Persuasive Design Techniques in the Attention Economy: User Awareness, Theory, and Ethics

by Devangi Vivrekar

A thesis submitted in partial fulfillment

of the requirements for the degree of

Master of Science

at Stanford University

June 6, 2018

To the Directors of the Program on Symbolic Systems: I certify that I have read the thesis of Devangi Vivrekar in its final form for submission and have found it to be satisfactory for the degree of Master of Science.

Signed Electronically
06/06/2018

Professor James Landay (Principal Advisor)
Computer Science

To the Directors of the Program on Symbolic Systems: I certify that I have read the thesis of Devangi Vivrekar in its final form for submission and have found it to be satisfactory for the degree of Master of Science.

Signed Electronically
06/06/2018

Professor Alia Crum (Second Reader)
Psychology

Acknowledgements

I would like to express my sincere gratitude to Professor James Landay, my primary advisor, for overseeing this thesis over the course of my coterm, and to Professor Alia Crum for being my second reader. I would also like to thank Grace Wu and Gobi Dasu who acted as co-authors on the study described in Chapter 3, and Professor Michael Bernstein and Geza Kovacs for their advice and feedback on the work done in Chapter 3 with the Chrome extension HabitLab. My sincere thanks to Damon Horowitz and Mikey Siegel, instructors in my Symbolic Systems coursework who deeply influenced my thinking early on, and Tristan Harris for inspiring me to study this topic. Finally, I would like to thank my parents and sister for their support throughout my academic journey.

Table of Contents

1 Introduction

The systematic study of persuasion has captured researchers' interest since the advent of mass influence mechanisms such as radio, television, and advertising. With the unprecedented growth of popular social media applications that are ubiquitously accessible on smart devices, consumers' attention, attitudes, and behaviors are constantly influenced by persuasive design techniques on platforms that profit by maximizing users' time spent on site. We take three different approaches - empirical, theoretical, and philosophical - to better understand the awareness, mechanisms, and ethicality of persuasive design techniques in the modern attention economy.

In **Chapter 2**, we review related work on the attention economy, persuasive technology, cognitive foundations of persuasion, the metaphysics of different modes of influence, persuasive strategies in social psychology, and identifications of persuasive design techniques in the attention economy.

Many users feel that these digital platforms draw them in and manipulate their time and behavior, but they lack a detailed awareness of persuasive design techniques in context of their use of the products. In **Chapter 3**, we discuss the creation and evaluation of a system that makes persuasive design techniques visible on social media. We compare the effectiveness of our system in improving user knowledge of such techniques to traditional methods of educating users, finding a significant improvement when our system is used.

Although there exists a rich social psychology literature on methods of persuasion and exploitable cognitive biases, we lack a mapping of the specific persuasive design techniques used by products like Facebook or LinkedIn onto this persuasive space. In **Chapter 4**, we take a theoretical approach to persuasive design with the goals of showing how design techniques interrelate to function as influential systems, and of providing a more nuanced vocabulary with which to discuss them, drawn from a range of relevant disciplines. We map our datasets of persuasive design techniques used on Facebook and LinkedIn onto existing theoretical frameworks in behavioral design, and also identify useful dimensions for future taxonomies.

Finally, we find it important to include a detailed discussion of the ethical questions surrounding persuasive technologies in the attention economy. In **Chapter 5**, we discuss philosophical approaches to the balance between persuasive structure and human agency, and compile useful ethical heuristics to make progress on the question of what constitutes ethical persuasive design. We discuss the method of value-driven design as a way to create products designed for intention-fulfillment rather than impulse-fulfillment, and we provide a sample phenomenological analysis of a modern product to showcase how philosophy of technology can add value to our understanding of what technology adds and subtracts from our lives. In **Chapter 6,** we discuss opportunities for future work and present our conclusions.

2 Related Work

We begin by reviewing related work on the emergence of the attention economy, the foundations of the field of study of persuasive technology, the effects of digital technologies on cognition, the distinctions between modes of influence related to persuasion, the social psychology account of persuasive strategies, and the identification of persuasive design techniques in digital products.

2.1 The Attention Economy

The abundance and rapid growth in popularity of information technologies over the past two decades has created a marketplace where consumers' attention is a scarce and valuable resource. Herb Simon first identified this phenomenon, now called the "attention economy," when he wrote about the "need to allocate that attention efficiently among the overabundance of information sources that might consume it" [1]. By using "valuable intangibles" in their design, such as immediacy, personalization, accessibility, and findability, popular digital applications constantly influence consumers' attention, attitudes, and behaviors [2]. By giving consumers the ability to share, comment on, and create their own content on social network services (SNS), these technologies empower users to curate and reach their own audiences, furthering information abundance and constraining the resource of attention on not only an individual level, but also the scale of entire social populations [3]. Persuasion scholar B.J. Fogg has called the influential reach of Facebook a social phenomenon of "Mass Interpersonal Persuasion" due to the culmination of persuasive experience, automated structure, social distribution, rapid cycle, social graph, and measured impact [4]. These massive platforms, exemplified by Facebook, YouTube, Snapchat, Twitter, and LinkedIn, among others, often aim to maximize users' time spent on their sites. The desire to maximize the metric of time spent on these platforms often stems from their bottom line: profit from advertisements. Advertisers pay the platforms to host their ads on the sites, and users' views and clicks of the ads become a "payment" of time and attention for their free use of the products. The digital landscape, then, has become situated in an "economy" where time and attention are the currencies. The focus of this work is to analyze the design of mass consumer digital products in the attention economy.

Users around the world spend large amounts of time on these products. 68% of American adults across a wide range of ages and demographics use Facebook, and 74% of them use the site daily [5]. Among 18-24 year old Americans, 78% use Snapchat, of which 71% visit the app multiple times a day [6]. 94% of 18-24 year old Americans use YouTube, 71% use Instagram, and nearly half (45%) use Twitter [6]. The median American uses at least three of the following eight products: YouTube, Facebook, Instagram, Pinterest, Snapchat, LinkedIn, Twitter, and WhatsApp [6]. Globally, Facebook has over 2 billion monthly active users and over 1.3 billion daily active users [8]. For comparison, this is just under the 2.3 billion followers of Christianity, the world's largest religion [9]. In addition to high frequencies of use, these technologies are characterized by long durations of use. For example, on mobile devices, YouTube viewing sessions last over 60 minutes on average, and 70% of time spent on YouTube is driven by automated recommendations, according to its chief product officer [10].

There are two broad categories of motivation due to which people engage with these products: instrumental and habitual. Instrumental use occurs when a user approaches and engages the technology with a specific purpose, and habitual use occurs when a user engages with the technology less intentionally and out of habit to merely pass time [11]. Hiniker et. al. report that this dichotomy of styles of use could

contribute to the level of meaning drawn from experiences engaging with products. Furthermore, they note that the same products can be used in both these modes of engagement: in 39% of samples, participants described their Gmail use as passively getting information, whereas Gmail is listed under the "Communication" category on the Google Play Store - a category that might suggest instrumental use [11]. Snapchat is listed under "Social" apps, which might suggest more habitual use, but in 66% of samples, participants reported active engagement rather than passive social media consumption. Almost all participants in a study of meaningful smartphone use reported passively scrolling on social media as a meaningless activity, and many reported a loss of autonomy, citing the habitual nature of their social media checking habits [11]. Even instrumental motivation has been found to degrade over the duration of use of these apps, leading to an "erosion of intention" by the very products that exist to help us intentionally consume information and communicate [11].

Technologist and design ethicist Tristan Harris describes the attention economy using the analogy of a city whose urban planners are the large corporations (e.g., Apple and Google) that lay the infrastructure for the rest of the applications in the market and function as monopolies that set the standard for what other applications must do to remain competitive in harnessing user attention [12]. When all applications, even meditation apps that intend to instill calm, must compete with social media apps for morning attention, they join what Harris calls the "arms race for attention" [12]. This proliferates techniques that "race to the bottom of the brainstem" [12].

The effects of such techniques can be seen indirectly in user sentiment and the general epidemic of overconsumption and underfulfillment - a phenomenon that has been present in many social trends including obesity, heart disease, industrial pollution, social isolation, diabetes, credit card debt, and now Internet addiction [13]. Many users are dissatisfied with the time they spend using these apps, especially on their phones [11]. In a study done by Google, vice president of product management Sameer Samat reported that "Over 70 percent of users had told Google that they wanted help striking a balance between their digital life and real-world interactions" [14]. This has recently prompted the addition of new features to its latest operating system "Android P" that reflect design ethicist Harris's calls for ensuring that time spent on these products is "time well spent" [14]. While Facebook has not replicated this feature, it has started changing its algorithm to promote interpersonal interactions and downrank viral videos that elicit passive consumption [15]. Nevertheless, the problems that arise for users of these products, as well as for the societies in which the products are fundamentally embedded, are far from solved. Thus, the attention economy has become an important subject of research for persuasion scholars who study the design of the technology that has these far-reaching effects.

2.2 Persuasive Technology

The subset of persuasive technology we focus on in this work is that of digital technologies in the attention economy that intend to persuade. Digital persuasive technologies differ from traditional persuasive media in that they act as "persuasive intermediaries" between the persuader and the target of persuasion, and, unlike billboards or static content, they interact dynamically with the target [16]. The technologies we analyze are all adaptive persuasive technologies that use both end-adaptive and means-adaptive strategies, meaning that they learn about the user over time to better tailor their persuasion. End-adaptive strategies are

those where the knowledge gained of the user has minimal application to future persuasive attempts in other domains, whereas means-adaptive strategies may apply in other domains [17,18].

The "intermediary" role of these technologies is the subject of study of the field of captology - an abbreviation for "computers as persuasive technologies." In the spirit of captology, we focus on the planned persuasive effects of computer technologies, and study both the attitudes and behavior changes intended by designers (where intent is endogenous as opposed to exogenous) [19]. However, as Atkinson warns, "if the program results in unforeseen, induced behavioural changes, it is appropriate for the discipline to assume responsibility and to name such phenomena" [20]. The two broad kinds of persuasion conducted by digital technologies include "macrosuasion," the overall persuasive intent of a product, and microsuasion, the "smaller persuasive elements to achieve an overall goal" [19]. Microsuasion is embodied through elements of visual and interaction design such as dialogue boxes, icons, interaction patterns, or reward strategies [19]. It is these more granular microsuasive techniques that are of interest in the majority of this work.

2.3 Cognitive Effects

We now discuss prior work on the effects of products in the attention economy on cognition. By virtue of competing in the attention economy, these products capture not only users' time but also their attention, or "focused mental engagements" on particular sets of information [21]. For example, many of these products use structured notifications to draw users in, interrupting other behaviors. This, in some part, results in people unlocking their phones over 150 times a day [7]. The distraction cost of receiving a phone notification during attention-demanding tasks has been shown to be comparable in magnitude to that of actively using a mobile phone for texting or making a voice call [22]. When the interruption is contextually different from the task, it causes increased stress, alters work strategies and mental states, and increases task reorientation time [23]. Interruption and distraction from digital technologies can have more serious consequences as well. Case studies have documented shallow breathing and temporary suspensions of breath while checking email - a phenomenon called email apnea [24]. In safety-critical domains, an ill-timed notification could alter response times in high-stakes situations, causing loss of life or catastrophic damage [25].

The dual-process model of persuasion describes two ways of processing persuasion: systematically and heuristically. In systematic processing, people differentiate strong and weak arguments, and are not affected by variables outside of the substance of the message, such as message length or attractiveness of the message source [26]. Systematic processing is triggered by things like forewarning that a persuasive message will be presented. In heuristic processing, people are susceptible to factors external to the substance of message. Heuristic processing is triggered by the presence of heuristic cues such as social proof, which is a technique commonly used by SNS [26]. The design of technologies in the attention economy often includes such stimuli that act as "demandances," which, unlike the utility-based concept of affordances (clues about how something should be used), are psychological "pulls" to execute a specific action, like an "itch" that drive goal-directed behavior [27].

These effects on users' attention have an impact beyond control of momentary awareness; they are also related to goal-directed achievement and higher-order aspects of cognition relevant to self-identity. These three levels of cognition have been described as the "spotlight," "starlight," and "daylight" of attention [28]. The "spotlight" refers to perceptual attention towards task-salient features of the environment and

interaction rules pertinent to those features. "Starlight" refers to goals relevant to who a person wants to be, and how current tasks relate to those higher goals. And "daylight" refers to metacognition, or defining the higher-level goals in the first place. By influencing our momentary attention, affecting our daily attitudes, and changing our behaviors over time, persuasive technologies affect all three of these cognitive modes. For example, video advertisements on Instagram direct users' perceptual attention towards sponsored content, while its aesthetically filtered photo feed leads to deeper attitudes like "envy spiraling," where users compare themselves both subliminally and explicitly to glamorous images. This, in turn, can lead to longer-term changes in self-image, and by one user's account, negatively affect mental health and self-worth issues [15,29]. Persuasion ethicist James Williams similarly categorizes types of distraction in terms of the effect they have on cognition: functional, which interferes with doing what you want to do and affects cognitive processes like task or goal alignment and intention awareness; existential, which interferes with being who you want to be and affects value and identity alignment, workstyle, regret, and time management; and epistemic, which interferes with wanting what you want to want and affects reflection, willpower, and reasoning [30]. The compounding effect of these products on higher-level cognition can then be extended to impact the very basis of our "freedom, wellbeing, and even the integrity of the self," which demonstrates how these seemingly small design techniques can have large cognitive impacts when interacted with so frequently and intimately [31].

Persuasive design techniques often engage or exploit cognitive biases and psychological limits, misperceptions, and fallacies [32]. They trigger the same cognitive processes that, for example, cause people to spend more on dinner when a restaurant is called Studio 97 instead of Studio 17, or to estimate higher athlete performance for players with higher jersey numbers [33]. These biases have been extensively studied and documented in the field of social psychology, which laid the groundwork for the field of behavioral economics. We now list some of these common biases and fallacies. Biases that stem from lacking full meaning in a situation include confabulation, anthropomorphism, authority bias, bandwagon effect, halo effect, reactive devaluation, denomination effect, illusion of transparency, spotlight effect, self-consistency bias, hindsight bias, time-saving bias, and impact bias (overestimating the length or intensity of future feeling states) [34]. Biases stemming from the need to make decisions quickly include social desirability bias, illusion of control, risk compensation, actor-observer bias, loss aversion, endowment effect, status quo bias, ambiguity bias, decoy effect, hyperbolic discounting, irrational escalation of commitment, appeal to novelty, identifiable victim effect, social comparison, decision fatigue, ratio bias, regret aversion, Elaboration Likelihood (when people rely on peripheral rather than central processing), and forecasting errors [34,35]. The overabundance of information can also trigger biases such as the availability heuristic, bizarreness effect, anchoring/framing, priming, confirmation bias, bias blind spot, choice overload, ostrich effect, diversification bias, and partitioning/mental accounting [34]. A final category of biases is due to memory constraints: the spacing effect, fading affect bias, suffix effect, primacy effect, peak-end rule, Google effect, and the next-in-line effect [34].

Users' mental models of these platforms vary greatly. In studying people's stories about how the Facebook algorithm ordered posts in their newsfeed, Eslami et. al. found that people had a range of often false "folk theories," that determined their mental model and interpretive stance towards the platform [36]. If users' mental models do not align with the design, or if technical constraints arise behind the scenes of the product, then designers of these platforms often employ some degree of what Adar et. al. call deception in the user experience [32]. For example, when Netflix's servers are overwhelmed, it seamlessly switches from a personalized recommender system to a simpler one based on general popularity heuristics, unbeknownst to

the user due to the lack of perceived change in the visual experience [32]. Hiding or abstracting away critical information is an important design technique that can be used both to improve users' perception as well as detract from it. We now turn to what makes a particular design technique persuasive as opposed to influential in another, perhaps more negative way.

2.4 Metaphysics of Related Modes of Influence

The study of "influence" is defined as the study of the cause of human change in belief, attitude, and behavior [37]. There exist several forms of symbolic transaction besides persuasion in which one tries to influence others' behaviors and attitudes: for example, manipulation and coercion. Our definition of persuasive technology ultimately casts a wide net across several of these related concepts. However, for completeness, we review here the ways they have been defined in the literature, drawing from social psychology, political philosophy, and philosophy of action [38].

2.4.1 Compliance. Historically, compliance has been defined as inducing change in a behavior [39]. Rhoads et. al. describe it as a "quick-fix solution to a social problem" since it does not require user agreement at a rational level and merely demands the performance of a behavior [39].

2.4.2 Conviction. Some distinguish techniques that achieve persuasion from those that achieve conviction, saying that persuasion relies more on "symbolic strategies that trigger the emotions of intended persuadees," whereas conviction relies on "strategies rooted in logical proof and reasoning" and that appeal to persuadees' reason and intellect [26]. In other words, persuasion is said to cater more to irrationalities whereas conviction caters to rationality.

2.4.3 Education. Education has been defined as inducing a change in beliefs, and is similar in content and technique to propaganda; if we already believe in the target beliefs, we tend to refer to it as education, whereas if we do not, we call it propaganda [39].

2.4.4 Deception. One working definition of deception as it applies in human-computer interaction is when "an explicit or implicit claim, omission of information, or system action occurs that is mediated by user perception, attention, comprehension, prior knowledge, beliefs, or other cognitive activity, and creates a belief about a system or one of its attributes that is demonstrably false or unsubstantiated as true, where it is likely that the belief will affect behavior or a substantial percentage of users" [32]. In philosophy, deception has been defined as "outright lying to those manipulated, including making false promises to them, but also misleading them without actually misrepresenting anything, such as by encouraging false assumptions or fostering self-deception that is advantageous to the manipulator's ends," or getting the target to "interpret the situation in a light favorable to the manipulator's purposes" [38].

2.4.5 Coercion. Coercion is distinguished by its degree of constraint; when a performed action is characterized by the strong lack of an acceptable alternative, a user is coerced (or "compelled" or "forced") to do it [38]. However, mere constraint in itself is different from coercion, since constraint can exist even when users have several acceptable options, such as during automatic speed limit enforcement [40].

2.4.6 Seduction. Verbeek, a philosopher of technology, defines seductive technologies as those that "do not so much coerce people or persuade them to act in a certain way at the cognitive level, but that simply make some actions more attractive than others" [40].

2.4.7 Manipulation. Manipulation has been characterized both as "deceptive non-coercive influence" and as "non-rational influence" [38]. Barnhill defines manipulation as directly influencing someone's beliefs, desires, or emotions such that she falls short of ideals for belief, desire, or emotion in ways typically not in her self-interest or likely not in her self-interest in the present context" [38]. When influence leads to beliefs or desires that are in someone's self-interest, it is then not considered manipulative. Manipulation can also be indirect. For instance, evoking an emotion that does not directly induce a self-harming action but rather evokes a state of mind that leads to decisions that are not directly in self-interest is also manipulative: "if [a manipulator] appeals to a cold, competitive malice that's under his control, then she does not manipulate him, because feeling controlled, competitive malice, and making decisions on the basis of it is likely to be in his self-interest. However if she stokes a hot competitive malice that's not tightly controlled, then she does manipulate him, since feeling and acting on not-tightly-controlled malice is typically not in one's self-interest" [38]. Barnhill further classifies types of manipulation as intricate vs. blunt, paternalistic vs. non-paternalistic, covert vs. overt, targeting beliefs vs. emotions, and changing a situation vs. a person.

These modes of influence are interrelated in interesting ways; the use of some can amplify or lead to the use of others, and some can be seen as more intense applications of others. Wood considers coercion and manipulation to be on a continuous spectrum: "coercion destroy[s] free choice" by rendering all but one option unacceptable, whereas manipulation merely influences choice without removing it" [38]. However, some see the two as fundamentally different, since a coercer typically alters the world or the product, such that seeing the altered worldview makes it rational to do what the coercer wants, "by the target's own lights" [38]. Manipulation, by contrast, aims to alter the target's viewpoints (beliefs, desires, and values) "while being indifferent to whether the alterations reflect what is true or desirable," - acting more as a "puppeteer [...] messing with our heads" [38,41].

Faden and Beauchamp contrast persuasion and manipulation by saying that "persuasion improves someone's understanding of her situation, but manipulation does not" [38]. They even define manipulation in terms of a lack of the cognitive process of understanding as "any intentional act that successfully influences a person to belief or behavior by causing changes in mental processes other than those involved in understanding" [38]. Mills claims that manipulation disguises itself as good persuasion by appearing to offer valid reasons, but in fact offers faulty ones, even though the manipulator knows the reasons to be faulty. He defines manipulations as a "persuasion manqué, as an attempt at internally directed and non-physically-based influence that deliberately falls short of the persuasive ideal" [38].

