consumer Research*, Vol. XXVIII, ed. MC Gilly and J Meyers-Levy, Salt Lake City, UT: Association for Consumer Research.

Chintagunta, Pradeep K., Shyam Gopinath, and Sriram Venkataraman (2010), "The Effects of Online User Reviews on Movie Box Office Performance: Accounting for Sequential Rollout and Aggregation Across Local Markets," *Marketing Science*, 29 (5), 944-57.

Christopherson, Kimberly M. (2007), "The Positive and Negative Implications of Anonymity in Internet Social Interactions: "On the Internet, Nobody Knows You're a Dog"," *Computers in Human Behavior*, 23 (6), 3038-56.

Cialdini, R. B. (1993), *Influence: Science and Practice*, New York: HarperCollins.

Connolly, T., LM Jessup, and JS Valacich (1990), "Effects of Anonymity and Evaluative Tone on Idea Generation in Computer-Mediated Groups," *Management Science*, 36 (6), 689-703.

Dellarocas, Chrysanthos (2006), "Strategic Manipulation of Internet Opinion Forums: Implications for Consumers and Firms," *Management Science*, 52 (10), 1577-93.

--- (2010), "Online Reputation Systems: How to Design One That Does What You Need," *MIT Sloan Management Review*, 51 (3), 33-37.

Diekmann, A. and D. Wyder (2002), "Trust and Reputation in Internet-Auctions," *Kölner Zeitschrift für Soziologie und Sozialpsychologie*, 54 (4), 674-93.

Diener, Edward, Scott C. Fraser, Arthur L. Beaman, and Roger T. Kelem (1976), "Effects of Deindividuation Variables on Stealing Among Halloween Trick-or-Treaters," *Journal of Personality and Social Psychology*, 33 (2), 178-83.

Dion, Karen, Ellen Berscheid, and Elaine Walster (1972), "What is Beautiful is Good," *Journal of Personality and Social Psychology*, 24 (3), 285-90.

Donath, Judith S. (1999), "Identity and Deception in the Virtual Community," in *Communities in Cyberspace*, ed. Marc Smith and Peter Kollock, New York: Routledge, 29-59.

Dubrovsky, V. J., B. N. Kiesler, and B. N. Sethna (1991), "The Equalization Phenomenon: Status Effect in Computer-Mediated and Face-to-Face Decision-Making Groups," *Human-Computer Interaction*, 2 (2), 119-46.

Duffy, JL (1999), "Initials: Why are authors so keen on them?," http://news.bbc.co.uk/2/hi/special_report/1999/02/99/e-cyclopedia/486501.stm.

Eagly, AH, W Wood, and S Chaiken (1978), "Causal Inferences about Communicators and Their Effect on Opinion Change," *Journal of Personality and Social Psychology*, 36 (4), 424-35.

Festinger, L., A. Pepitone, and T. Newcomb (1952), "Some Consequences of Deindividuation in a Group," *Journal of Abnormal and Social Psychology*, 47 (382-389).

Flanagin, A. J., V. Tiyaamornwong, J. O'Connor, and D. R. Seidold (2002), "Computer-Mediated Group Work: The Interaction of Member Sex and Anonymity," *Communication Research*, 29, 66-93.

Glowniak, Jerry (1998), "History, Structure, and Function of the Internet," *Seminars in Nuclear Medicine*, 28 (2), 135-44.

Guadagno, Rosanna E. and Robert B. Cialdini (2007), "Gender differences in impression management in organizations: A qualitative review," *Sex Roles*, 56 (7-8), 483-94.

Hann, Il-Horn, Kai-Lung Hui, Tom S. Lee, and Ivan P. L. Png (2007), "Overcoming Information Privacy Concerns: An Information Processing Theory Approach," *Journal of Management Information Systems*, 24 (2), 13-42.