How, then, has persuasion been historically characterized? Social psychology traditionally defines persuasion as the process that changes attitude [37]. Persuasion attempts to win "the heart and mind" of the target, which involves affective change; it is also characterized by more long-lasting effects than other forms of influence such as compliance, since the target accepts and internalizes the message to a greater degree [37]. In human-computer interaction, persuasion has been described as "a voluntary attempt to change attitudes or behaviors or both," whereas coercion implies force and deception implies misinformation [19]. Persuasion comes from someone else's conceptual framework" [20] and "involves specific intent from an outside agent"

to cause the target to adopt "previously untenable beliefs, attitudes, or behaviors that are foreign to their own conceptual and behavioral repertoire" [20]. In this spirit, we define persuasion as follows:

> ***Persuasion:*** *the mode of influence in which one agent (for our purposes, a digital technology acting as a vehicle for the designer) intentionally attempts to change another agent's (the user's) attitudes or behaviors.*

James Williams has identified even more modes of influence that persuasive technologies can have over users, which vary in degree of goal alignment and constraint; these include technologies that "demand, drive, tempt, guide, invite, suggest, and direct" [42]. Moving forward, we do not discard persuasive techniques that fall under metaphysical categories besides persuasion, and instead use the term persuasion as a broad umbrella to capture techniques across these related modes of influence. We default to the singular term of "persuasion" because many of the other modes of influence have built-in normative connotations, and we aim to provide a descriptive and relatively non-judgemental account of persuasive design, especially in **Chapter 4**. However, we revisit the distinction between different kinds of influence in **Chapter 5**, where we discuss the questions of what constitutes ethical persuasion and how to design for more ethically permissible forms of persuasion.

2.5 Persuasive Strategies

Designers of our physical world often employ persuasive strategies to encourage particular attitudes or behaviors. Although the approach of using interpersonal or non-digital persuasion to understand digital persuasion has largely focused on computer-mediated communication rather than human-computer interaction [32], it is still insightful to review examples of real-world persuasion, since their analogues are present in digital design as well. For example, architectural elements in casinos such as maze-like paths, the illusion of small, secluded spaces, and lack of apparent exits persuade people to stay and continue spending money [32]. Visual cues indicating food portion size affect intake of items like Campbell's soup and popcorn; by removing visual cues that indicate how much has been consumed or when to stop consuming, people engage more in activities that are considered mindless [43]. Park benches with central armrests discourage overnight occupation by the homeless [44]. "Placebo buttons" such as those found in crosswalks, elevators, and thermostats often provide the illusion of control without actual functionality [32]. Dual-button toilet flush controls steer users to choose between two behaviors - one that conserves water and one that does not [32].

Underlying these non-digital persuasive design techniques are well-known persuasive strategies that have been studied and categorized in social psychology. Kelton Rhoads conducted a review of landmark compilations of persuasive strategies published since the topic gained prominence due to the advent of mass media [39]. What constitutes a persuasive strategy? Marwell and Schmitt define it as "the reduction of the multitude of possible behaviors into meaningful clusters" such that a strategy is a "group of techniques towards which potential actors tend to respond similarly" [45]. Rhoads found varying numbers of strategies; however, these distinctions are often not rigid and quantitative but rather stem from the differences in level of social psychology abstraction. For example, Marwell & Schmitt identify several strategies in their 1967 taxonomy of the dimensions of compliance-gaining behavior: reward, punishment, positive expertise, negative expertise, liking/ingratiation, gifting/pre-giving, debt, aversive stimulation, moral appeal, positive

self-feeling, negative self-feeling, positive altercasting, negative altercasting, positive esteem of others, and negative esteem of others [46]. Sequential-request compliance theory groups strategies into two kinds: foot-in-the-door (FITD): small initial request followed by a larger, but still reasonable, second request, and door-in-the-face (DITF): first outrageously large request, followed by a smaller, more moderate request [47]. Levine & Wheless identify 53 tactics in a 1990 review of previous taxonomies [48], listed in **Figure 1**.

Robert Cialdini has taken a social psychology and communications approach to persuasion, breaking down the psychological laws that govern persuasion into "pre-suasion" - the factors that establish trust and set the tone for the persuasion - and six psychological principles (reciprocation, liking, social proof, authority, scarcity, and consistency) that "represent certain psychological universals of persuasion" [49].

FIGURE 1 Compliance-gaining Tactic Origins

Tactic	*Tactic*
1. Allurement	28. Justification for action
2. Altercasting (–)	29. Liking
3. Altercasting (+)	30. Manipulation
4. Altruism	31. Modeling
5. Audience use	32. Moral appeal (–)
5. Aversive stimulation	33. Moral appeal, (+)
6. Coercion	34. Negotiation
7. Communicator image (+)	35. Non-negotiation
8. Cooperation	36. Normative rules
10. Debt	37. Personal responsibility
11. Direct request	38. Personal rejection
12. Disclaimer	39. Pre-giving
13. Duty	40. Pressure for action
14. Emotional appeal	41. Promise
15. Empathic understanding	42. Punishment from others
16. Esteem (–)	43. Reason
17. Esteem (+)	44. Receiver image (–)
18. Expertise	45. Receiver image (+)
19. Expertise (–)	46. Relationship (–)
20. Expertise (+)	47. Relationship (+)
21. Extended expertise	48. Self-feeling (–)
22. Explanation	49. Self-feeling (+)
23. Feedback	50. Simple statement
24. Flattery	51. Suggestion for negative alternatives
25. Guilt	
26. Inaction	52. Threat
27. Ingratiation	53. Warning

Figure 1. Levine & Wheless's 1990 review of prior taxonomies of persuasive strategies lists 53 different compliance-gaining tactics (figure adapted from [48]).

The pre-suasion strategies include:

1. **Priming via context:** focusing recipients initially on concepts that are aligned, with forthcoming information, such as playing a German song in a store to make people buy more German products.
2. **Associations irrelevant to goals:** using priming to elicit unsavory behavior, such as exposing subjects to violent language to increase the number of shocks administered.
3. **Mood contamination:** using earlier information to influence affective state relevant to later questions or tasks, such as how asking a question about happiness about social life influences response to overall happiness levels.
4. **Agenda-setting theory:** drawing on the fact that we assume earlier information points out what to consider important in future information.
5. **Investigatory reflex:** using distinguishing factors and self-relevance (e.g., photos of self), unfinished arguments, and unsolved problems to draw attention.

The six key psychological factors in persuasion [49] include:

1. **Reciprocation:** our tendency to want to repay the favor.
2. **Consistency:** our need for internal self-understanding, or our tendency to observe our own behaviors to learn about ourselves, which can trump rationality, and is often enforced through commitment.
3. **Social proof:** using the actions of others to decide on behavior for ourselves, especially when we view others as similar to us.
4. **Authority:** deferring to authority, even in outrageous situations, such as administering shocks to other experimental participants due to the instructions of an experimenter.

5. **Scarcity:** the fear of losing out on access to a scarce resource.
6. **Liking:** attractiveness, which often leads to a massive halo effect where associated aspects of an attractive prospect seem more attractive. For example, aesthetically pleasing designs and faces, including online avatars, elicit more positive response than others, regardless of function [52].

From an HCI perspective, Adar categorizes types of deception by functional, behavioral, and mental model strategies which respectively misrepresent performance information, take advantage of psychological or sensory limits of perception (e.g., Fitt's Law is applied when a drop-down menu bar that is programmed not to roll back as the user moves a couple pixels outside of the box), and use metaphors to mislead users into thinking something works differently than it actually does [32]. Once users develop a "relationship" with the product, they become more susceptible to truth bias (believing the computer to be truthful), which opens the door to further deception; this is because users are less adept at placing themselves in the "mindset" of a technological system as they might be with another human [32].

There is also empirical work on persuasion that tests the success of different methods on controlled populations. For example, computer-generated graphics are more persuasive than text as a form of computer-mediated communication [50]. In environmental conservation settings, focusing persuasive messages on descriptive normative information increased undesired behavior, whereas focusing messages on injunctive normative information suppressed it [51]. High source credibility has been shown to lead to more favorable attitudes than low source credibility when people agree with the message or perceive it as having stronger arguments without knowing the source, but high source credibility has the opposite effect when people initially disagree with the message [52].

2.6 Persuasive Design Techniques

For better or worse, products in the attention economy employ persuasive techniques in their design that are similar to the techniques that are used in the non-digital world. Digital products use these techniques virtually everywhere, from the font colors to the text that frames anything from profile completion cues to promoted posts [53]. These techniques have been called subtle and "Pavlovian," and include things like "the red badge on Facebook's app that signals a posted comment or message; YouTube and Netflix videos that automatically cue up to play one after the other; and the Snapchat orange fireball emoji that signals a streak of daily chats between friends" [54]. Now entire products exist to help websites optimize for where to place their content to get the most clicks and shares [55].

While these techniques add utility for users to some degree, some techniques are considered more insidious than others. One collection of more unsavory techniques called "dark patterns" includes "bait and switch, confirmshaming, disguised ads, forced continuity, friendspam, hidden costs, misdirection (purposefully focusing your attention on one thing in order to distract your attention from another), price comparison prevention, sneak into basket, and trick questions" [124]. In fact, LinkedIn's old version of the onboarding flow to "add a contact" spammed the user's contacts without explicit permission, which led to the class action lawsuit *Perkins v. LinkedIn* in 2015 and allowed affected users to submit a claim to get a payout of around $10 each [124]. But some techniques are less contested and even more subtle: "by making use of just-noticeable differences (JNDs), developers can create the illusion that an action has, or has not, been influenced" [32]. Slight variations of these techniques appear throughout all products in the attention

economy, due to the nature of the attention economy itself. For instance, examining the Twitter home screen reveals techniques that are present across all SNS (**Figure 2**). Once one product starts, for example, autoplaying videos, it is more likely to increase users' time spent on its platform, and since attention and time are limited resources, competing products adopt successful design tricks from each other.

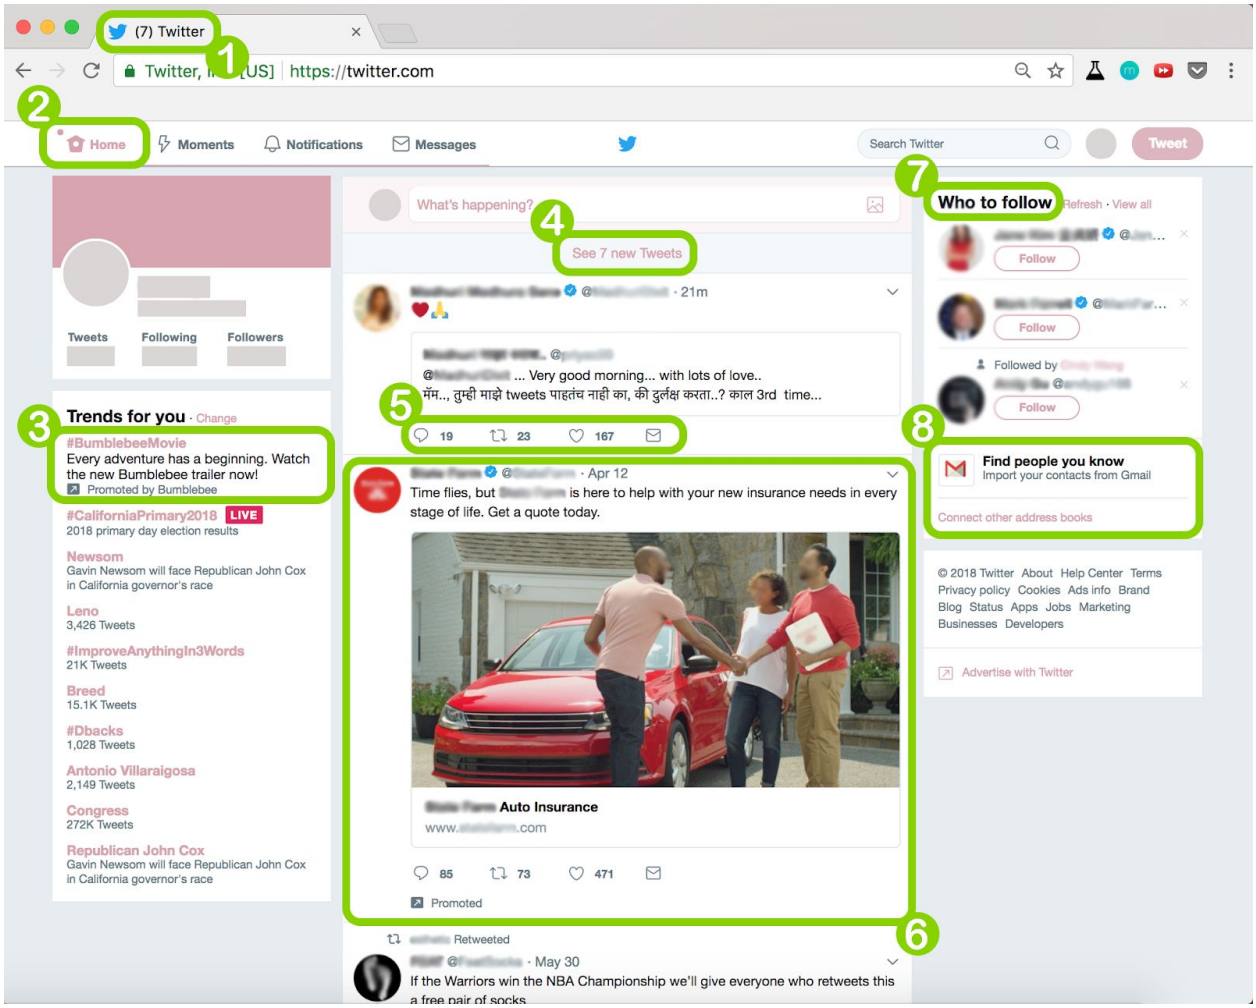

Figure 2. The home screen of Twitter, annotated with common persuasive design techniques that are also used in other products: 1) counts of unread messages, 2) asymmetric dot buttons on the "Home" tab, 3) artificially promoted trends at the top of the trends list, 4) text quantifying the amount of content to catch up on, 5) animated buttons that fill with color as users click them to share content, 6) large ads formatted to look similar to other posts, 7) customized suggestions of more people to follow, and 8) one-click contact imports from other SNS.

It is interesting to consider the question of how many persuasive design techniques there are. We gather data on what users of these products consider persuasive design techniques in **Chapter 3**. It is even more interesting to consider how these techniques are related. We address these questions in **Chapter 4**.

3 Nudget: System Design for User Awareness

Modern social media websites routinely employ persuasive design to orchestrate specific user actions. To raise Facebook users' awareness about the behavioral design techniques and tricks used on their cognitive resources, we developed a system, Nudget, that makes such techniques visible during Facebook use. Nudget (named loosely after the idea of "nudging" user awareness through micro-interventions) annotates desktop Facebook feeds in real-time, offering information about relevant persuasive theory at play. We conducted a 5-day trial in which 15 participants used Nudget while 6 participants used traditional state-of-the-art print material to learn about persuasive design techniques. We found that Nudget users displayed a significant difference in post-intervention scores of learning about the techniques compared to the control ($p=0.01$). We discuss the implications of our system on the effort to educate social media users about the web design calculated to guide their online behavior.

3.1 Introduction

To set and accomplish goals, we must have control over our own attention. Modern social media users lose agency (the capacity to act completely independently and make free choices) because social media platforms implement persuasive design techniques to orchestrate target behaviors that improve the websites' bottom line. In particular, they employ bottom-up visual and content-driven techniques to drive up their engagement metrics [8]. It is important to recognize, measure, and think critically about such manipulations and their effect on cognitive resources, but we lack modern vocabularies, user awareness, and transparency to do so. This motivates us to design a system that elucidates the underlying social and psychological tricks driving users' interaction with social media. By drawing back the curtain to show users how websites conceive of the purpose of their design features, we aim to educate users about such techniques to make them more aware for future social media use.

Past studies have shown human susceptibility to social media persuasion, and a number of systems exist to encourage users to take control of their time spent on social media. MeTime was a system that spread user awareness by offering users a glimpse at their pie chart of time spent on various SNS – i.e., how big of a slice of their pie each social media company owns [57,58]. The MeTime system showed that pure awareness helped users better manage their time and grow stronger against social media manipulations. Another system LivingSmart, showed that it is possible to use technology like focus-supporting browser extensions and task management applications to prevent digital attention manipulation even with participants with attention deficit disorder [59]. These systems demonstrate the viability of using technology to empower users against the platforms that reduce their agency to complete their own tasks and achieve their personal goals.

We extend this work by expanding the design space of our interventions beyond post-use reflection of time spent on social media sites. Current systems can tell people post hoc that they have been manipulated and provide them with graphical displays of their time consumed by corporations. Existing solutions can also block entire sites or remove news feeds, but this option is not feasible for many social media users. These systems are not able to tell users when and how they are being manipulated in real time, even though it is not the sites themselves but rather the attention-grabbing tactics they use that lead to user attention manipulation. Our system is granular enough to highlight manipulation in real time and in visual and

temporal proximity to the very manifestations of persuasive theory that manipulate user attention. By giving users' real-time, relevant insight into the underlying persuasive theory behind the features they use, we aim to improve knowledge retention about what happens behind the scenes of what they see on the screen. Although we believe educating users about the inner workings of persuasive technology platforms is a valuable end in itself, such interventions can also provide us with intuition about which persuasive design techniques users find most memorable and problematic.

3.2 System

Nudget is a Google Chrome extension that annotates users' Facebook feeds on their desktop computers, showing informational popup messages about persuasive theory at play in-situ. We discuss the design space of our Nudget intervention messages, the design principles we considered when building our display, and the implementation of our system.

3.2.2 Persuasive Theory Design Space

We now describe the subset of persuasive theories that inform the text of our Nudget interventions. Instead of exposing issues about time distributions or mental health disorders on the site as covered in previous work, we focus on situating our interventions in the larger space of persuasive technology design. We draw upon Robert Cialdini's research on the effects of persuasive framing on behavior [49] as well as Nir Eyal's and B.J. Fogg's persuasive technology design loops [60,61]. Taken together, these theories suggest that all the manipulations we are interested in occur at some point along the following model that describes user behavior: *motivation, ability and prompt* [64]. Motivation in a digital context refers to sensation, anticipation, and belonging (emotional needs), ability refers to reducing the barrier to access for the target behavior, and prompt refers to deadlines or cues to make the target behaviors urgent. For example, getting and giving likes, as well as the "Show Comments" feature is an instance of *motivation* – the social and emotional factors relating to the need to be heard and make themselves heard, which motivates people to generate and share content. Autoplaying a video from a post is an example of lowering *ability* needed to complete the target behavior of viewing the video, because it reduces the number of clicks needed from 1 to 0. The scrollable ever-updating side feed is a *prompt* – a deadline for reading information now before it is forever swept away to disappear. The "Home" and "Notification" buttons are landing points for prompts (such as the bright red numerical count of unread updates) because they are affixed to the top of the page and always noticed. The action flow from motivation to ability to prompt gets people to repeat the behavior on a fixed schedule. This is described by the following cognitive flow: acquisition using external prompt, creating desire, and affixing the internal prompts; in other words: attention comprehension, elaboration, integration, and enduring attitude change. For example, people come to Facebook out of habit, boredom, or loneliness (prompts external to the site), and after the login, have desires created by the motivators present onsite (creating desires/elaboration), have these desires fulfilled by manipulations onsite (integration), which completes the cycle of these created desires being affixed to internal prompts, leading to enduring attitude change over time.

3.2.3 Design Principles

To teach users about these complex persuasive theories, we broke down the interconnected webs of manipulation into feature-specific instances of persuasion. Please see **Appendix A** for our final list of informational Nudget messages (called "Nudgets"). In designing our messaging, we wanted to balance a desired level of knowledge with short attention spans and the rapid scrolling behavior users display on Facebook. We also did not want to disrupt the natural Facebook experience by demanding extensive reading, to the extent that users would uninstall our extension. Via informal usability testing, we determined that users would tend to read ~3 Nudgets on one screen. We also realized that the optimal text length of Nudgets would be about Tweet-sized (<280 characters), and we included a two-word tagline on each Nudget to allow for rapid reading. A drawback inherent in our design is that we, too, use a bottom-up attention-grabbing technique when displaying these bright pop-up messages. However, we found in usability testing that without the "pop-up" animation or the bright yellow color, users would not even notice our messages, due to the variety of elements already on their Facebook feeds as well as their habit of focusing intensely on the central feed itself. Thus, we employed our own subset of persuasive design techniques to even be noticed in the vicinity of Facebook's display. We anticipated that this could make Nudgets slightly disruptive for some users even while educating them.

3.2.4 Implementation

We built the Nudget system on top of the HabitLab project out of the Stanford HCI Group [62]. HabitLab itself is a Chrome extension that helps users gain control of their browsing. We implemented our custom intervention in Javascript and jQuery, packaged it into a fork of HabitLab, and deployed it to the Chrome store for our participants to install.

Upon a user scroll event on the Facebook feed, the system runs "Nudget" as long as "Nudget" has not already run within the last 4 hours. The system tracks divs of the various Facebook manipulations and annotate those divs with 3 Nudget descriptions of the manipulative technique. These bite-sized chunks of information appear with arrows pointing to the element on the screen to which the message refers. The annotations appear for 15 seconds, one by one, in the form of open popup boxes, and then close up and move to the top left corner of the screen, where the arrows in the messages change into numbers to avoid spatial mislabelling as users scroll. Nudget annotations are clickable so that users can toggle an annotation between open and closed states.

Some examples of Nudgets are shown in **Figure 3** and **Figure 4**. The first uses jQuery selectors to find the notification button and then display the following text next to it: "Quantified FOMO: Facebook shows a number to tell you exactly how many "important" things you're missing out on. What % is actually important?" (**Figure 3**) The second uses selectors to find the small access point to settings in the top right of the Facebook page and annotate the following text there: "Don't Leave: Compare how easy it is to see your feed vs. deactivate your account in settings. Facebook purposefully makes it hard for you to exit." (**Figure 4**). For a close-up of another contextual Nudget, see **Figure 5**. To try Nudget on Google Chrome, download it here: bit.ly/nudget.

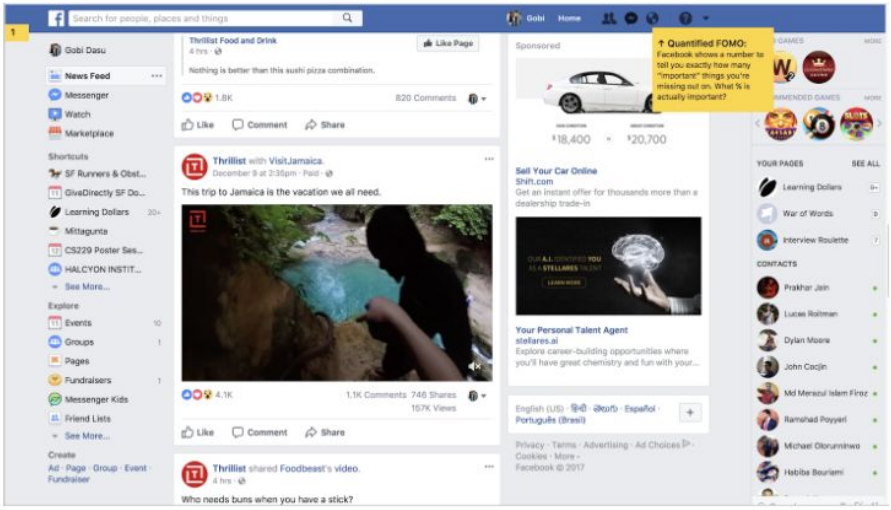

Figure 3. (left) An example Nudget that displays the following text below the notifications icon: "Quantified FOMO: Facebook shows a number to tell you exactly how many "important" things you're missing out on. What % is actually important?"

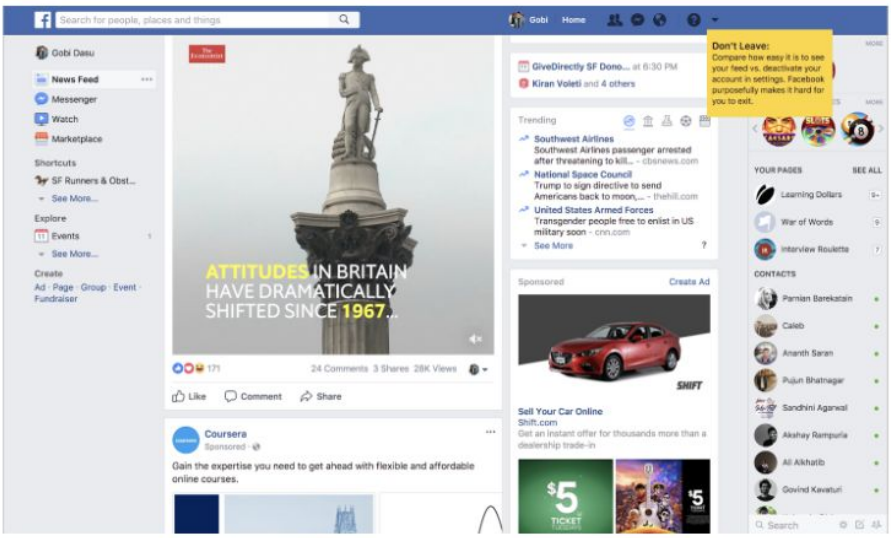

Figure 4. (left) An example Nudget that displays the following text below the top drop-down menu "Don't Leave: Compare how easy it is to see your feed vs. deactivate your account in settings. Facebook purposefully makes it hard for you to exit."

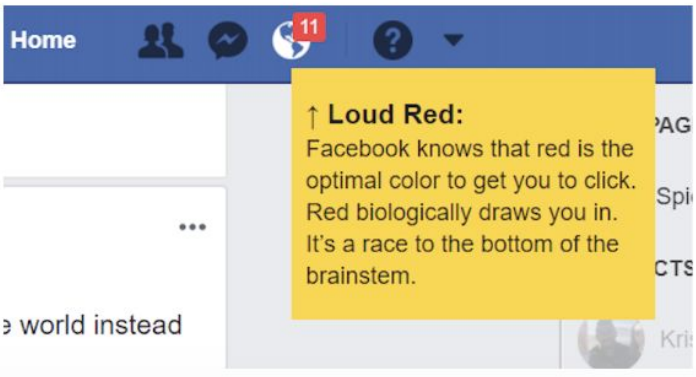

Figure 5. (left) A close-up of how a Nudget would appear to a user in context as they scroll through their Facebook feed.

3.3 Evaluation

To evaluate the effectiveness of the Nudgets as a system, we conducted a 5-day between-subjects study with 15 participants in the treatment group and 6 participants in the control group. We recruited participants through our personal networks and HCI class. Both the treatment and control groups completed two questionnaires at the beginning of the study. The first questionnaire was a survey about their usage and opinions of social media websites. The second questionnaire (see **Appendix B**) was an open-ended quiz that asked participants to identify manipulation techniques on three screenshots of LinkedIn, a social networking site that overlaps with Facebook in the design techniques used.

After taking these surveys, the treatment group used the Nudget system for a week on Facebook while the control group read articles on persuasive techniques used by social media websites. At the end of the study, both the treatment and control groups were given the same two questionnaires again. We hypothesized that if our system were successful, it would increase the number of techniques identified by people in the treatment group after the intervention significantly more than the number identified by people in the control group.

We used the first questionnaire to control for pre-existing attitudes towards social media sites (explained further in the next section). We used the second questionnaire to measure changes in knowledge of persuasive techniques. Specifically, for each participant, we calculated the difference in the number of relevant persuasive techniques they identified across the screenshots of LinkedIn before and after the intervention; we used this difference to calculate two scores of awareness: 1) the Nudget-specific awareness assessment score, which we obtained by dividing the difference by 22 (the size of the set of persuasive techniques covered in Nudget interventions), and 2) the general awareness assessment score, which we obtained by dividing the difference by 171 (the size of the pool of all 171 unique persuasive design techniques identified by the participants).

In this way, since the participants were instructed about the use of the techniques on Facebook and tested on examples from LinkedIn's, our calculated scores parameterized "transfer knowledge" - the knowledge that transfers across platforms. The Nudget-specific score allowed us to grade participants in context of their learning with the system, and the general score allowed us to get a sense of their performance when graded against a more complete gold standard set of possibilities. See **Appendix D** for the full list of 171 techniques.

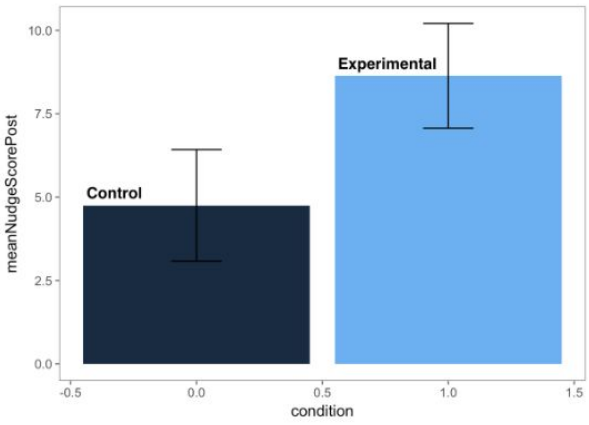

Figure 6. A t-test comparing post-intervention transfer scores between experimental and control conditions. The x-axis represents the mean score, and the error bars are 95% confidence intervals.

3.4 Results

We now present both quantitative and qualitative results measuring the effect Nudget had on our participants' awareness of design techniques.

3.4.1 Quantitative Evaluation

We performed two statistical evaluations to determine whether condition was a significant predictor of transfer knowledge scores.

3.4.1.1 Transfer Knowledge Evaluation I

We found that the condition participants were placed in, experimental or control, mattered significantly in our assessment of transfer knowledge. Namely, a t-test yielded a significant difference (p=0.01) between the post-intervention transfer scores of the experimental group, which was exposed to Nudget, and the control group (**Figure 6**).

We obtained the post-intervention scores used in this test by grading against knowledge of manipulations presented by Nudget, not against the general answer pool. While we also computed a post-intervention score graded against the general answer pool, the system was not able to significantly improve performance for the experimental group. The t-test between post-intervention general transfer scores of the two conditions had a p-value of 0.4.

3.4.1.2 Transfer Knowledge Evaluation II

We also performed a regression to predict participants' post-intervention Nudget-specific awareness assessment score using the following controls: age, how much users liked/disliked Facebook, condition (experimental vs. control), pre-intervention general score, and pre-intervention Nudget-specific score. We found that condition was indeed a significant predictor (p=0.023, **Figure** 7). However, when we regressed post-intervention general awareness assessment scores, we did not find that condition was a significant predictor (p=0.056, **Figure 8**).

	Estimate	Std. Error	t value	Pr(>|t|)	
(Intercept)	-1.29739	6.49184	-0.200	0.8466	
age	0.09795	0.27097	0.361	0.7271	
condition	3.12818	1.11943	2.794	0.0234	*
facebook_likingLike	-2.19125	1.29826	-1.688	0.1299	
facebook_likingNeither like nor dislike	0.09573	1.47942	0.065	0.9500	
GenScore.pre	0.11772	0.05014	2.348	0.0468	*
NudgeScore.pre	0.58947	0.22616	2.606	0.0313	*

Figure 7. (left) Regression predicting post-intervention Nudget transfer score, with control variables including age, condition, pre-existing like or dislike of Facebook, and both pre-intervention assessment scores.

	Estimate	Std. Error	t value	Pr(>|t|)	
(Intercept)	-4.5865	20.6464	-0.222	0.829766	
age	0.5192	0.8618	0.603	0.563503	
condition	7.9476	3.5602	2.232	0.056091	.
facebook_likingLike	-0.5054	4.1289	-0.122	0.905596	
facebook_likingNeither like nor dislike	-1.8191	4.7051	-0.387	0.709122	
GenScore.pre	0.8057	0.1595	5.053	0.000986	***
NudgeScore.pre	-1.3225	0.7193	-1.839	0.103255	

Figure 8. (left) Regression predicting post-intervention general transfer score, with the same control variables as above.

We note that in **Figure 7** and **Figure 8**, the variable *facebook_liking* was split by our statistics software into two variables — one indicating whether the user liked Facebook or not (before the intervention) and the other indicating whether or not the user had no feelings toward Facebook.

3.4.2 Qualitative Evaluation

Participants had positive responses to the Nudget system. When asked if they would rather learn about the topics through Nudgets or reading articles, 66.7% responded that they preferred the Nudgets. Many participants said the interactive and personal aspects of Nudgets made it more relevant and enjoyable to learn about the topics. They thought that seeing concrete examples of persuasive design techniques in a product they used daily was more useful than reading an article. One participant remarked in the post-questionnaire, *"Having the Nudgets appear on Facebook made the experience more interactive, and it resonated on a more personal level for me. Sometimes when I read an article about these things, I tend to distance those experiences from my own experiences. Seeing the Nudgets appear on my personal account allowed me to see just how affected I personally am by Facebook's persuasive design techniques."*

However, many participants also had suggestions for ways to improve the Nudgets. Many said that the Nudgets were repetitive and reported seeing the same Nudgets many times. Others suggested that the Nudgets should be more context-aware, saying that sometimes the text of the Nudget did not correspond to where and when the Nudgets showed up. Finally, some said that using the Nudgets for an extended period would be a nuisance and would have *"diminishing returns"* after a few days of use because they had already read most of the Nudgets before.

Participants reported mixed results in regard to Nudgets changing their opinions and use of social media sites. Many people said that they already knew that social media sites used persuasive techniques but Nudgets emphasized and enhanced that knowledge, which is not surprising given that our participant pool was highly trained, and, in several cases, proficient in human-computer-interaction. Participants said that Nudgets made them more aware of the pervasiveness of these design techniques and how intentional every design choice is to maximize time spent on social media. However, most participants were unsure of how this knowledge would translate to action. 33% of participants in the treatment group said they would spend less time on social media; the rest said that the Nudgets prompted them to think more intentionally about social media, but they were unsure if they would actually spend less time on it.

3.5 Discussion

We found that Nudget was able to teach users about the manipulations that it was designed to cover (with significant results from both a t-test and regression); however, Nudget did not facilitate statistically significant generalizable awareness transfer learning. The fact that condition was not a significant predictor for generalizable knowledge transfer may have been because our system did not teach enough interventions common to Facebook and LinkedIn effectively. First, perhaps the duration of our longitudinal study was not long enough to repeat concepts enough times for users to gain higher level knowledge, process, meta-analyze, and retain. The second reason may be that even though we coded our system to choose from 20 manipulations at random, participants reported recalling only between 3 and 9 unique nudges. Certain divs like autoplaying videos or posts with comments open may not have showed up enough times to be

picked for annotation. A third reason may be because such a system has to run on more than one platform (not just Facebook) for participants to be able to generalize knowledge. Finally, the large size of the gold standard answer pool size may have made increases in awareness a tiny effect size.

Our results suggest that Nudget is most effective as an education tool that is used for a few days. It is clear that the Nudgets did increase awareness and knowledge of persuasive design techniques, validating our goal for creating Nudget. Nudget effectively demonstrated persuasive design examples in real time on Facebook, a website that the participants regularly visited. However, after the first few days of the study, many participants complained that the Nudgets were annoying and repetitive; for example, one remarked *"The Nudgets were super annoying because of where they were on the screen, their color, and their unavoidability."* Additionally, while Nudgets were an effective tool to increase awareness of persuasive design techniques, this did not translate into direct action. Many participants said that they did not foresee themselves spending less time on social media as a direct result of Nudgets. For instance, one noted *"It's hard to change such behaviors. I just feel more guilty about it now. Maybe that's a start."* Therefore, more work can be done to investigate how increased awareness can inspire action. We discuss additional future work to be done in **Chapter 6**.

3.6 Summary

In summary, we built and evaluated a system that makes persuasive design techniques visible on Facebook. We found that Nudget users displayed a significant difference in transfer learning scores compared to those who used traditional print and video materials. Given that Nudget improves transfer knowledge, it would be useful to investigate what other kinds of behavior change it can drive, especially compared to other interventions. Creating interventions like Nudget also highlights the need to develop a theoretical taxonomy of persuasive design techniques, which we will discuss in the next chapter.

4 Taxonomy of Persuasive Design Techniques

Having empirically studied user awareness about persuasive design techniques, we turn to a more theoretical question about these techniques: how do they interrelate to form a system of persuasion? We use the datasets of Facebook's and LinkedIn's persuasive design techniques gathered in our Nudget user study, along with theoretical insights from the literature, to make progress on this question. Specifically, we map persuasive design techniques onto a two-dimensional behavioral grid to situate them along different parameters pertaining to human behavior. We then identify additional dimensions of interest to flesh out a fuller design space within which to situate persuasive design techniques.

4.1 Motivation and Goals

James Williams notes that an area ripe for impact is "differentiating between different types of technology-driven persuasion by how well they are aligned to our goals and how much they constrain our actions" [42]. He points out the need for a better vocabulary in the attention economy: "*asserting* our freedom of attention means developing its conceptual and linguistic foundations" [63]. There exists a rich behavioral and social psychology literature on persuasion (discussed in **Chapter 2**), but we lack a mapping of modern digital persuasive design techniques onto this theoretical space [64]. Thus, in this chapter, we take a descriptive theoretical approach with the following goals in contributing to a taxonomy of persuasive design techniques. We want to more systematically categorize the different types of persuasive pulls we feel when we interact with these massively popular mass consumer products to:

1. Show how persuasive design techniques interrelate and function as a system to affect people.
2. Provide a more nuanced vocabulary for persuasive design techniques, drawing from the relevant academic disciplines.
3. Lay the foundation for pumping intuition about which kinds of persuasion we feel are problematic or ethical.

In this chapter, we focus on the first two goals; we address the third in **Chapter 5.**

4.2 Mapping Persuasive Design Techniques onto the Fogg Behavioral Grid

Although interest in taxonomies of influence has existed since Marwell and Schmitt's taxonomy of 16 techniques in 1967, no one has mapped the specific techniques used by mass consumer products in the attention economy onto this persuasive space. B.J. Fogg's 2009 paper "The Behavior Grid: 35 Ways Behavior Can Change" (later updated in 2012) identifies a specific two-dimensional theoretical space that categorizes different digital and non-digital user behaviors based on how often they occur and how novel or frequent they are [64]. Fogg notes that we lack a mapping of both persuasive theories and persuasive design techniques onto his behavior grid, and that this is an important task for future work. We focus here on filling this gap by mapping persuasive design techniques onto the grid. We focus on the design techniques themselves and not persuasive strategies discussed in **Section 2.5**, since as he notes, a single theory (e.g., operant conditioning) might extend across multiple rows and columns of his grid, reducing the usefulness of the analysis. For example, such strategies might include "allurement," "emotional appeal," or "guilt" from

Levine & Wheless's 1990 review of 53 tactics [48]. A major ambiguity in dealing with persuasive strategies would be accounting for all the different ways that a persuasive strategy can be used.

Our grid adapts from Fogg's 2012 updated 15-box grid [65] as well as his earlier 2009 version of a 35-box grid [64]. Fogg identifies seven rows that capture the schedule on which a behavior occurs: a one-time behavior, a one-time behavior that leads to an ongoing obligation (such as adopting a puppy rather than just playing with a puppy once), behaviors that last for a period of time, behaviors that occur on a predictable schedule, behaviors that occur on an irregular cue rather than a fixed schedule, behaviors performed at will, and behaviors that are always performed. He identifies five columns that capture the type of behavior change performed: a new behavior, an existing behavior, and increase in frequency or duration of the behavior, a decrease in frequency or duration of the behavior, and a behavior that stops being performed.

We removed the "behavior is at will" row, which Fogg defines as containing behaviors that the user can perform at any moment, primarily because these behaviors depend on the screen that the user is currently on. Due to the abundance of screens and user journeys on both Facebook and LinkedIn, clicking on a single button can take a user to a completely different page; the only thing that remains constant is the horizontal toolbar at the top of the web page that allows options such as "Home," and "Notifications," as well as the search bar, which allows users to independently type and navigate to their desired page. Many of these "at will" user behaviors are captured in the row that includes behaviors on a "predictable schedule." More importantly the one- and two-click user journeys allow many different behaviors to occur "at will." Furthermore, the use of "will" conflates the line between autonomous action and action that is prompted by a persuasive technique, for which there are no clear metaphysical distinctions due to the ultra-persuasive worlds into which such platforms place users. Fogg's non-digital examples of "behavior for a period of time" include exercising for 30 minutes and flossing for a longer duration. In the context of Facebook and LinkedIn, we took "a period of time" to refer to behaviors that take longer than a few seconds (e.g., scrolling for five seconds as opposed to clicking a button for less than a second).

What constitutes a persuasive design technique that we might map onto this grid? There is no standard definition for what a persuasive design technique is, and opinions vary on how many there are [66]. As opposed to persuasive strategies, of which there appears to be a finite set based on past literature, the number of design techniques is "large, perhaps uncountable" [66]. Every successful interface is designed to achieve a purpose, and guides the user to some degree to achieve a task. For this reason, from one perspective, almost anything could be considered persuasive as long as the design incorporates an argument or usage intention [44]. Following the standard definition from captology, a persuasive design technique occurs when a designer designs with intent to change an attitude or behavior [19]. Since it is difficult to survey the designers of these platforms for their intent about each specific design element without revealing potentially sensitive information about the inner workings of their platforms, our dataset of persuasive design techniques consists of 1) techniques we identified on Facebook while designing Nudget, and 2) techniques that subjects in our Nudget user study (described in **Chapter 3**) perceived to be persuasive on LinkedIn. These are listed fully in **Appendix A** and **Appendix C**. Fogg later defined persuasive design techniques to be "a specific implementation of a persuasive experience" [64]. The design techniques identified by our study participants are along the lines of what Fogg has in mind, as seen in the example he provides of Amazon's technique to offer free shipping for purchases over $25: "the timing of Amazon's offer and how it is worded is a design technique" [64].