Haraway, Merrill D. (1990), "A Manifesto for Cyborgs: Science, Technology, and Socialist Feminism in the 1980's," in *Feminism/Postmodernism*, ed. L. J. Nicholson, London: Routledge.

Hayne, S. C. and R. E. Rice (1997), "Attribution Accuracy When Using Anonymity in Group Support Systems," *International Journal of Human—Computer Studies*, 47, 429-52.

Hinduja, Sameer (2008), "Deindividuation and Internet Software Piracy," *Cyberpsychology & Behavior*, 11 (4), 391-98.

Hoffman, DL, TP Novak, and MA Peralta (1999), "Information Privacy in the Marketspace: Implications for the Commercial Uses of Anonymity on the Web," *Information Society*, 15 (2), 129-39.

Holtgraves, T. (2004), "Social desirability and self-reports: Testing models of socially desirable responding," *Personality and Social Psychology Bulletin*, 30 (2), 161-72.

Huang, JH and YF Chen (2006), "Herding in Online Product Choice," *Psychology & Marketing*, 23 (5), 413-28.

Jeppesen, Lars Bo and Lars Frederiksen (2006), "Why Do Users Contribute to Firm-Hosted User Communities? The Case of Computer-Controlled Music Instruments," *Organization Science*, 17 (1), 45-63.

John, Leslie K., Alessandro Acquisti, and George Loewenstein (2011), "Strangers on a Plane: Context-Dependent Willingness to Divulge Sensitive Information," *Journal of Consumer Research*, 37 (5), 858-73.

Joinson, Adam N. (2001), "Self-Disclosure in Computer-Mediated Communication: The Role of Self-Awareness and Visual Anonymity," *European Journal of Social Psychology*, 31 (2), 177-92.

Kelly, A. E. and K. J. McKillop (1996), "Consequences of Revealing Personal Secrets," *Psychological Bulletin*, 120, 450-65.

Kiesler, S., J. Siegel, and T. W. McGuire (1984), "Social Psychological Aspects of Computer-Mediated Communication," *American Psychologist*, 39, 1123-34.

Kiesler, S., D. Zubrow, A. M. Moses, and V. Geller (1985), "Affect in Computer-Mediated Communication: An Experiment in Synchronous Terminal-to-Terminal Discussion," *Human—Computer Interaction*, 1 (77-104).

Klar, Y. and E. E. Giladi (1997), "My Group Can Be Below the Group's Average: A Robust Positivity Bias in Favor of Anonymous Peers," *Journal of Personality and Social Psychology*, 73 (5), 885-901.

Landy, David and Harold Sigall (1974), "Beauty is Talent: Task Evaluation as a Function of the Performer's Physical Attractiveness," *Journal of Personality and Social Psychology*, 29 (3), 299-304.

Lane, J. D. and D. M. Wegner (1995), "The Cognitive Consequences of Secrecy," *Journal of Personality and Social Psychology*, 69, 237-53.

Leary, MR, JB Nezlek, D Downs, J Redford-Davenport, J Martin, and A McMullen (1994), "Self-Presentation in Everyday Interactions: Effects of Target Familiarity and Gender Composition," *Journal of Personality and Social Psychology*, 67 (4), 664-73.

Lee, Doo-Hee, Seunghee Im, and Charles R. Taylor (2008), "Voluntary Self-Disclosure of Information on the Internet: A Multimethod Study of the Motivations and Consequences of Disclosing Information on Blogs," *Psychology & Marketing*, 25 (7), 692-710.

Lee, Edward (2008), "Warming Up to User-Generated Content," *University of Illinois Law Review*, 2008 (5), 1459-548.

Liu, Yong (2006), "Word of Mouth for Movies: Its Dynamics and Impact on Box Office Revenue," *Journal of Marketing*, 70 (3), 74-89.

Lwin, May O. and Jerome D. Williams (2003), "A Model Integrating the Multidimensional Developmental Theory of Privacy and Theory of Planned Behavior to Examine Fabrication of Information Online," *Marketing Letters*, 14 (4), 257-72.