We take a relatively liberal approach to defining what constitutes a persuasive design technique. In our dataset, anything from a monetization strategy (technique #10) to the color of a button to the text used under an icon could be considered the visual conduit of persuasion, under a broad notion of what persuasion is metaphysically. We call the visual element through which the persuasion occurs the *vehicle*. In our analysis, each persuasive design technique must have a vehicle as well as a *method of persuasion*: how it persuades. The same visual element can have multiple methods of persuasion, which we count as two different persuasive design techniques (e.g., techniques #6 and #7).

4.2.1 Case Study: Facebook

In **Figure 9**, we present the behavior grid populated with Facebook's persuasive design techniques, where the dark gray row labels indicate rows present in Fogg's 2012 condensed version of the behavior grid. We sort the persuasive design techniques identified while designing Nudget into boxes on the grid. We supplement these with a few additional techniques observed by tracing the vehicles of the Facebook behaviors that Fogg lists in his grid [64]. An example of one of these additional techniques is illustrated in **Figure 10**.

Facebook's design techniques appear to encourage primarily one-time behaviors and behaviors that occur on a cue. Since Facebook has such a massive culture of heavy daily active use, many of the persuasive design techniques employed by the website upon opening it prompt behaviors for users that automatically occur both on a predictable schedule in order to form habits with the product, as well as behaviors that occur on an unpredictable cue, which relies on the behaviorist theory of variable reinforcement rewards [67].

Creating this mapping also teaches us more about the behavior grid and its suitability for persuasive analysis. Many of the design techniques could belong in different boxes in the grid; for example, showing a number overlayed over the notifications icon at the top toolbar serves to remind users on a predictable schedule (every time they open the website or app, which, for 2 billion users is a habitual part of their monthly routine [8]). But during continued use of the website or app, the same notification style reappears on the top toolbar on a variable rewards schedule, which serves to increase the frequency, and also presumably the duration of users' notification-checking behavior. Clicking on a notification, in turn, takes them to other kinds of posts and events on Facebook, where they then get transported to the boxes in the behavior grid relating to liking, commenting, and other engagement behaviors. Thus, the entire user experience can be modelled as a journey across the behavior grid, where the initial behaviors determine the initial conditions of the trajectory, and the rest becomes a random walk that is a function of the elements and design techniques that the user is then exposed to, as well as any overriding behaviors the user performs outside of the grid to direct their own behavior agentially onsite in spite of the pulls of persuasive design.

Another example of a persuasive design technique that could lead to behaviors spanning multiple rows is the design choice of auto-enabling comments below each post and displaying all the comments sorted by an algorithm that places the most active conversations above those that are less active. This technique can have multiple behavioral effects: prompting a behavior of scrolling through comments, which could be thought to either always occurs (bottom row) since every post enables comments, or occur for a span of time greater than a few instantaneous seconds. Similarly, notifications that prompt you to wish your Facebook friends a happy birthday trigger a behavior that occurs on cue (the prompt is the notification); however, this is also a behavior that creates an ongoing obligation, due to the human tendency to reciprocate birthday wishes. The mysterious newsfeed feed algorithm seems to rank posts by usefulness, which could be seen as decreasing the behavior of browsing useless posts; however, it also keeps people scrolling via variable

reinforcement by providing mildly interesting posts on an unpredictable schedule, which increases time spent on site.

Legend - Figure 9	
Color	Count
	1-5
	6-10
	11-15
	16-20
	20-25

		A	B	C	D	E
	Facebook	perform new/unfamiliar behavior	perform existing/familiar behavior	increase behavior intensity, frequency, or duration	decrease behavior intensity, frequency, or duration	stop a behavior
1	one time behavior	collection of personal and network information during onboarding, ad placement and display uniform with friends' posts, display of friends who have played each game, "You are now connected on Messenger" chat notification with new friends	"turn on chat" call-to-action, pre-populated "write a comment" text boxes adjacent to profile picture below posts	displaying posts that friends are tagged in		
2	one time behavior that leads to ongoing obligation/cost	profile-completeness cues (self-representation leads to envy spiraling), one-click "confirm friend" menu adjacent to newsfeed	comments as verbal approval, counts of likes/emoji reactions as social approval, friends' "interest" in events as social proof			buried menu to change security settings
3	behavior for a period of time	two-click option to play games with friends		unauthorized autoplay, always enabling comments, ability to display multiples chat boxes towards bottom of the screen, minimizing video to the right side of newsfeed to allow you to keep watching		
4	behavior on a predictable schedule		slot machine feed (variable reinforcement rewards), "read" timestamps on Messenger	glued reminders at the top menu bar, glued search bar		
5	behavior is on cue		poking reciprocity, intentional interruption via notification, tagging as social approval, one-click event join button	quantified FOMO, red color overlay alerts, notifications prompting birthday wishes		
6	behavior is always performed		mystery algorithm, micro-personalization, auto-population login information and profile picture	endless comments, "breaking news" side feed, bottomless bowl newsfeed	incompleteness of options in left menu	

Figure 9. Persuasive design techniques from Facebook mapped onto the Fogg Behavioral Grid; color legend above.

The variability in each individual user's familiarity with the product might also affect the box in the behavioral grid into which we map a persuasive design technique. For example, the row "behavior is always performed" differs based on how deeply the user has matched their mental model of the affordance the

platform provides with the exclusivity of the means to achieve it. For instance, teens using Snapchat often start to measure the strength of their friendships through the app's "streak" feature; they subscribe to the subculture of the platform, believing that the number of consecutive messages exchanged with a contact parametrizes depth of connection [88]. This extreme adoption of the platform's metrics can turn behaviors that, for some less invested users, might only be performed on cue (e.g., tagging your friend in a photo), into mandatory behaviors, due to the expectations of the social group with which they interact through the platform (e.g., 'liking' your friend's profile picture as a knee-jerk reaction to let them know you've seen it).

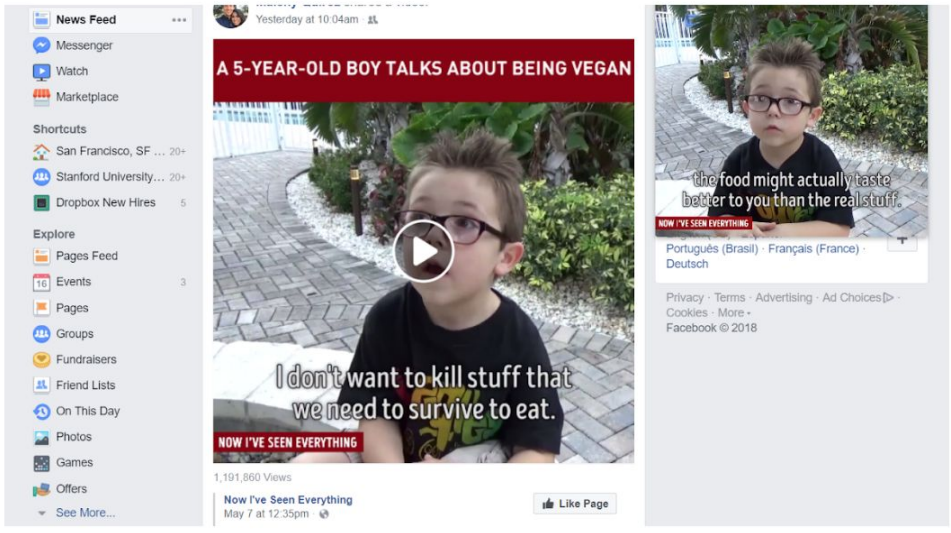

Figure 10. (left) An example of a persuasive design technique on Facebook: as you play a video and scroll past it midway, the video minimizes into a smaller window that stays fixed to the right of the feed, so that you can continue scrolling and seeing more content while also maximizing your watch time of the video.

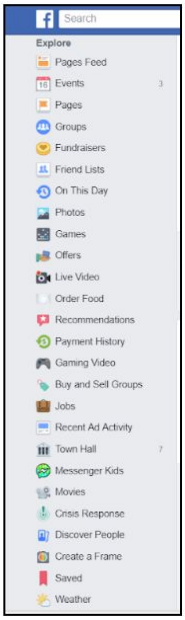

Figure 11. (right). The "Explore" menu pinned to the left side of the Facebook home screen includes over 25 possible actions to take. If a product becomes synonymous with a particular means for a user (e.g. Facebook as the means of entertainment or the means of interpersonal connection), then the range of options (or lack thereof) in such menus can more easily become synonymous with the scope of actions the user is conditioned to consider taking in order to reach the desired end.

The last two columns of Fogg's matrix raise two interesting questions; what does Facebook allow us to do less of, and what behaviors does Facebook stop? Fogg notes that "Facebook offered relatively few design techniques to persuade people to decrease or to stop a behavior, as shown by the lack of items in Columns D and E" [64]. As Fogg suggests in box D2, when users become flooded with notifications that are repeatedly subpar in quality, the persuasive design technique of providing notifications for a range of Facebook activities and updates may actually have the effect of training users to ignore certain notifications. Similarly, by including over 25 possible buttons to click in the left column navigation menu, pictured in **Figure 11**, the platform trains users to not direct their attention to the less used buttons, which frees up more attentional tokens to focus on the most useful (presumably the ones at the top). The menus that social media platforms include provide different means for connecting with people or engaging in activities together, but by reducing users' frame of possibility to the options on the menu, they decrease other

behaviors that users might take to connect with others, such as using other platforms to send messages, make calls, or hold face-to-face meetings.

4.2.2 Case Study: LinkedIn

We now map the 171 persuasive design techniques identified by Nudget subjects on LinkedIn onto our adaptation of the Fogg behavioral grid. Since they are generated via a close reading of three specific screenshots integral to a typical user flow on LinkedIn, they achieve a level of granularity that our more general observations of Facebook might not. As expected by the social nature of LinkedIn and its competitiveness in the attention economy, LinkedIn uses similar persuasive design techniques to Facebook. However, one important difference in LinkedIn's design is the subscription business model called LinkedIn Premium, in which users can pay for additional insights into their profile as well as expanded messaging capabilities, among other exclusive features [68]. This motivates LinkedIn to use strong persuasive pulls to convince users to actually put down money and buy their subscription in addition to using their product for free and paying via their time and attention. We studied the persuasive design techniques across three different screens on LinkedIn, the "Home" screen, the "My Network" tab, and the "Notifications" tab. The data we use here are only the persuasive design techniques that our Nudget study subjects noted and reported on their transfer assessments (numbered and listed in **Appendix C**), so they represent the elements most salient to users both before and after our system intervention. We divide the screens into "regions" along the natural grid lines of the web page for clarity in discussion, and we map the persuasive design techniques from each of the three screens onto the Fogg behavioral grid (screen 1: **Figures 12-14**, screen 2: **Figures 15-17**, screen 3: **Figures 18-20**).

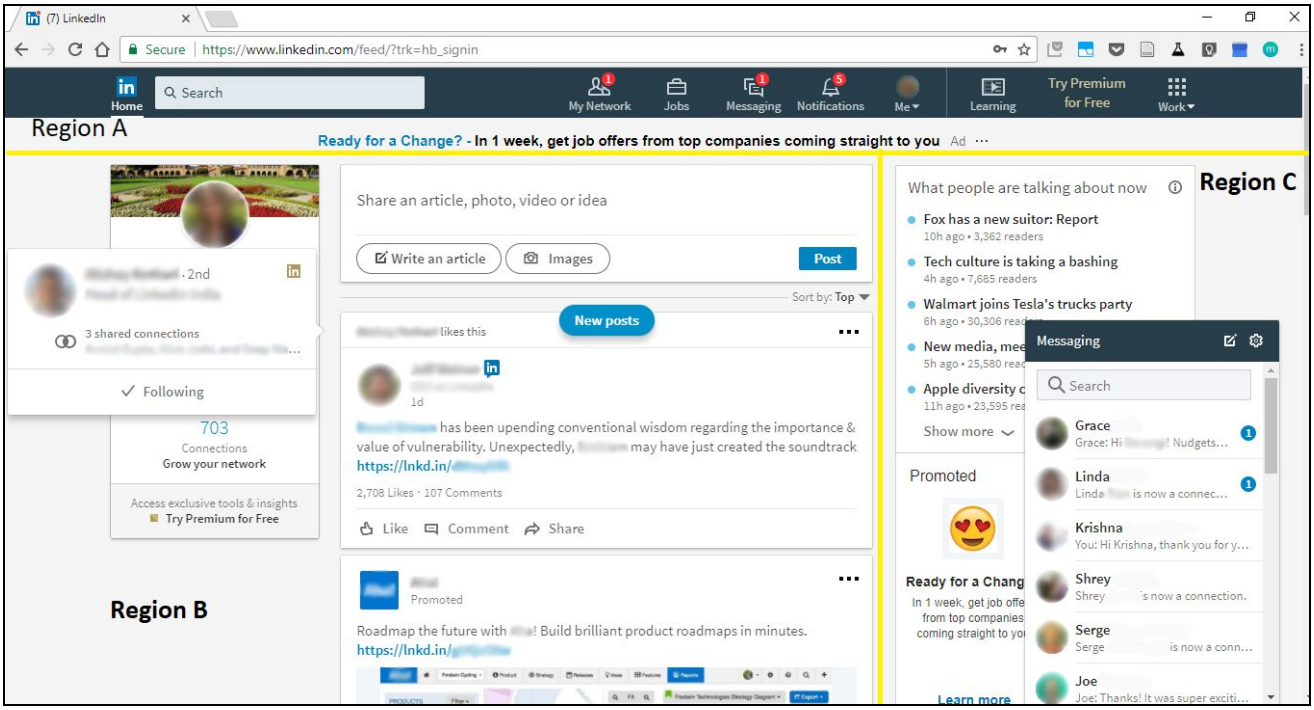

Figure 12. LinkedIn screen 1 appears upon login, and is comprised of a profile summary, social feed, news links, and messaging box.

Legend - Figures 13, 16, 19	
Color	Count
	1-3
	4-6
	7-9
	10-12
	13-15

		A	B	C	D	E
	Screen 1	perform new/unfamiliar behavior	perform existing/familiar behavior	increase behavior intensity, frequency, or duration	decrease behavior intensity, frequency, or duration	stop a behavior
1	one time behavior	5, 17, 21, 22, 24, 25, 41, 42, 44, 58, 65, 87, 89, 92	13, 14, 15, 16, 18, 29, 39, 40, 45, 53, 69, 73	34, 52, 57	11	
2	one time behavior that leads to ongoing obligation/cost	6, 7, 8, 9, 59, 60, 62, 63, 64	3, 19, 20, 30, 31, 36, 55, 76	32, 70		
3	behavior for a period of time		33	12, 35, 43, 68		
4	behavior on a predictable schedule		66, 67, 72, 77, 78, 79, 83	23, 80		
5	behavior is on cue	90, 91	1, 2, 4, 48, 51, 74, 75, 81, 82			
6	behavior is always performed	28, 71, 85, 88	37, 47, 49, 50, 54, 56, 84, 86	10, 27, 46, 61	26, 38	

Figure 13. Persuasive design techniques from LinkedIn Screen 1 mapped onto the behavioral grid; color legend above.

#	Persuasive Vehicle	Method of Persuasion
1	Notification badges on the horizontal toolbar for "notifications," "messages," and "network"	Makes you want to click and see new notifications (arouses curiosity)
2	Red color of notification badges on the horizontal toolbar	Stands out/catches your attention/indicates urgency in order to redirect your clicks to other people's or companies pages.
3	Number on the notification badges on the horizontal toolbar	Makes it feel like a to-do list and makes you want to get the numbers to 0 (arouses our "base desire for having order instead of chaos"
4	Intermittent variable notifications	The delivery schedule of notifications is varied and intermittent, which keeps it changing and thus interesting.
[refer to Appendix C for full list]		
89	Promoted ad	Encourages you to click on content that will make the site money.
90	Use of huge emoji in promoted ad	Distracts you from other useful information to encourage you to click on the promoted content.
91	Blue "learn more" text below promoted ad	The color blue makes it enticing to click on the promotion.
92	"In 1 week, get job offer…"	Enticing (and likely false) promises in advertisements encourage clicking.

Figure 14. A sample of the list of persuasive design techniques from LinkedIn screen 1.

As with Facebook, it is interesting to see the rare persuasive design techniques that fall into the last two columns. In region 1A, the "work" feature makes users think it will help them find work to do (design technique #11), but it turns out to be a collection of other LinkedIn features to use. Thus, it decreases the behavior of looking for jobs, which is what an unknowing user might think they are setting out to do, and instead redirects it to other features such as "learning," "advertising," and "slideshare." In region 1B, LinkedIn does not allow you to custom curate your feed (technique #38). Even though you can follow people (just as you can on Twitter and Facebook), you do not get to decide whose posts you want to see when, so you are forced to let the platform decide which posts you see first. This decreases the behavior of controlling which people or companies' posts you get to see, and does not let you actively prioritize what really interests you.

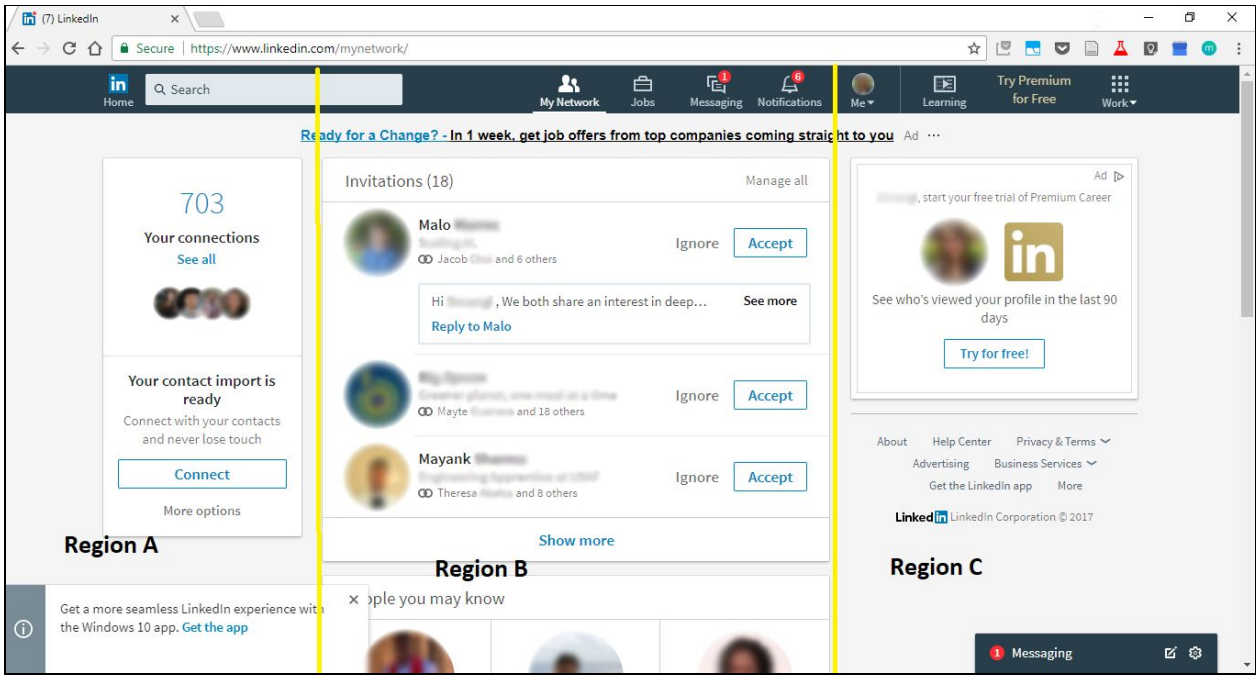

Figure 15. LinkedIn screen 2 appears upon clicking "My Network" on the top horizontal toolbar, and is comprised of a network summary box, pending invitations box, suggested connections feed, and Premium advertisement region.

		A	B	C	D	E
	Screen 2	perform new/unfamiliar behavior	perform existing/familiar behavior	increase behavior intensity, frequency, or duration	decrease behavior intensity, frequency, or duration	stop a behavior
1	one time behavior	116, 117, 123, 129, 131, 132	105, 109, 110, 114	111	122	97
2	one time behavior that leads to ongoing obligation/cost	94, 95, 96, 98, 100, 101, 136, 137, 139, 142	93, 103, 104, 112, 113, 127, 128, 140	121, 124	125	
3	behavior for a period of time	99	106			
4	behavior on a predictable schedule				133	
5	behavior is on cue	118, 119, 135	126, 130, 134		102	
6	behavior is always performed	138	107, 108, 141	115, 120		

Figure 16. Persuasive design techniques from LinkedIn Screen 2 mapped onto the behavioral grid; color legend on P.30.