Ma, Meng and Rita Agarwal (2007), "Through a Glass Darkly: Information Technology Design, Identity Verification, and Knowledge Contribution in Online Communities," *Information Systems Research*, 18 (1), 42–67.

Madrigal, Alexis (2011), "Facebook Tells Salman Rushdie He Has to Go By His Given Name, Ahmed Rushdie," *The Atlantic*.

Mann, L. (1981), "The Baiting Crowd in Episodes of Threatened Suicide," *Journal of Personality and Social Psychology*, 41 (4), 703-09.

Mayzlin, D. (2006), "Promotional Chat on the Internet," *Marketing Science*, 25 (2), 155-63.

McKenna, Katelyn Y. A. and John A. Bargh (1998), "Coming Out in the Age of the Internet: Identity "Demarginalization" Through Virtual Group Participation," *Journal of Personality and Social Psychology*, 75 (3), 681-94.

Messinger, Paul R., Xin Ge, Eleni Stroulia, Kelly Lyons, Kristen Smirnov, and Michael Bone (2008), "On the Relationship between My Avatar and Myself," *Journal of Virtual Worlds Research*, 1 (2).

Messinger, Paul R., Eleni Stroulia, Kelly Lyons, Michael Bone, Annie Niu, Kristen Smirnov, and Stephen Perelgut (2009), "Virtual Worlds — Past, Present, and Future: New Directions in Social Computing," *Decision Support Systems*, 47, 204-28.

Milberg, Sandra J., Sandra J. Burke, H. Jeff Smith, and Ernest A. Kallman (1995), "Values, Personal Information, Privacy, and Regulatory Approaches," *Communications of the ACM*, 38 (12), 74.

Mizerski, RW (1982), "An Attribution Explanation of the Disproportionate Influence of Unfavorable Information," *Journal of Consumer Research*, 9 (3), 301-10.

Moses, Asher (2011), "Trending to a Tragedy: How Viral Fame Can Burn the Young," http://www.smh.com.au/technology/technology-news/trending-to-a-tragedy-how-viral-fame-can-burn-the-young-20110325-1c9cd.html.

Olson, Ingrid R. and Christy Marshuetz (2005), "Facial Attractiveness Is Appraised in a Glance," *Emotion*, 5 (4), 498-502.

Orenga, V., A. M. Zornoza, F. Prieto, and J. M. Peiró (2000), "The Influence of Familiarity Among Group Members, Group Atmosphere and Assertiveness on Uninhibited Behavior Through Three Different Communication Media," *Computers in Human Behavior*, 16, 141-59.

Parks, Malcom R. and Kory Floyd (1996), "Making Friends in Cyberspace," *Journal of Computer-Mediated Communication*, 1 (4).

Parrish, Kevin (2011), "Blizzard Introduces BattleTag as Real ID Alternative," http://www.tomshardware.com/news/Real-ID-BattleTag-Battle.net-Blizzard-Diablo-3,14301.html.

Pedersen, D. M. (1997), "Psychological Function of Privacy," *Journal of Environmental Psychology*, 17, 147-56.

Pegoraro, Rob (2010), "World of Warcraft users blast Blizzard's 'Real ID' rule (update: Blizzard retreats)," http://voices.washingtonpost.com/fasterforward/2010/07/world_of_warcraft_real_names.html.

Porter, Constance Elise and Naveen Donthu (2008), "Cultivating Trust and Harvesting Value in Virtual Communities," *Management Science*, 54 (1), 113-28.

Poster, M. (1990), *The Mode of Information: Poststructuralism and Social Context*, Chicago: Polity.

Postmes, T. and R. Spears (1998), "Deindividuation and Antinormative Behavior: A Meta-analysis," *Psychological Bulletin*, 123 (3), 238-59.