#	Persuasive Vehicle	Method of Persuasion
93	Showing your number of notifications in the tab on Google Chrome	Attracts attention towards the LinkedIn tab by making you wonder about the significance of the number
94	Toast prompting user to download Windows app	Opens the door to moore notifications even while not on LinkedIn, allowing the site to pull you away even if you're doing something else and spend more time on LinkedIn
95	Text in toast prompting user to download Windows app	Use of words like "seamless" makes it seem like users will have a better experience if they get the app.
96	Placement of toast prompting user to download Windows app	Layout breaks the grid of LinkedIn since the toast crosses over into the next region, which gets us to notice it.
[refer to Appendix C for full list]		
139	Repeating the "Try for free" call to action	The recurring nature of this call to action (repeated left to right across the screen) makes it seem more urgent.
140	Echoed stylization of "try for free" button	The "try for free button" uses the blue and white stylization of the "accept" and "connect" buttons, which have already been associated with positive actions.
141	Default status of messaging tab as closed	Lets you see articles and ads and allow their persuasion to attract you first so you click on them
142	Highlighting the "golden" color associated with the premium membership	Brings to mind the "status" that the site has associated the color with

Figure 17. A sample of the list of persuasive design techniques from LinkedIn screen 2.

In screen 2, the majority of techniques serve to persuade users to engage in one-time behaviors that lead to ongoing obligations. Since screen 2 is all about growing users' network and amplifying the network effects that boost both the individual use of and the overall number of users of the product, there are many techniques in Row 2 to give users a reason to come back and use the site more often. For example, in region 2A, the use of persuasive language shows the ways LinkedIn appeals to key human needs on Maslow's hierarchy: "Never lose touch" (technique #102) and variations of the word "connect" and "contact" (technique #108). These phrasings guide users to equate LinkedIn with human connection, whereas true human connection is only a small part of what the site persuades users towards. By using the design pattern of highlighting the "accept" button in blue and the "ignore" button in gray (the latter using a word that alludes to a negative real-world behavior), the site dissuades the behavior of clicking "ignore" (technique #121). By making the count of notifications seem like a count of the number of action items left, it decreases the behavior of ignoring notifications and not acting upon them (technique #125).

Since the website's sources of income come from ad views and paid subscriptions, the design techniques that are plentiful in screen 2 are also the ones that have users perform new one-time behaviors: clicking on various creatively positioned and formatted ads that use your own profile photo (technique #90), juxtaposing an ad with your profile picture, and placing single lines of text permanently at the top of the feed). By de-emphasizing and making certain links less symmetric (e.g., "Help Center" and "About") LinkedIn makes clicking on them less appealing and less likely (technique #133). The same visual element can persuade

in multiple different ways. For example, placing your profile picture next to the company logo serves initially as a hook to draw your eye in, since we are naturally narcissistically attracted by photos of ourselves (technique #134). Once your attention has been directed to the right side of the screen, you process the two symbols - photo and logo - on a deeper symbolic level, imagining the conceptual blending of the self and the premium membership. Once you get used to seeing who has viewed your profile, the inherent curiosity and self-consciousness makes this a valuable piece of information that draws you repeatedly back to the platform (technique #136).

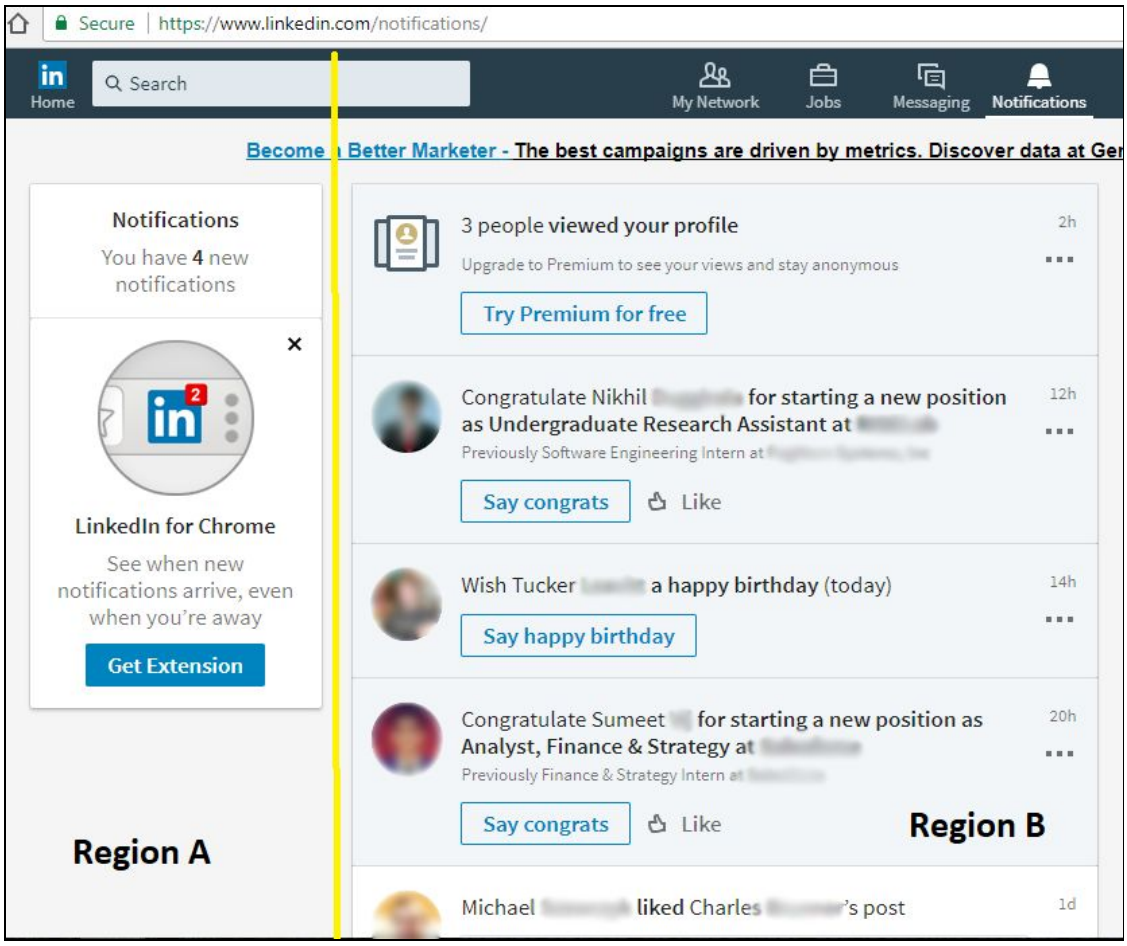

Figure 18. LinkedIn screen 3 appears upon clicking "Notifications" on the top horizontal toolbar, and is comprised of a Chrome extension advertisement and a notifications list.

		A	B	C	D	E
	Screen 3	perform new/unfamiliar behavior	perform existing/familiar behavior	increase behavior intensity, frequency, or duration	decrease behavior intensity, frequency, or duration	stop a behavior
1	one time behavior	150, 151	153, 158, 160	169, 170		
2	one time behavior that leads to ongoing obligation/cost	143, 145, 147	146, 154, 155, 156, 157, 159, 163, 164, 165		149	
3	behavior for a period of time		152	167		
4	behavior on a predictable schedule					
5	behavior is on cue		144, 148	168, 171		
6	behavior is always performed		162	161		

Figure 19. Persuasive design techniques from LinkedIn Screen 3 mapped onto the behavioral grid; color legend on P. 30.

#	Persuasive Vehicle	Method of Persuasion
143	Extension download message	Opens to the door to more & faster notifications even if not onsite, which creates more opportunities to pull you back to the platform if you're online doing something else.
144	Including the number of notifications ("4" written and "2" overlayed on the app icon)	Gets you to look at the ad
145	Glimpse of what the Chrome extension looks like	Helps you imagine what it would look like, which makes it easier to make a decision to get it.
146	"Home" button	Provides an easy exit to more of the feed, which is architected solely by LinkedIn
[refer to Appendix C for full list]		
168	"X liked Y's post" notification	Normalizes "liking" behavior and encourages you to do the same.
169	White color of previous notifications	The softer shade of notifications below the fold makes our eyes more comfortable looking at the notifications, encouraging us to spend more time in this section at the end.
170	White color of previous notifications	Contrast with the blue background color of newer notifications, which makes you want to take action to make these more uniform with the others.
171	Including the timestamp of a notification	Conveys the urgency of reacting by counting the number of hours elapsed.

Figure 20. A sample of the list of persuasive design techniques from LinkedIn screen 3.

In screen 3, the majority of techniques lead to behaviors that create an ongoing obligation due to the nature of messaging; a simple pre-populated one-click reply action by the sender will actually initiate an official chat conversation with the receiver. The receiver is not made aware in the moment that the message was generated by a LinkedIn button, and has no way to distinguish between an organic, independent

message and one generated by the platform. The initial message then opens the door to a chat response, which sparks a longer-term communicative obligation that invests users in the platform. Technique #149 uses the phrasing "even when you're away" to paint the default user state as not being "away," and to reassure users that in the non-ordinary case when they are away from the platform, notifications can bring them back.

		A	B	C	D	E
	LinkedIn	perform new/unfamiliar behavior	perform existing/familiar behavior	increase behavior intensity, frequency, or duration	decrease behavior intensity, frequency, or duration	stop a behavior
1	one time behavior	5, 17, 21, 22, 24, 25, 41, 42, 44, 58, 65, 87, 89, 92, 116, 117, 123, 129, 131, 132, 150, 151	13, 14, 15, 16, 18, 29, 39, 40, 45, 53, 69, 73, 105, 109, 110, 114, 153, 158, 160	34, 52, 57, 111, 169, 170	11, 122	97
2	one time behavior that leads to ongoing obligation/cost	6, 7, 8, 9, 59, 60, 62, 63, 64, 94, 95, 96, 98, 100, 101, 136, 137, 139, 142, 143, 145, 147	3, 19, 20, 30, 31, 36, 55, 76, 93, 103, 104, 112, 113, 127, 128, 140, 146, 154, 155, 156, 157, 159, 163, 164, 165	32, 70, 121, 124	125, 149	
3	behavior for a period of time	99	33, 106, 152	12, 35, 43, 68, 167		
4	behavior on a predictable schedule		66, 67, 72, 77, 78, 79, 83	23, 80	133	
5	behavior is on cue	90, 91, 118, 119, 135	1, 2, 4, 48, 51, 74, 75, 81, 82, 126, 130, 134, 144, 148	168, 171	102	
6	behavior is always performed	28, 71, 85, 88, 138	37, 47, 49, 50, 54, 56, 84, 86, 107, 108, 141, 162	10, 27, 46, 61, 115, 120, 161	26, 38	

Figure 21. The full dataset of LinkedIn's persuasive design techniques mapped onto the adapted version of the Fogg Behavioral Grid. Most techniques prompt one-time behaviors, one-time behaviors that lead to ongoing costs, behaviors on cue. Both new and existing behaviors are prevalent, whereas changing behavioral habits on the product (especially decreasing or stopping behaviors) are de-emphasized.

These mappings of Facebook and LinkedIn's design techniques help elucidate how the seemingly small individual techniques are interconnected to form a web of persuasive guidance. One time behaviors that are higher up vertically in the grid lay the foundation for more habitual behaviors lower in the grid. For example, once users give in to the triggering message of the red notification count, clear their notifications, and experience the satisfaction of removing the jarring red overlay and restoring visual symmetry over the icon, it lays the foundation for them to want to do so every time they login. Moving from columns A to C illustrates the typical progression of user behavior for tasks important to the platform's bottom line such as engagement actions and ad views. For example, comments, likes, and emoji reactions are an important measure of user engagement; the creators of these platforms could plausibly aim for new users to like or comment on their first post, continue to do so after they see how easy and satisfying it can be, and then increase their frequency of commenting and liking. A persuasive design technique like personalizing and micro-targeting ads causes behaviors that are one-time actions, such as ad clicks; over time, as the platform learns which topics to present the user with at the optimal times, they affect the users' mental models of and stance towards the platform as well. Each individual user's behavior grid also influences that of other users due to the interactive nature of social media; the more comments under a post, the more likely that other

users will spend more time with a post or see a comment that then compels them to comment. The finite limit on users' time and attention in the attention economy connects the Fogg behavioral grid for LinkedIn to that of Facebook and every other platform that seeks user attention. Increasing the time spent on one platform (column C) inherently decreases the time available to spend on another platform (column D). One limitation of the grid, however, is that when placing persuasive design techniques into these buckets, due to the difference in what is novel to some users and not others, it creates multiple mappings of the same techniques, which is an issue even Fogg runs into in his original paper, where he includes certain behaviors such as "create a group" and "create a page" under multiple boxes (A2 and B2).

4.3 Design Space Dimensions

Fogg identifies two dimensions of interest when studying persuasive design: the type of behavior change (novelty and increase/decrease) and the schedule on which the behavior change occurs (ranging from one-time to permanent). We identify the following additional dimensions (listed in **Figure 22)**, synthesized primarily from the persuasion theory, psychology, and behavioral economics literature. These ultimately provide a richer multidimensional theoretical space that can account more granularly for more of the psychological concepts of interest in persuasion theory.

The first category of dimensions deals with the temporal process of persuasion. The Fogg behavioral model [64] can be used to categorize design techniques based on which part of the workflow of accomplishing a desired behavior they intervene in: motivation, ability, or prompt. Motivators (pleasure/pain, hope/fear, and acceptance/rejection) create the desire to engage in a behavior or have an attitude; ability factors (time, money, physical and mental effort, social deviance, and routineness) determine the ease of performing the behavior, and prompts (sparks, facilitators, and signals) offer opportune moments to perform them. In this theory, with the right level of motivation, followed by sufficient ability, followed by an opportune prompt, the persuasive design technique will be successful. The second dimension consists of the 8Es of exploratory learning when using a technology: experiment, expectation (what you predict will happen), evidence (gathering), evaluation (comparing experiment with expectation), explanation (analysis of gathered evidence), editing (revision of theory/model), exploration, and exercise [119]. Third, we list the cognitive components of the problem-solving process when a user interacts with technology: forming an interpretation of the technology, defining goals and sub-goals, identifying and using operators, stringing operators into methods, and selecting between methods using selection rules [119]. We can classify persuasive design techniques based on which cognitive component of this process they target. The third dimension incorporates uses of technology products that are not mindless and are more goal-driven (e.g., someone logging into Facebook to send a specific message or search for a particular group rather than engage in the more mindless open-ended scrolling behavior). The second dimension accounts for techniques that evoke both structured and unstructured behaviors on the platform. Since both these types of behaviors are documented in the HCI literature [11], a strength of our list of dimensions is that they account for both.

The fourth category of dimensions deals with the style of delivery of the persuasive techniques. Fogg identifies a functional triad of three different functions computers can play when interacting with people; they can function as tools (increasing the user's capabilities), as mediums (providing an experience), or as social actors (creating a relationship) [19]. When a persuasive design technique is delivered by a computer (or technology product) acting as a tool, it can be recognized by the way it performs tasks like the following:

reduction (simplifying), tunneling (guiding completion through a sequence of tasks), tailoring (using personally customized information as part of the design), suggestion (timed intervention at an opportune moment where both motivation and ability are present), self-monitoring (making behavior visible through tracking), surveillance (making behavior observable by others), and conditioning (reinforcing target behaviors). When persuasive design techniques are enacted through a computer that functions as a medium, we can recognize them as offering exploration and insight via simulations and using procedural rhetoric [19]. Finally, when persuasive design techniques are enacted through computers that function as social actors (an approach studied in the field of computer-mediated communication), we can recognize them via their reliance on physical cues, psychological cues, persuasive language, social dynamics, and adoption of social rules.

Label	Category	Description of Dimension	Possible Values	Number of Possible Values
1	Temporal process	Function of technique in Fogg behavioral model	motivation, ability, prompt [64]	3
2	Temporal process	Exploratory learning process targeted	exploratory learning process targeted: experiment, expectation, evidence, evaluation, explanation, editing, exploration, exercise [119]	8
3	Temporal process	problem-solving substep targeted	interpretation, goals, operators, methods, selection rules [119]	5
4	Delivery style	interpersonal role of the computer	tool, medium, social actor [19]	3
5	Delivery style	personality-mediated dimensions	authoritative/non-authoritative instruction style, cooperative/competitive social feedback, extrinsic/ intrinsic motivation type, positive/negative reinforcement [69]	4
6	User's cognition	Maslow's hierarchy of needs targeted	self-actualization, esteem, belonging and love, safety, physiological [70]	5
7	User's cognition	degree of affective cognition	receiving phenomena, responding to phenomena, valuing, organizing values into different priorities, internalizing values [71]	5
8	User's cognition	type of memory targeted	procedural, semantic, episodic, perceptual, transfer memory [120]	5
9	User's cognition	psychological biases engaged	stems from: not enough meaning, need to act fast, too much information, memory constraints, choice architecture [34]	5

Figure 22. Our list of nine additional dimensions along which to classify persuasive design techniques based on their occurrence during the temporal process of persuasion, their delivery style, and the aspects of the user's cognition engaged. We list the various dimensions along with the possible values along each dimension.

The fifth dimension incorporates findings from Halko and Kientz in the personality-based moderators of successful persuasive design. They identify four parameters that can take on opposite values to be more successful at persuading people with different personality types: persuading using an authoritative vs. non-authoritative instruction style (through an agent), persuading via the notion of cooperative social feedback and teamwork vs. encouraging competition, persuading using extrinsic motivators like trophies or badges vs. intrinsic motivators like feeling fulfilled, and persuading via negative reinforcement (removing aversive stimuli) vs. positive reinforcement (adding positive stimuli) [69]. The sixth dimension refers to Maslow's hierarchy of needs, a widely accepted ordering of the basic psychological, social, and physical needs of humans from the highest-order to the lowest-order: self-actualization, esteem, belonging and love, safety, and physiological needs [70]. We can categorize persuasive design techniques based on which needs they primarily fulfill. Next, the seventh dimension in the user cognition category is taken from Bloom et. al.'s 1956 taxonomy of the affective domain, which has to do with the stages of acquiring, processing, and internalizing values. Along this dimension, persuasive design techniques can fall into the buckets of having to do with receiving phenomena, responding to phenomena, valuing, organizing values into different priorities, or internalizing values [71]. The eighth dimension sorts persuasive design techniques by the type of memory they target: procedural (which stores skills and procedures), semantic (declarative memory that stores facts), episodic (which stores events), perceptual (which stores sensory data), or transfer memory (the interference of old and new memories) [120].

Finally, there exist several psychological limits and behavioral economics biases, and the ninth dimension categorizes persuasive design techniques by which biases they involve. For example, "whether an action is described in terms of risk or gain," can involve the illusion of scarcity, sunk cost fallacy, and relative memories to intentionally persuade [32]. These biases have been grouped under different categories and can arise when there is not enough meaning (e.g., the halo effect where impressions in one area influence those in another), when fast action is required, when there is too much information, when memory processing constraints are triggered [34]. We list some of these specific biases in **Chapter 2**. In this dimension we also include persuasive design techniques that are a direct result of choice architecture manipulation, a design method that asks what optimal users would truly want, and then tries to minimize the costs imposed on those who might "satisfice" (non-optimally plan and conduct behaviors), in order to "nudge" them towards the optimal behavior" [72].

Other dimensions that would be harder to parameterize accurately might be the utility level of the persuasive direction for the user's end goal: is it helping you do something you wanted to do, or is it persuading you to do something you didn't want to do? Using additional data sources that would be accessible to the makers of these sites, such as click-through-rates or deviance in time spent due to controlled introduction of a new feature, we could categorize design techniques based on quantitative measures of their performance in the product. However, identifying these nine dimensions makes progress on the second goal of the taxonomy listed at the beginning of this chapter, of providing a more expansive and nuanced vocabulary with which to describe these techniques. Although these dimensions are not perfectly orthogonal (e.g., the motivation parameter in dimension 1 might include needs from Maslow's hierarchy in dimension 6), we could make a simplifying assumption that each persuasive design technique must take on a value along each of the dimensions.

This allows us to ask questions like: which persuasive design technique serves to provide motivation (dim. 1) as the user is engaged in the exploration part of the learning process (dim. 2) that is building the user's interpretation of the platform (dim. 3) while the computer is acting as a social actor (dim. 4) with an

authoritative feedback style providing extrinsic motivation (dim. 5)? Taking the product of the number of possible values across all dimensions gives rise to a space of 900,000 possible persuasive design techniques. Conceptualizing different techniques in this space gives us the opportunity to discover new techniques and identify the proper context in which to use them based on a deeper understanding of the theoretical underpinnings from which they arise. It also allows us to reverse-engineer a design choice we find persuasive and explain temporally and psychologically how and why it functions. Our approach thus far does not directly address the third goal of the theoretical approach - moralizing persuasion and pumping intuition about how to determine which persuasive techniques are ethical in which contexts. How might we determine what patterns maximize agency and minimize unauthorized persuasion? This is a question we address in the next chapter using a philosophical approach.

5 Implications, Ethics, and the Philosophy of Technology

"If we make explicit how certain technologies shape our lives, we can create the distance we need to be able to relate to these forces. This generates the space to experiment with the use of technology, keeping a sharp eye on the quality of the practices resulting from them, and based on the realization that every practice in which a technology is used shapes our own subjectivity as well." - Peter-Paul Verbeek [40]

Our discussion so far has helped "create the distance" to which philosopher of technology Verbeek refers, such that we can step back and analyze persuasive design techniques and the way they interlock to form digital systems of persuasion. We can now turn to the third goal of the theoretical approach to persuasive design techniques mentioned in **Chapter 4**: gaining a stronger intuition about which techniques are considered ethical in different contexts. Philosophical concepts like intentions and values will play a large role in this discussion, and they will perhaps contribute more to changing the way we frame questions rather than providing direct answers. However, by raising these questions in light of the past two chapters, we can make philosophical progress. In this chapter, we begin by taking a descriptive approach and surveying past work on agency and ethics. We then transition from this philosophical review to a prescriptive argument, in which we analyze a digital product through the lens of philosophy of technology. By doing so, we demonstrate a philosophical approach to measuring the merits of a technology.

First, why does the question of what constitutes ethical persuasion arise? Many contemporary accounts of ethical issues in digital technology focus on information security and privacy. However, there is an underlying and sometimes overshadowed factor driving the design choices that lead to information crises: the fact that information abundance in the attention economy has created the need to profit by maximizing the capture of attention using digital products. Williams notes that "the core challenge of the Internet is that it optimizes more for our *impulses* than our *intentions*," which creates important and sometimes unsalutary threats to attentional freedom [73]. Indeed, as our relationship with technology has evolved, we have slowly allowed persuasive design to chip away at our agential control and more readily allowed platforms to steer us, perhaps similar to the way we will eventually cede control to self-driving steering and took our hands off the wheel. When is the user steering, and when is the platform? To determine the ethicality of persuasion, we must first begin by understanding the nature of human agency and the balance between it and technological structure in guiding behaviors and changing attitudes.

5.1 Structure and Agency

Philosophers and social scientists have long studied the question of whether human behavior is determined by social and contextual structure external to individuals, or whether humans act "freely" of their own agency. The literature on this debate between "structure" and "agency" defines structure as "the recurrent patterned arrangements which influence or limit the choices and opportunities available" and agency as "the capacity of individuals to act independently and to make their own free choices" [74]. On one hand, theorists emphasize the power that agents have in constructing the world. However, persuasion theorist Kelton Rhoads claims that our lack of awareness about persuasive pulls creates an illusion of freedom: "I'd contend we only enjoy the limited amount of freedom that remains between the powerful influences that largely determine our everyday lives. Because we are not consciously aware of those

influences, we perceive freedom" [66]. Products in the attention economy tread the fine line between designing for users' autonomous choices and persuading them to do things they did not autonomously decide or want to do before coming onto the platform. Overall, it is undoubtable that they do persuade us, but we can also acknowledge that our level of control over our behavior plays a role in mediating the effects of these products in our lives.

This is why taking a middle ground stance is reasonable. Modern theorists view structure and agency as coexisting and coevolving: "structure influences human behaviour, and humans are capable of changing the social structures they inhabit" [75–77]. Along this vein, Verbeek proposes a hybrid model of agency, in which we must view the human and the persuasive technology as a hybrid system, because the defining questions of ethical persuasion don't exist without one or the other: "when people develop connections with technology, these connections form the places where freedom must be located [...] freedom is therefore also a hybrid affair, distributed over people and artefacts" [40]. Freedom, then, is not the lack of outside influence, but the capacity to develop a relationship with these influences. Developing this idea, we posit that freedom of attention is the capacity to have full agential control over actions in a digital setting, such that reflection on our relationship with these technologies yields a satisfactory balance between the influence of the self and of the platform. Verbeek further explains that every persuasive technology is grounded in normative claims or assumptions, and that in analyzing persuasive technology, "a design may never be seen purely as instrumental, but always as mediatory" [40]. Whether the normativity is rooted in a financial bottom line from the attention economy or in a non-monetary desire for behavior change, it is important to consider the technology as mediating between the human persuader and the human user.

5.2 Ethical Persuasion

To answer which forms of persuasive design are ethical, it is often easier to begin by identifying forms that are not. There exists a spectrum of different ways to persuade (and, in fact, a spectrum of modes of influence discussed in **Section 2.5**), including education, compliance, coercion, manipulation, seduction, and deception. Correspondingly, there should exist a spectrum of ethicality rather than well-defined brightlines. Consider the difference between persuading a dieting friend to eat a donut either by mentioning donuts to them or by placing a box of donuts by their bedside every morning [12]. Intuitively, the latter seems more unethical than the former, yet they seem to fall along a context-dependent continuum. To take a modern technological example, something seems similarly unethical about data analytics firm Cambridge Analytica's micro-personalized targeted political ads that were ripe with disinformation and were delivered on Facebook using personal information often gathered without consent [78]. The recent indictment of foreign adversaries that used platforms like Facebook and Twitter exactly as intended in order to microtarget, spread propaganda, and incite protests similarly seems wrong [79]. Perhaps one reason why these persuasive operations feel wrong is that people are used as a means to an end (and if one considers the persuaders' ends unethical, then the operation becomes even more unethical). In fact, both Facebook's former president and former head of user growth have admitted "tremendous guilt" over the product's divisive role in the public sphere and its addictive effects on children [80].

In addition to harnessing intuition about the opposite of ethical persuasion, we can look to past work, which has commented on the ethicality of specific types of persuasive design techniques. For example, consider the question of how we should evaluate the strategic design of choice architectures (e.g., persuasive

design techniques like opt-ins and opt-outs in various contexts, rather than forced constraints on choice altogether). According to Thaler and Sunstein, such "nudges" are the most reasonable middle ground between rigid structure and free-form agency, especially since they do not coerce by prohibiting ability to take the alternate course of action. However, the line between persuasion and coercion is not clear; Miller would argue that "much persuasive discourse is indirectly coercive; that is, the persuasive effectiveness of messages often depends heavily on the credibility of threats and promises proffered by the communicator" [26]. Sunstein and Thaler contend that "the sheer complexity of modern life, and the astounding pace of technological and global change, undermine arguments for rigid mandates or for dogmatic laissez-faire" and that "emerging developments should strengthen, at once, the principled commitment to freedom of choice and the case for the gentle nudge" [72].

According to Wood, the persuasion present in advertisements bypasses rationality, undermines agency, and manipulates us without blatantly deceiving us, which is "even more appalling" than pure deception [38]. He says that even advertisements with beneficial ends, such as giving up smoking, do so by unethical means, since they "corrupt the root of rational communication [and] preclude the possibility of any free human community" [38]. Fogg takes a more moderate view on the use of emotion in persuasion, intuiting that it is only unethical "when its intent is to exploit users or when it preys on people's naturally strong reactions to negative emotions or threatening information expressed by others" [19]. He suggests that techniques like operant conditioning or extreme surveillance should raise red flags [19]. Lockton suggests considering intended commercial benefit, intended social benefit, and helpfulness to the user as orthogonal dimensions of what he calls the DwI (Design with Intent) space in HCI [44]. For example, the design choice of making it difficult to put a TV on standby is intended to be environmentally beneficial but could inconvenience individual users. However, he admits that measures along these individual dimensions will remain controversial, which will make it even harder to scope out a space that will be ethically useful [44].

The risk inherent in making general intuitive claims about different categories of design techniques stems from the difficulty in accounting for the intentions of the designer, the intentions of the users, the specifics of the means used, and the actual ends achieved. However, we synthesize a list of useful ethical heuristics and questions that users, designers, and ethical analysts can ask when developing a normative judgement about a persuasive technology. First, we should consider the ultimate outcome of the persuasion as a factor in determining ethicality. Baron has advocated considering the utility, or worthiness of the end, arguing that manipulation is less objectionable if "the worthiness of any end, taken by itself, is sufficient to justify manipulation as a means to it" [38]. Berdichevsky and Neuenschwander further argue for the importance of considering the end, since "if something is unethical for you to do of your own volition, it is equally unethical to do when someone persuades you to do it" [16].

Second, we should consider the means when evaluating normativity of influence, and should specifically consider if a form of influence can undermine rationality and agency by its very existence. Harris poses the following thought experiment: what if YouTube knew the perfect video to play for you at a given time that would be completely entertaining and helpful, and it autoplayed that perfect video? [81]. Something about the act of autoplaying the video and taking over your control of the dimension of time still makes it feel like structure is dominating agency and skewing the balance of the hybrid model of agency. Even if a technique steers a user back to a course that a "reflective and rationally self-governing person would have followed without needing to be steered," what matters most according to Wood is "the way the manipulation undermines and demeans the person manipulated, by violating and disrespecting his rational capacities to choose for himself how to live" [38].

Adar et. al. tie together the consequentialist and deontological considerations by proposing a model of ethical analysis of design based loosely on criminology, which deals with intent as well as means and ends. This three-part model includes "motive (why it happens), means (how to do it), and opportunity (when it works)" [32]. Measuring motive can be tricky, and some have cautioned against it, saying that "choice is such a messy thing to dive deep into, because then you realize that nobody knows what it means to choose" [82]. But taking into account intentionality when determining ethicality seems unavoidable. If we merely rely on a probabilistic, operational picture of which laws of human behavior make certain outcomes more likely, or which means feel the most problematic, we miss out on the underlying motivations and intentions that ground the entire existence of the persuasion in the first place. A useful tool in incorporating judgement of intentions in analysis of persuasion is Fogg's "stakeholder cost-benefit value analysis," which asks the questions of which stakeholders have the most to lose and gain [19]. Mismatches in stakes could indicate mismatches in motivation, which would tilt the ethical equation.

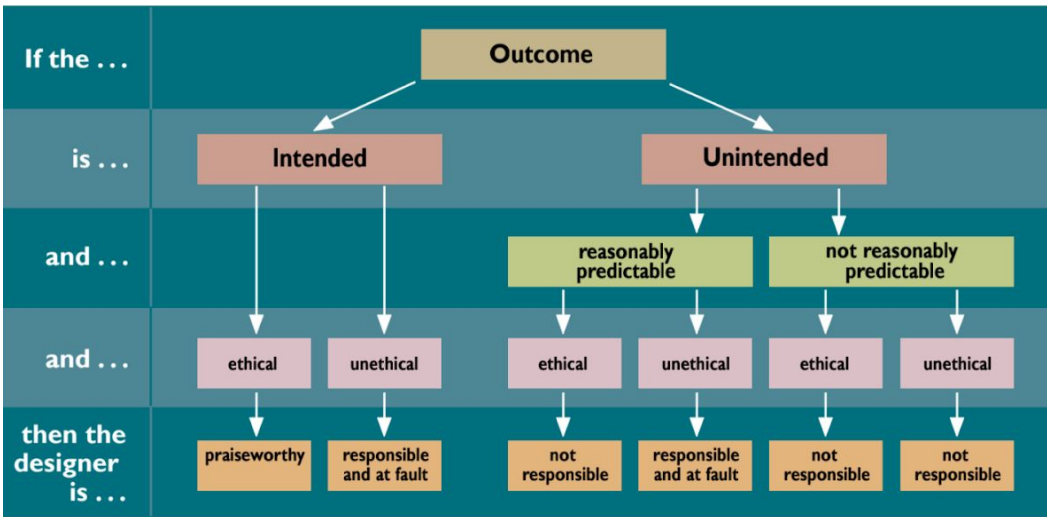

Figure 23: This flowchart, adapted from Berdichevsky and Neuenschwander [16] illustrates one quick method of gaining intuition for designer culpability or praiseworthiness based on the intentionality, ethicality, and predictability of the outcome of persuasion. However, instead of serving to provide a definitive framework, such a chart reveals several of the key unanswered questions relating to ethical persuasion: how should we gauge both user and designer intention? How do we explore the ethicality of the means and ends of the persuasion without collapsing it into a binary decision? And how do we attribute culpability given that persuasion is a symbolic transaction between two agents?

One of the most important and interesting categories of ethical heuristics for persuasive design is centered around the ideas of disclosure and transparency. These ideas appear under many names throughout the literature, but they all boil down to a counterfactual transparency test of sorts: whether or not the persuasive technique still works if the persuasive strategy and vehicle are revealed to the user. If revealing the persuasion technique causes the opposite effect (i.e., if people are so vigilant of not falling prey to the target behavior that they actively perform the opposite behavior), then perhaps there is a case to be made for why it is not ethical. John Rawls called this the publicity principle in 1971: "In its simplest form, the publicity principle bans government from selecting a policy that it would not be able or willing to defend publicly to its own citizens. [...] If it adopts policies that it could not defend in public, it fails to manifest that respect. Instead, it treats citizens as tools for its own manipulation" [72]. This is supported by Adar's HCI definition of "benevolence" in deception, which says that "if the end-user would prefer an experience based on the deceptive interface over the experience based on the "honest" one, we consider the deception benevolent" [32]. This notion of truth is important yet elusive, since designers themselves could be unaware of the truth,

or could be targets of a form of deception themselves (e.g., noisy user data or miscalculated user intentions) [32]. Knowing this, it is important to be aware of the slippery slope of increased social acceptance of persuasion that initially seems unethical [83]. Designers should maintain the essence of their initial ethical intuitions about human agency even once their technology has become ubiquitous.

One challenge is that there is no apparent way to systematically crowdsource user choice about each persuasive design technique in a timely manner without greatly interrupting the user experience. Furthermore, it is difficult to predict whether users will consent to or be able to continue being persuaded by a technique once they have reverse-engineered its mechanisms [16]. Nevertheless, Berdichevsky and Neuenschwander claim as an ethical principle (what they call the "disclosure principle") that creators of persuasive technologies ought to "disclose their motivations, methods, and intended outcomes, except when such disclosure would significantly undermine an otherwise ethical goal" [16]. The Nudget system described in **Chapter 3** provides an extreme example of what such transparency might look like if inbuilt on social networking sites. Efficient disclosure, however, should aim for direct, concise, and standardized communication rather than Nudget's educational goals of generating interest and longer-term knowledge retention.

It is interesting to imagine what user response would be if each major design technique had to display the reasoning behind its design onscreen, or report time spent on site due to that technique. These ideas are beginning to be put into practice. For instance, viewing total time spent on site is currently functionality implemented by HabitLab, a Chrome extension that provides behavioral interventions to reduce distraction by products in the attention economy [62]. Google announced its forthcoming implementation at a platform level of time spent statistics in its latest "Android P" operating system, and so did Apple in the "Digital Health" initiative incorporated in its newest operating system, iOS 12 [84,121]. However, adopting mass transparency will still be difficult since "means-adaptive" persuasive systems are likely weakened by disclosing how they adapt to individuals, whereas end-adaptive systems are more likely to benefit from disclosure [17]. Ploug et. al. explain that when Amazon, for example, makes personalized book sale recommendations, it is functioning as an end-adaptive system that benefits from presenting reasons why it recommended certain products, whereas when it tweaks language to include superlatives like "international bestseller," it is means-adaptive, since the persuasive technique targets the means of persuasion [17]. Drawing back the curtain on the latter kind of technique could reduce the success of the target behavior [85]. So, this kind of disclosure of means or intent would then affect the achieved end of the behavior, which would feed back into our ethical equation of the motive, means, and opportunity.

Another important ethical heuristic is the golden principle based on Rawls's veil of ignorance. In "A Theory of Justice," Rawls argues that we would only agree to obey the ethical rules that benefitted us regardless of who we were in society had we no idea who we were [16]. So, the golden principle applied to persuasive design states that "the creators of a persuasive technology should never seek to persuade anyone of something they themselves would not consent to be persuaded of" [16]. Furthermore, the "accuracy principle" states that designers ought not to lie, especially because humans have a harder time distinguishing computers lying than humans due to the lack of access to behavioral and bodily cues. If human lying is considered unethical, then, by extension, so is computer lying.

Bell and Whaley identify "hiding" and "showing" as two kinds of deception [32], which suggests that we might begin an ethical analysis by asking ourselves what is hidden and what is shown, and with what intention? It is also worthwhile to consider the extent to which the persuasion undermines your autonomy (ability to perform the non-desired behavior)? If the persuasion completely prevents the ability to perform

the non-desired behavior, then perhaps it is more constraining, coercive, and unethical than one where the non-desired behavior can readily be performed. Furthermore, what is the level of reciprocal equality? That is, can the user exert the same forms of influence over the technology or the particular technique in question [19,32]? If yes, then that points to a balance in Verbeek's hybrid agency structure. Lieto and Vernero pose the question of whether the persuasive argument is reducible to a logical fallacy [86]. If yes, then that could point to an undermining of rationality. Fogg also points out the question "if a human were using this strategy to persuade me, would it be ethical?" [19]. This can help focus on the ethicality of the means, since presumably the motivation and outcome would be the same. Finally, we pose the question of how the persuasive technology evolves a user's original intention. Is the change in intention (e.g., from 'send a quick message on LinkedIn' to 'catch up on all my notifications') considered a corruption of intention or an elevation of intention? How does this change in mental framing of intention align with the user's ultimate goals?

Answering these questions and checking different persuasive designs against these heuristics will provide support and reasoning to back our intuition about each persuasive design technique in its own context. These questions and heuristics can also help answer questions of responsibility and blame, as Berdichevsky and Neuenschwander illustrate via the flowchart in **Figure 23**. Verbeek complicates their picture by saying that due to the hybrid nature of agency between humans and persuasive technology, "ambient intelligence and persuasive technologies therefore also interfere in the moral responsibility of people for actions arising in interaction with them" [40]. What does this mean practically in terms of our institutions of liability as well as the design practices that give rise to such situations? Some say that "designers of persuasive technologies should be held responsible only for reasonably predictable outcomes" [16]. But what should be the threshold of reasonability given the platforms' ability to conduct large-scale behavioral experiments that can track the impact of changing every pixel on the screen? And what should the responsibility of these platforms as a whole be in bringing about more ethically palatable forms of persuasion and discarding problematic ones? Before considering the possible design-level and institution-level changes that could help, we consider a few counterarguments.

5.3 Counter-Arguments

Some argue that the effects of persuasion and the call for more ethical persuasive design in the attention economy is overblown. One counter-argument revolves around the use of advertisement blockers; since ads are the linchpin of the attention economy, critics of the ethical awakening in persuasive design argue that using ad blockers significantly reduces possible agency-undermining effects of these platforms. However, as Williams explains, not viewing ads still exposes you to the "design that exploits your non-rational psychological biases in ways that work against you [...] a product or service does not magically redesign itself around your goals just because you block it from reaching its own" [31].

Proponents of the "agency" side in the agency vs. structure debate claim that instead of focusing on the problem of how to make persuasion more ethical, we should focus on empowering people to have more self-control. While self-control is an important part of responsibly using technology that is meant to addict us, we should remember the power of design techniques borrowed from settings like casinos and bottomless buffets [43] that constantly pull us into the digital world of the attention economy. With the massive micro-personalization efforts and artificial intelligence underlying technology, we are subjected to algorithms that have been trained to observe and adapt to our behavior, personality, and most intimate

preferences. For example, we should remember that the placement and wording of many advertisements occurs through sophisticated techniques such as ad format selection via contextual bandits [87]. While some argue that products in the attention economy are just the latest iteration of mass influence technology such as print advertising, radio, or television, it is their ubiquitousness and daily active use pattern ubiquitousness that make their influence different and often more insidious. Portraying the problem as one in which we just need to be more mindful of our interaction with apps can be likened to saying we need to be more mindful of our behavior while interacting with the artificial intelligence algorithms that beat us at chess; equally sophisticated algorithms beat us at the attention game all the time [88]. When we remember that hundreds of engineers and designers predict and plan for our every move on these platforms, it seems more justified to shift the focus of the discussion towards ethical persuasion. Williams maintains that asking users to "adapt" to distraction, become more media literate, and choose options that aren't on a platform's menu when they don't align with the user's values or goals comprises a "pessimistic and unsustainable view of technology" [89]. Understanding that the problems with persuasive technology are not simply of information management and user overuse but fundamentally of our attention is an important realization.