Rains, Stephen A. and Craig R. Scott (2007), "To Identify or Not to Identify: A Theoretical Model of Receiver Responses to Anonymous Communication," *Communication Theory*, 17 (1), 61-91.

Ren, Yuqing, Robert Kraut, and Sara Kiesler (2007), "Applying Common Identity and Bond Theory to Design of Online Communities," *Organization Studies*, 28 (3), 377-408.

Schul, Yaacov, Ruth Mayo, and Eugene Burnstein (2008), "The Value of Distrust," *Journal of Experimental Social Psychology*, 44 (5), 1293-302.

Sen, Shahana and Dawn Lerman (2007), "Why Are You Telling Me This? An Examination into Negative Consumer Reviews on the Web," *Journal of Interactive Marketing*, 21 (4), 76-94.

Sengupta, Somini (2011), "Rushdie Runs Afoul of Web's Real-Name Police," *New York Times*, November 14, 2011.

Siegel, J., V. Dubrovsky, S. Kiesler, and T. W. McGuire (1986), "Group Processes in Computer-Mediated Communication," *Organizational Behavior and Human Decision Processes*, 37, 157-87.

Singer, J., C. Brush, and S. Lublin (1965), "Some Aspects of Deindividuation: Identification and Conformity," *Journal of Experimental Social Psychology*, 1 (4), 356-78.

Sproull, L. and S. Kiesler (1986), "Reducing Social Context Cues: Electronic Mail in Organizational Communication," *Management Science*, 32 (11), 1492-512.

Sproull, L. and S. Kiesler (1991), *Connections: New Ways of Working in the Networked Organization*, Cambridge, MA: MIT Press.

Sui, Jie and Chang Hong Liu (2009), "Can Beauty Be Ignored? Effects of Facial Attractiveness on Covert Attention," *Psychonomic Bulletin & Review*, 16 (2), 276-81.

Suler, John (2004), "The Online Disinhibition Effect," *Cyberpsychology & Behavior*, 7 (3), 321-26.

Taylor, Chris (2011), "How Salman Rushdie Used Twitter to Defeat Facebook," http://mashable.com/2011/11/14/salman-rushdie-facebook/.

Tice, DM, JL Butler, MB Muraven, and AM Stillwell (1995), "When Modesty Prevails: Differential Favorability of Self-Presentation to Friends and Strangers," *Journal of Personality and Social Psychology*, 69 (6), 1120-38.

van 't Wout, M. and Alan G. Sanfey (2008), "Friend or Foe: The Effect of Implicit Trustworthiness Judgments in Social Decision-Making," *Cognition*, 108 (3), 796-803.

von Hippel, W., C. von Hippel, L. Conway, K. J. Preacher, J. W. Schooler, and G. A. Radvansky (2005), "Coping with stereotype threat: Denial as an impression management strategy " *Journal of Personality and Social Psychology*, 89 (1), 22-35.

Werner, C. M., I. Altman, and B. B. Brown (1992), "A Transactional Approach to Interpersonal Relationships: Physical Environment, Social Context, and Temporal Qualities," *Journal of Social and Personal Relationships*, 9 (2), 297-323.

Yee, Nick and Jeremy N. Bailenson (2007), "The Proteus Effect: The Effect of Transformed Self-Representation on Behavior," *Human Communication Research*, 33 (3), 271-90.

Yee, Nick, Jeremy N. Bailenson, Mark Urbanek, Francis Chang, and Dan Merget (2007), "The Unbearable Likeness of Being Digital: The Persistence of Nonverbal Social Norms in Online Virtual Environments," *CyberPsychology & Behavior*, 10 (1), 115-21.

Zebrowitz, LA., L. Voinescu, and M.A. Collins (1996), ""Wide-Eyed" and "Crooked-Faced": Determinants of Perceived and Real Honesty Across the Life Span," *Personality and Social Psychology Bulletin*, 22 (12), 1258-69.