Finally, some say that it is difficult to reject the unethical side effects of persuasive design techniques without allowing for the benefits they introduce into our lives. The impact of these platforms has been far from completely negative. Products like Twitter and Facebook have brought the world together, helped organize resistance against totalitarian governments, uplifted the voices of oppressed populations, and facilitated movements countering police brutality [90]. But at what cost to the individual's cognitive resources and freedom of thought and attention? As long as there remains a mismatch in the bottom line goals of the platforms and the deepest goals in our lives, the way the technologies influence us will be far from optimal and could often be unjustified. By trying to hold our technology products to a higher standard of ethical persuasion, we open the door to more harmonious designs that continuously empower us instead of distracting and frustrating us. But what kinds of changes are needed in the way we make and use technology to move towards this vision? We now review and discuss improvements that could be made at different levels in the attention economy.

5.4 Improvements

My view of helpful steps to move persuasive technology in the attention economy away from problematic modes of influence on users include: 1) more context-awareness of users when considering the behavior patterns for which we design in the attention economy, 2) device-level changes to better inform and empower users, 3) better metrics for designers to design for what truly matters and 4) possible institutional changes for wider-reaching impact.

5.4.1 Awareness of User's Ambient and Cognitive Context

As Mark Weiser famously remarked in "The Computer for the 21st Century," "There is more information available at our fingertips during a walk in the woods than in any computer system, yet people find a walk among trees relaxing and computers frustrating. Machines that fit the human environment instead of forcing humans to enter theirs will make using a computer as refreshing as taking a walk in the woods." Past work has demonstrated the successful design of a context-aware cell phone application SenSay that uses multiple sensors to determine the optimal time and way to send notifications [91]. Tsikerdekis et. al. have shown the successful use of nonverbal detection techniques to prevent context-specific,

identity-based persuasion to safeguard online communities from deception [92]. Hiniker et. al. have described how designers can detect when a product is being used intentionally or mindlessly, and can change its interface to adapt to the type of use by foregrounding the primary use case and backgrounding others, "unbundling" single common patterns of use, and designing for disengagement once the primary purpose is achieved (e.g., when Gmail displays a congratulatory message over an empty inbox) [11]. When products become more intentional about realizing the user's purpose and style of engagement, then they can design for minimization of overreaching manipulation and elimination of all distraction besides what is necessary to fulfill the primary intention at hand.

5.4.2 Device-level Changes

Positive change in the movement towards ethical persuasion should also come at the device level. The very design of home screens and default notification structures drives many persuasive techniques that can be disrespectful of our attention by placing us in a "catch-up" mindset [13]. We also need better standards for labeling: in a critical review of captology, Atkinson advocates labelling persuasive elements such as advertisements with consistent labels across apps, which can bolster users' rights to accept or reject the content of persuasion [20]. Harris suggests that the Apple App Store should reward apps that lead to a greater proportion of "time well spent" (determined by users through questionnaires) instead of solely rewarding the ones that are trending or have high download counts [12]. In other words, the menu for what users can choose from should be sorted by what's good for you, like placing "healthier" options at arm's reach on a shelf. Another design change he advocates is the practice of having your phone or computer routinely giving you detailed breakdowns about your time spent on various sites, asking you whether this is how you want to be spending your time, and guiding you through changing your habits if it is not [12]. Overall, architects of the platforms that host apps in the attention economy should design more intentionally to prevent the possible unsavory side effects that their persuasive design causes, and they can equip users with higher levels of information about their use of the platforms, healthier defaults, and ways to continue improving their digital habits. Changes that come from the makers of smartphones and operating systems will have a wide reach in self-regulating the attention economy and holding it to more consistent ethical standards across the board.

5.4.3 Better Metrics

Metrics are important because they play a key role in the feedback loops that drive design; we improve and iterate our designs based on what we can measure. They also play a role in user's mental models of their purpose as they enter and spend time on these platforms; often, we adopt as our personal metrics those the site deems important (e.g., Facebook likes, Tinder swipes, and Snapchat streaks). As designers, we parameterize a user's needs and values through measured behavior using such metrics, but we lose information about their deeper driving desires. Users do not come to these products with the intention of solely spending time or money or sending a high number of messages or comments, yet that is often what is measured and optimized for. This demonstrates the importance of creating more meaningful metrics that parameterize what users truly care about, as well as encapsulating the root of the needs and goals that bring them to these platforms. How might we move towards a more phenomenological understanding of questions like "Which YouTube videos helped you learn an instrument? What comedy made you laugh? What encouraging videos helped you face your fears?" [93]. We address a possible approach to these questions in the next two subsections. It might initially seem that answering such questions would not create metrics that

link usefully to measuring revenue for the sites; however, metrics that align better with user values are not always contrary to the long-term business profits of companies in the attention economy; they actually pose a market opportunity. Holmen explains that "people increasingly want to spend time well, not spend more of it [...] If it's our shallowest self which is reflected to us every time we open Facebook, Instagram and YouTube, the best business opportunity around might be to begin to cater for our aspirational selves" [94]. Metrics that help companies encourage the formation of healthy digital habits could actually inspire more loyalty and customer satisfaction. Analogous historical evidence for this can be seen in the sales successes of car companies that offered seat belts almost two decades before they were federally mandated in response to consumers' desire for a safer user experience [122]. Due to the nature of the attention economy, such a change made by one product could spur ripples of change across other products. Finally, better metrics enable more introspective designers who can more easily determine the ethicality of particular forms of persuasion with clearer answers about their users' deeper values and the experiential and cognitive effects of design choices.

5.4.4 Institutional Changes

Legal regulation of persuasion is rarely an optimal route for the fast-moving technology industry. Nonetheless, in light of the recent national impacts of technologies like Facebook, Twitter, and YouTube in our fundamental democratic institutions (e.g., foreign interference in the 2016 presidential election [78,95,96]), people have begun to call for regulation at an institutional level to reduce the social fallout of these technologies by moderating the techniques they employ. Proponents of such change compare regulation of the attention economy to regulation of the addictive tobacco industry, regulation of extraction economies like coal, or the rise of the user-demand-driven organic food movement [12,97]. In this way, design ethicists call for the attention economy to turn to green solutions that are regenerative and replenishing, since attention has become the modern mined resource. Some have even called for the need for a "Silent Spring" awakening in the attention economy that leads to a "Digital Environmental Protection Agency" [98], transitioning control from the private sector's market regulation to the public sector's policies. Techno-sociologist Zeynep Tufekci explains why the time is ripe for public discourse about political action regarding these platforms, drawing on an analogy with the automobile industry: "Facebook is only 13 years old, Twitter 11, and even Google is but 19. At this moment in the evolution of the auto industry, there were still no seat belts, airbags, emission controls, or mandatory crumple zones. The rules and incentive structures underlying how attention and surveillance work on the internet need to change. But [...] there are few solutions to the problems of digital discourse that don't involve huge trade-offs—and [...] these are deeply political decisions" [99].

Undoubtedly, more innovative business models will play a major role in restructuring the motivations that compel companies to exploit psychological weaknesses in the name of profit. Roger McNamee, an early investor in Facebook and a personal mentor to its creator Mark Zuckerberg, recently laid out the argument for a subscription-based ad-free model of Facebook, claiming that the lack of ads and sponsored clickbait could improve the user experience enough to charge a subset of users a small fee [100]. However, switching to a monthly payment or pay-per-use business model would pose new challenges (such as people's valuation of the product being highly variable due to the endowment effect, where users in one survey said they would pay an average of about $7 per month to use Facebook but would need $75 per month to stop using it [101]). There have also been creative ideas such as monetizing attention in escrowed bonds at the outset of persuasion, for a more decentralized, market-based solution regulated by consumers

themselves based on their judgement of the outcome of the persuasion [102]. However, as we have discussed earlier, removing ads in lieu of other revenue streams will not solve all the issues that arise with psychologically manipulative design. McNamee has also called for greater user control over the social graph, including the power to export it more easily to other social networks, which could give smaller platforms access to network effects and make them more competitive with the monopolistic companies that currently dominate the attention economy [100]. Regulating the large monopolies that set the standard for design in the attention economy could level the playing field and reset design standards for all apps.

These institutional changes will require time and consideration to become enactable. However, there are more immediate steps to be taken in the shorter term. One such step is the wider adoption of the process of value-driven design, an approach based on the principle of foregrounding deeper user values in design. We now briefly discuss this approach along with some of its challenges.

5.5 Value-Driven Design

> *"On every platform, a person who wants to be attentive to their friends can find themselves in a state of frazzled distraction. As users, we end up acting and socializing in ways we don't believe in, and later regret. We act against our values: by procrastinating from work, by avoiding our feelings, by pandering to other people's opinions, by participating in a hateful mob reacting to the news, and so on [...] The coded structure of push notifications makes it harder to prioritize a value of personal focus; the coded structure of likes makes it harder to prioritize not relying on others' opinions; and similar structures interfere with other values, like being honest or kind to people, being thoughtful." -Joe Edelman* [103]

How might we keep persuasion that leads to a target behavior unaligned with the users' deeper values as low as reasonably achievable? One approach is value-driven design, which maintains that we must gauge and then design to support each user's true values, not just goals, and that we must distinguish true values from ideological commitments (broader socialized values), norms (standards chosen due to social consequences), and goals or fears [103]. Even direct values themselves include allegiances, situational values, and performative values - a range that stems partly from the fact that people have tiers of preferences about what they want, and what they want to want, which social psychologists have encapsulated in theories of multiples selves [104,105]. In the longer term, this value-centered design might involve rethinking traditional separations between the user and the designer in the design process, and wider adoption of design methods like Value-Sensitive-Design (VSD) and Participatory Design (PD) [106,107,125].

But an added challenge comes from the fact that persuasive technology influences human intentions and values itself. Not only does it condition human behavior, but it also helps "shape the interpretations on the basis of which human beings make intentional decisions" [40]. In an extreme case of such influence, our capacity for moral reflection can be "swapped for a voluntary exposure to influence from technology" [40]. So then how do we truly gauge values and design for the values that will actually lead to organically defined well-being? What mental states or behavior patterns denote true human flourishing, and how might we strive towards those? These are important questions for designers of persuasive technologies of the future, who are creating interfaces for users tired of being pulled and cajoled to take micro-actions that aren't aligned with their deepest intentions. Getting people to understand their own values and goals will be a challenge, given the disparity between the experienced and remembered self [108]. But there are things that we know

we cherish and want in life when we are in touch with our highest, deepest, most reflective selves, and finding the vocabulary to describe and create those experiences will prove beneficial.

From an analytical perspective, we can move towards an experiential understanding of what we lose and gain when we interact with a specific technology by drawing from foundational work on phenomenology in the philosophy of technology. We now analyze a specific technological product using a phenomenological lens, and we demonstrate the emergence of subtle and important insights about the experiential impacts of technology.

5.6 Case Study: Phenomenological Analysis of Google Maps

Google Maps is one of the most prolific navigation technologies in the world, and over the thirteen years since its launch, it has come to play an integral role in the daily lives of users worldwide [109]. Recent studies have shown that using Google Maps is linked to the weakening of particular cognitive functions related to navigation. Such findings shed light on the broader personal impact of such a seemingly innocuous, benevolent technology. This aspect of Google Maps motivates questions about the user experience that can be approached in unique ways using foundational concepts and opinions from the philosophy of technology. Although Google Maps is not directly a competitor to the mainstream products in the attention economy, its cognitive underpinnings make it insightful for this kind of analysis. Phenomenology provides a mode of inquiry into " the conditions of what makes things appear as such," and lets us probe the co-constitutive nature of humans and technology and the complete experience, in all its sensory and cognitive fullness, of interacting with a technology [110]. We now deconstruct the technology through the lens of two philosophers of technology and their phenomenological perspective towards technology's impact on our minds and lives. We draw from Heidegger's 1977 essay "The Question Concerning Technology" and Albert Borgmann's 1984 essay "Focal Things and Practices." By analyzing this technology from the lens of philosophy, we demonstrate a method of analysis that can be applied to products in the attention economy, in order to encapsulate some of the more abstract and experiential positives and negatives with which they are associated.

One of Google Maps's core features is providing voice and text-based navigational guidance. Using location data, satellite imaging, crowdsourced reviews, and live traffic data, it enables users to get accurate, real-time directions for the best route to a selected destination. There exist many similar GPS-based routing services like Apple Maps, Waze, and Hopstop, and several of the critiques we discuss here could apply to them. However, we choose to focus on Google Maps since it is one of the oldest and most far-reaching solutions, hosting maps of over 28 million miles since 2005, of locations ranging from North Korea to the Great Barrier Reef [109,111]. It also outranks its competitors in numbers of users by millions, and now has over a billion monthly users [112]. We focus our analysis on the mobile use case, an illustration of which is captured in **Figure 24**.

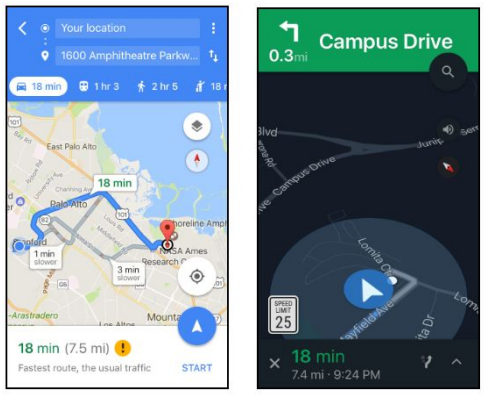

Figure 24. Example screens from the Google Maps mobile application.

Now we discuss the recent work in neuroscience and psychology that sheds light on how navigational aids like Google Maps adversely impact cognition. More specifically, when using GPS technology to navigate, we "put less work into generating our own internal picture of the world" [113]. These internal representations are called "mental maps" and consist of relative positional relations between worldly landmarks. According to one researcher, "They are very individual...The things which matter to you might be completely different to those that matter to your wife or your children" [113]. These mental maps allow people spontaneity and flexibility while traveling; without them, "you can never deviate from the route you know, look for shortcuts, or improvise if the situation calls for it" [113]. Furthermore, additional findings suggest that people who use a set of directions rather than traditional paper maps draw their route and surroundings less accurately after the trip, and are more susceptible to unknowingly taking redundant or retraced paths [113]. Neuroscientists have even demonstrated that English taxi drivers, a seasoned group of manual navigators, have more gray matter in the posterior hippocampus of their brains [123]. Although not an inherent cognitive benefit in itself, this represents a lack of risk of psychiatric disorders associated with lower levels of gray matter.

These cognitive findings draw attention to the link between Google Maps and changing subjective experience. As philosophers of technology, we can then ask the question of what reality is like for someone using Google Maps. *What qualities characterize the experiences of a Google Maps user, and what representations of the world are available to them?* Due to the phenomenological nature of the questions, it seems natural to call upon Heidegger to provide a formal approach. His concept of "enframing" represents one of the main tenets of his philosophy of technology: that the essence of technology has to do with its way of revealing our notions about our relationship to technology itself [114]. For example, Heidegger explains how upon construction of a dam, our conception of a river changes from that of a majestic flow of water to that of a standing-reserve of energy waiting to be tapped [114]. By disentangling our own attitudes, we can use our understanding of our enframements of technology as windows into understanding the phenomenological role of the technology. Since society has conceived and developed Google Maps to supplant navigational cognitive pathways and neural mapping capabilities, we have signed away the capacity to experience the fullness of self-constructed direction sense. This loss is not merely cursory but indicates the loss of a deeper part of our reality. As physicist John Huth puts it, "It's almost like depriving yourself of music, or a conversation with another person. There's a richness that you're missing out on" [113]. In this way, understanding what aspects of our brain we "conceal" via our collective enframing of Google Maps reveals what we miss out on due to repeated usage of the tech.

However, it is undisputed that things are also gained phenomenologically. Google Maps enables us to feel the comfort of setting out fearlessly to new destinations with the guarantee of accurate rerouting in case of misstep. But is this enough to extinguish the concern of missing a deeper richness? As Google executive Dan Sieberg puts it, the technology allows us to "arrive anywhere – Edinburgh, Cologne, Tokyo – and within moments know our way to our hotel, have a list of the best-rated restaurants and know the best route to take on the metro" [111]. It arguably reduces the friction of having to make manual searches, phone calls, and physical map consultations to iron out these details ourselves. It is not hard to imagine that being forced to do this for every potential destination would be an unbearable burden on our time. But are there times when undertaking such actions would not only contribute to our own cognitive good but also return a more encompassing splendor to our quality of experience? We believe Heidegger would say yes. Heidegger might point out that regardless of the alleviation of certain mental efforts, using the technology makes it so that reality becomes more transactional. The user interface constantly visually highlights the path to take,

and a voice regularly blurts out the next mental step a user should make, or the next street name or turning point to selectively seek. This transforms a smooth, continuous journey into a choppy, discrete set of miniature destinations, one for each instruction in the route. Users' attention is focused so intently on the directions from point A to B that it becomes hard to meaningfully observe let alone remember the current surroundings.

Another aspect of the human travel experience that is concealed when using Google Maps is the spontaneity and serendipity of whimsicalness - taking the long way just because you feel like it, or even thinking to turn down an unexplored road. In addition, there is a risk that the foresight afforded by Google Maps "could take away the magic of seeing [a place] for real for the first time," [111]. The integration of third-party reviews of points of interest along the way could even make it so that "everyone will head for the same well-reviewed destinations" [111], creating a sameness that would detract from the thrill of stumbling upon exotic-little known locations that become backdrops for unique memories only you have. Overall, the same affordance of Google Maps - the ability to always have a perfect, virtual travel guide in your pocket - both adds and takes away from aspects of what is required to make travel completely and subjectively fulfilling phenomenologically.

Apart from the question of what is lost and gained when using the app is the question of the nature of the embodied sense of presence felt by users of Google Maps. Here it will be useful to draw upon Borgmann's characterization of "focal practices" and their associated sense of wholesome presence, which we will now recount. According to Borgmann, focal practices are central to living the good life. They consist of "a few simple things," and are "liberating and invigorating…deeply touching and fleeting…where effort and joy are one; split between means and ends" with a "unity of achievement and enjoyment, of competence and consummation" [111,115]. He concretizes this by using the example of an outdoor runner: "running is simply to move through time and space, step-by-step. But there is splendor in that simplicity," he says, describing the wholeness of the experience of delighting in both the process of undertaking an activity as well as completing it.

Applied to Google Maps, this concept spurs the following question. If the purpose of travel is to undertake a journey qua journey, what are the implications of orienting the user in a such a way as to hide key aspects of how the journey came to be? Indeed, the technology is designed such that the backend work required to chart a path, observe traffic, adjust the path correspondingly, and recommend suggested stops or leisure activities is abstracted to the point where the user is never exposed to the nuanced decision-making behind the scenes. This disables the full, end-to-end understanding that characterizes a focal practice. Heidegger's idea that everything you reveal conceals something else is also apt here. Why reduce the possible enjoyment and enrichment a user would experience in understanding the inner workings of the *means* of travel by foregrounding the *ends*?

Another relevant aspect of the technology under a focal practice lens is the "egocentric" reorientation of the map to put users in the center. Studies have shown that following egocentric cues leads to lower recall of landmarks passed on the route as compared to traditional maps [113]. In fact, in one recent experiment, researchers studied the varied effectiveness in acquiring "landmark-referenced (allocentric) knowledge" and "view-referenced (egocentric) knowledge" [116]. They found that learning navigation using cartographic maps led to faster improvements in judgements of relative direction, while learning using routes led to more accurate results in scene and orientation-dependent pointing. This sheds light on the fact that using a route-based guide like Google Maps foregrounds aspects of the individual's perceived role in navigation, and gives users a quantifiably stronger sense of how the world points out away from them. This contrasts with

the allocentric presentation inherent in paper maps, which "forces you to plan and frame your route within a meaningful context: towns, forts, universities, parks, and natural features named for local heroes and history or local flora and fauna" [117]. Although Google Maps does display these features when used as just a map, it strips away all extraneous details as soon as users enter a destination [117]. In this way, using the technology lessens the potential to truly and focally make the practice of travel one's own.

One could argue that Google Maps's recent integration of more three-dimensional street views and intelligent recommendations about midway stops represents attempts to make the technology more allocentric. Nevertheless, just like Borgmann's concept of the outdoor runner, there is something more organically satisfying gained from disengaging with the technology to undertake the crux of the practice of travelling. One journalist who decided to navigate around Boston without Google Maps poignantly recounted such an experience: "That evening, as I drove down the Pike, my window down and my phone buried deep in my pants pocket, the city snapped into shape around me. Suddenly I was not just a guy who had learned a set of moves. I was a guy who knew his way" [113].

Viewing Google Maps through the lens of Heidegger and Borgmann allows us to articulate more precisely what aspects of the human experience are at stake when we use the technology. Thinking about these issues as agents for technological change brings up important design questions to consider as we shape the technology in coming years. For example, how might we design Google Maps to support mental mapping instead of supplant it? Might we include quizzes or route-based games to keep cognitive skills sharp? Could we create different modes of usage corresponding to people's different intentions behind travel, allowing more transactional interactions for urgent, target-oriented journeys versus more serendipitous, free-flowing travel experiences for leisure trips? This could take a form ranging from a simple question requiring user input about how leisurely their travel is before giving them a route, to a redesign of the route-generation algorithms to increase room for seemingly suboptimal suggestions to parameterize serendipity. Furthermore, how might we reduce the egocentric worldview the current technology induces? One writer notes that "You can redirect most smartphone navigation apps to align with the magnetic compass instead of your direction of travel" [117]. Might we make this an easier, more quickly accessible choice on the screen? Might we remove the screen altogether and have the guiding voice be more mysterious and whimsical?

Even more ideas arise when responding to the insights gained from viewing Google Maps through the focal practice lens. How might we preserve the focal nature of travelling without ridding ourselves of the ease of the tech? Borgmann takes a step beyond Heidegger to say that "we do not have to seek out pretechnological enclaves to encounter focal things," so might it be possible to use the global reach of Google Maps to not just support the idea of travel as a focal practice but fully transform travel into a focal practice? Considering Borgmann's ultimate view of focal practices as providing a "profounder commerce with reality" that "deepens charity and compassion," could we include prompts or incentives for travelers to engage constructively with others along their journeys rather than being lone arrows on the screen? Could we achieve this by not providing statistical insights gained from others' travels to allow for more personalization? What if we made the maps annotatable to allow users to scrawl personalized nicknames for landmarks? As a front-end engineer or user experience designer working on Google Maps, or even a designer using the Google Maps API for my own app, such questions would be very relevant to my work. Framing them using the rhetoric of philosophical analysis provides a unique window into such design choices.

In the end, looking at Google Maps through the lens of Heidegger allows us to critically examine the nature of the reality available to users, showing both the experiential gains and losses. Examining Google Maps through Borgmann's concept of focal practices helps elucidate how abstracting information and portraying egocentric symbols contributes to and detracts from wholesomeness of travel as a focal experience. As philosophers and technologists, such analysis allows us to propose new directions for even the most established technologies. In a world where innovation often outpaces critical discourse, bringing the voice of philosophers into modern discussions is especially worthwhile as we strive to gain a more complete understanding of the positives and negatives of technologies on the human experience.

6 Future Work and Conclusion

In the past three chapters, we have taken a system design approach, a persuasive theory approach, and an ethics and philosophy of technology approach to studying persuasive design techniques. We began by creating and evaluating a system that makes persuasive design techniques visible on Facebook feeds, finding that our system could improve transfer knowledge of such techniques significantly more than traditional methods of educating users. We then mapped specific persuasive design techniques from Facebook and LinkedIn onto the behavior grid and theoretically analyzed how they interconnect to function as influential systems, offering a standard behavioral vocabulary with which to classify them. Finally, we synthesized several normative heuristics to provide techniques for designers to intuit what constitutes ethical persuasion, providing a sample phenomenological analysis of Google Maps to advocate for experiential and value-driven perspectives in the development of future products in the attention economy.

There is important future work to be done in all three of these domains. On the empirical side, we measured the effects of a system like Nudget on user knowledge, but how might such a system affect user behaviors, engagement styles, or time spent on the platform? We also hope to more granularly measure the impact that design techniques in different categories of persuasive theory have on user engagement, attitudes, and cognitive resources on social media. Furthermore, our participant pool was highly trained, which leaves scope for future work to answer the question of how a more representative sample would respond to a system like Nudget, and how results would change over the course of a longer longitudinal study. Using the insights from the user feedback on systems like Nudget can help create better design tools for the architects of mass consumer social media sites, by giving them more effective ways to understand how their persuasive design techniques affect people's cognition.

From a persuasion theory standpoint, we can continue to probe deeper into the Fogg Behavioral Grid by more closely examining techniques in individual grid boxes of interest. Using the additional dimensions we identify in Chapter 4, it would prove incredibly fruitful to use a methodology like morphological analysis to expand the theoretical space that persuasive design techniques occupy. This would help us discover new methods of persuasion as well as boundary conditions and extreme cases of persuasion, which would further develop our ethical intuition about persuasion [118]. We could use such a taxonomy to discover new persuasive techniques, and to better understand how to benefit from the same cognitive weaknesses and abnormalities that sometimes lead us to act against our goals by turning them into creative ways to benefit from our flaws and perform behaviors we deeply want.

We are at a critical point in our relationship with popular technology products in the attention economy. In the coming years, understanding the ways in which they persuade us will be a crucial problem to get right, because the stakes and influence of these technology is higher than ever before. These products shape our thoughts and actions both when we are on and off their platforms. Via personalization and filter bubbles, they change the spread of information and subsequent social and cognitive stratification of society. Maintaining a demand for ethical persuasion and freedom of attention demonstrates the need for more constructive design, institution-level conversations, and shifts towards ethical persuasion. We should ultimately continue to find, test, and implement ways for designers of these platforms to design for our deeper values and use technology to nourish rather than distract us, to encourage us to fulfill our goals rather than scroll mindlessly.

There is a slow but sure social awakening occurring in the technology industry and in the attention economy concerning the ethical responsibility of tech platforms to examine how they persuade, and there needs to be further work at the design level, platform level, and institutional level to make sure we maintain the right level of control over our cognitive resources. Understanding the way technology guides and shapes our behaviors and attitudes is crucial to understanding how we function in the world today. By empowering both everyday users, the creators of these technologies, persuasion scholars, and policymakers with better ways to understand, discuss, and mindfully design the persuasive effects of our platforms, we can make progress towards a world in which people feel more empowered with an understanding and control of their own attention, autonomy, and digital lives.

References

1. Simon HA. Designing Organizations for an Information-Rich World [Internet]. 1 Sep 1969 [cited 28 May 2018]. Available: http://zeus.zeit.de/2007/39/simon.pdf
2. Kelly K. "Better than Free" [Internet]. 5 Feb 2008 [cited 28 May 2018]. Available: https://www.edge.org/3rd_culture/kelly08/kelly08_index.html
3. Wu F, Huberman BA. "Novelty and collective attention." Proc Natl Acad Sci U S A. 2007;104: 17599–17601.
4. Fogg BJ. "Mass Interpersonal Persuasion: An Early View of a New Phenomenon." In: Persuasive Technology Lab [Internet]. [cited 29 May 2018]. Available: http://captology.stanford.edu/wp-content/uploads/2014/03/MIP_Fogg_Stanford.pdf
5. Social Media Fact Sheet. In: Pew Research Center: Internet, Science & Tech [Internet]. 5 Feb 2018 [cited 29 May 2018]. Available: http://www.pewinternet.org/fact-sheet/social-media/
6. Smith A, Anderson M. "Social Media Use in 2018." In: Pew Research Center: Internet, Science & Tech [Internet]. 1 Mar 2018 [cited 29 May 2018]. Available: http://www.pewinternet.org/2018/03/01/social-media-use-in-2018/
7. Shi E, Niu Y, Jakobsson M, Chow R. "Implicit Authentication through Learning User Behavior. Lecture Notes in Computer Science." 2011. pp. 99–113.
8. Matthieu. "The 50 numbers to know about social networks in 2018." In: Spitche [Internet]. 8 Jan 2018 [cited 29 May 2018]. Available: https://www.spitche.com/digital-marketing-en/50-numbers-know-social-media-2018/
9. Hackett C, McClendon D, Posts, Bio. "Christians remain world's largest religious group, but they are declining in Europe." In: Pew Research Center [Internet]. 5 Apr 2017 [cited 29 May 2018]. Available: http://www.pewresearch.org/fact-tank/2017/04/05/christians-remain-worlds-largest-religious-group-but-they-are-declining-in-europe/
10. Feldman B. "Most of the Time, You're Watching What YouTube Wants You to See." In: Select All [Internet]. 11 Jan 2018 [cited 29 May 2018]. Available: http://nymag.com/selectall/2018/01/most-of-your-time-spent-on-youtube-is-dictated-by-youtube.html
11. Lukoff K, Yu C, Kientz J, Hiniker A. "What Makes Smartphone Use Meaningful or Meaningless?" Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies. 2018;2: 1–26.
12. Harris, S. "Waking Up Podcast #71 — What is Technology Doing to Us? | Sam Harris." In: Sam Harris [Internet]. [cited 25 May 2018]. Available: https://wakingup.libsyn.com/71-what-is-technology-doing-to-us
13. Edelman J. "Is Anything Worth Maximizing? On Metrics, Society, and Selves" [Internet]. Available: http://nxhx.org/maximizing/
14. Roettgers J. "For Facebook, Change Is Political. For Google, It's Personal." In: Variety [Internet]. 8 May 2018 [cited 25 May 2018]. Available: http://variety.com/2018/digital/news/google-personal-wellbeing-online-abuse-1202803655/
15. Constine J. "Instagram CEO confirms upcoming "time spent" Usage Insights." In: TechCrunch [Internet]. 15 May 2018 [cited 25 May 2018]. Available: http://social.techcrunch.com/2018/05/15/instagram-usage-insights/

16. Berdichevsky D, Neuenschwander E. “Toward an ethics of persuasive technology.” Commun ACM. 1999;42: 51–58.
17. Ploug T, Hasle P, Oinas-Kukkonen H. “Persuasive Technology: 5th International Conference,” PERSUASIVE 2010, Copenhagen, Denmark, June 7-10, 2010, Proceedings. Springer Science & Business Media; 2010.
18. Kaptein M, Eckles D. “Selecting Effective Means to Any End: Futures and Ethics of Persuasion Profiling.” Lecture Notes in Computer Science. 2010. pp. 82–93.
19. Fogg BJ. “Persuasive Technology: Using Computers to Change what We Think and Do.” Morgan Kaufmann; 2003.
20. IJsselsteijn W, de Kort Y, Midden C, Eggen B, van den Hoven E. “Persuasive Technology: First International Conference on Persuasive Technology for Human Well-Being,” PERSUASIVE 2006, Eindhoven, The Netherlands, May 18-19, 2006, Proceedings. Springer Science & Business Media; 2006.
21. Davenport TH, Beck JC. “The Attention Economy: Understanding the New Currency of Business.” Harvard Business Press; 2002.
22. Stothart C, Mitchum A, Yehnert C. “The attentional cost of receiving a cell phone notification.” J Exp Psychol Hum Percept Perform. 2015;41: 893–897.
23. Mark G, Gudith D, Klocke U. “The cost of interrupted work. Proceeding of the twenty-sixth annual CHI conference on Human factors in computing systems.” CHI ’08. 2008. doi:10.1145/1357054.1357072
24. Stone L. “Just Breathe: Building the case for Email Apnea.” In: HuffPost [Internet]. HuffPost; 8 Feb 2008 [cited 28 May 2018]. Available: https://www.huffingtonpost.com/linda-stone/just-breathe-building-the_b_85651.html
25. Daniel C. McFarlane & Kara A. Latorella. “The Scope and Importance of Human Interruption in Human-Computer Interaction Design.” Human-Computer Interaction. 2009; Available: https://www.tandfonline.com/doi/abs/10.1207/S15327051HCI1701_1
26. Dillard JP, Pfau M, Chaiken S. “The Persuasion Handbook: Developments in Theory and Practice.” SAGE; 2002.
27. Shrager J. “Demandance” [Internet]. 2015. Available: http://arxiv.org/abs/1507.01882
28. Mole C. “Attention.” 2009; Available: https://plato.stanford.edu/entries/attention/
29. Hallden L. “Towards a Bra-free Instagram Experience.” In: NewCo Shift [Internet]. NewCo Shift; 30 Dec 2017 [cited 29 May 2018]. Available: https://shift.newco.co/towards-a-bra-free-instagram-experience-3e43273b611f
30. Williams J. “The philosophy and ethics of attention and persuasion” [Internet]. 28 Jun 2016 [cited 25 May 2018]. Available: https://onlinedialogue.nl/wp-content/uploads/2016/06/James-Williams-DPI-talk-in-Utrecht-June-2016.compressed.pdf
31. Williams J. “Why It’s OK to Block Ads. In: Practical Ethics” [Internet]. 16 Oct 2015 [cited 25 May 2018]. Available: http://blog.practicalethics.ox.ac.uk/2015/10/why-its-ok-to-block-ads/
32. Adar E, Tan DS, Teevan J. “Benevolent deception in human computer interaction.” Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI ’13. 2013. doi:10.1145/2470654.2466246
33. Cialdini R. Pre-Suasion: A Revolutionary Way to Influence and Persuade. Random House; 2016.
34. Benson B. “Cognitive bias cheat sheet – Better Humans.” In: Better Humans [Internet]. Better Humans; 1 Sep 2016 [cited 25 May 2018]. Available:

https://betterhumans.coach.me/cognitive-bias-cheat-sheet-55a472476b18
35. Petty RE, Cacioppo JT. "The Elaboration Likelihood Model of Persuasion." Communication and Persuasion. Springer, New York, NY; 1986. pp. 1–24.
36. Eslami M, Karahalios K, Sandvig C, Vaccaro K, Rickman A, Hamilton K, et al. "First I "like" it, then I hide it." Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems - CHI '16. 2016. doi:10.1145/2858036.2858494
37. Rhoads K. "Definitions: Persuasion, Compliance, Propaganda, Brain Washing, & Social Psychology" [Internet]. [cited 25 May 2018]. Available: http://www.workingpsychology.com/definit.html
38. Coons C, Weber M, Baron M, Barnhill A. "Manipulation: Theory and Practice." Oxford University Press; 2014.
39. Rhoads K, Cialdini R. Working Psychology: Applying Psychology to Influence, Persuasion & Teamwork [Internet]. [cited 25 May 2018]. Available: http://www.workingpsychology.com/
40. Verbeek PP. "Ambient Intelligence and Persuasive Technology: The Blurring Boundaries Between Human and Technology." Nanoethics. 2009;3: 231–242.
41. Coons C, Weber M. "Paternalism: Theory and Practice." Cambridge University Press; 2013.
42. Williams, J. "NEXT17 | Distraction by Design: Why the Attention Economy Is in a Moral Crisis" [Internet]. 2017. Available: https://www.youtube.com/watch?v=y32RWsZXCL4
43. Wansink B, Painter JE, North J. "Bottomless bowls: why visual cues of portion size may influence intake." Obes Res. 2005;13: 93–100.
44. Lockton D, Harrison D, Stanton N. "Design with Intent: Persuasive Technology in a Wider Context." Lecture Notes in Computer Science. pp. 274–278.
45. Rhoads, K. "Marwell & Schmitt's 16." [Internet]. [cited 5 June 2018]. Available: http://www.workingpsychology.com/marwell.html
46. Marwell G and Schmitt D. Dimensions of Compliance-Gaining Behavior: An Empirical Analysis on JSTOR. In: American Sociological Association [Internet]. [cited 31 May 2018]. Available: http://www.jstor.org/stable/2786181
47. Dillard JP. "The Current Status of Research on Sequential-Request Compliance Techniques." Pers Soc Psychol Bull. 1991;17: 283–288.
48. Rhoads K. "53 Persuasion Tactics: The Levine & Wheeless Taxonomy" [Internet]. [cited 25 May 2018]. Available: http://www.workingpsychology.com/levine.html
49. Cialdini RB. "Influence: The Psychology of Persuasion." Harper Collins; 1993.
50. King, W. "The persuasive effect of graphics in computer-mediated communication." Computers in Human Behavior. Pergamon; 1991;7: 269–279.
51. Cialdini RB, Demaine LJ, Sagarin BJ, Barrett DW, Rhoads K, Winter PL. "Managing social norms for persuasive impact." Social Influence. 2006;1: 3–15.
52. Tormala ZL, Briñol P, Petty RE. "When credibility attacks: The reverse impact of source credibility on persuasion." J Exp Soc Psychol. 2006;42: 684–691.
53. Kohavi R, Deng A, Longbotham R, Xu Y. "Seven rules of thumb for web site experimenters." Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining - KDD '14. 2014. doi:10.1145/2623330.2623341
54. Evangelista B. "Check your phone 86 times a day? Tech insiders say that's by design." In: San Francisco Chronicle [Internet]. San Francisco Chronicle; 4 May 2018 [cited 24 May 2018]. Available: https://www.sfchronicle.com/business/article/Check-your-phone-86-times-a-day-Tech-insiders-1288

8512.php
55. Contempo BETA [Internet]. [cited 25 May 2018]. Available: https://contempo.the-social-edge.com/
56. Rosen LD, Mark Carrier L, Cheever NA. "Facebook and texting made me do it: Media-induced task-switching while studying." Comput Human Behav. 2013;29: 948–958.
57. Tufekci Z. "Not This One." Am Behav Sci. 2013;57: 848–870.
58. Whittaker S, Kalnikaite V, Hollis V, Guydish A. "Don't Waste My Time." Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems - CHI '16. 2016. doi:10.1145/2858036.2858193
59. Moëll B, Kollberg L, Nasri B, Lindefors N, Kaldo V. "Living SMART — A randomized controlled trial of a guided online course teaching adults with ADHD or sub-clinical ADHD to use smartphones to structure their everyday life." Internet Interventions. 2015;2: 24–31.
60. Bator RJ, Cialdini RB. "New Ways to Promote Proenvironmental Behavior: The Application of Persuasion Theory to the Development Of Effective Proenvironmental Public Service Announcements." J Soc Issues. 2000;56: 527–541.
61. Fogg BJ. "The New Rules of Persuasion" [Internet]. [cited 25 May 2018]. Available: http://captology.stanford.edu/wp-content/uploads/2015/02/RSA-The-new-rules-of-persuasion.pdf
62. HabitLab for Chrome - Build better habits online! [Internet]. [cited 25 May 2018]. Available: https://habitlab.stanford.edu/
63. Williams J. "Stand out of our Light: Freedom and Resistance in the Attention Economy." Cambridge University Press; 2018.
64. Fogg BJ. The Behavior Grid: 35 ways behavior can change. Proceedings of the 4th International Conference on Persuasive Technology - Persuasive '09. 2009. doi:10.1145/1541948.1542001
65. Fogg, BJ. "The Behavior Grid: 15 Ways Behavior Can Change" [Internet]. [cited 24 May 2018]. Available: http://behaviorgrid.org/
66. Rhoads K. "Ethics of Influence & Persuasion" [Internet]. [cited 24 May 2018]. Available: http://www.workingpsychology.com/ethics.html
67. Skinner BF. "Operant behavior." American Psychological Association. 1963; doi:10.1037/h0045185
68. LinkedIn Premium [Internet]. [cited 2 Jun 2018]. Available: https://premium.linkedin.com/
69. Halko S, Kientz JA. "Personality and Persuasive Technology: An Exploratory Study on Health-Promoting Mobile Applications." Persuasive Technology. Springer, Berlin, Heidelberg; 2010. pp. 150–161.
70. "Maslow's hierarchy of needs." Oxford Reference. [Internet]. [cited 25 May 2018]. Available: http://www.oxfordreference.com/view/10.1093/oi/authority.20110803100138279
71. Forehand M. "Bloom's Taxonomy" [Internet]. 11 Jul 2012 [cited 2 Jun 2018]. Available: https://www.d41.org/cms/lib/IL01904672/Centricity/Domain/422/BloomsTaxonomy.pdf
72. Thaler RH, Sunstein CR. "Nudge: Improving Decisions About Health, Wealth, and Happiness." Penguin; 2009.
73. Williams J. "Orwell, Huxley, Banksy." In: Rough Consensus [Internet]. 24 May 2014 [cited 1 Jun 2018]. Available: https://blogs.oii.ox.ac.uk/roughconsensus/2014/05/orwell-huxley-banksy/
74. Barker C. "Cultural Studies: Theory and Practice." SAGE Publications; 2003.
75. Bourdieu P, Nice R. "Outline of a Theory of Practice." Cambridge: Cambridge University Press; 1977.
76. Bourdieu P. "The Logic of Practice." Stanford University Press; 1990.
77. Hurrelmann K. "Social Structure and Personality Development: The Individual as a Productive Processor of Reality." CUP Archive; 1988.

78. Rosenberg M, Confessore N, Cadwalladr, C. "How Trump Consultants Exploited the Facebook Data of Millions" New York Times. [Internet]. 17 Mar 2018 [cited 25 May 2018]. Available: https://www.nytimes.com/2018/03/17/us/politics/cambridge-analytica-trump-campaign.html
79. Internet Research Agency Indictment. In: U.S. Department of Justice [Internet]. 16 Feb 2018 [cited 1 Jun 2018]. Available: https://www.justice.gov/file/1035477/download
80. Meyer D. Facebook Is "Ripping Apart" Society, Former Executive Warns. In: Fortune [Internet]. [cited 2 Jun 2018]. Available: http://fortune.com/2017/12/12/chamath-palihapitiya-facebook-society/
81. Harris, Sam. "Waking Up Podcast #71 — What is Technology Doing to Us? | Sam Harris." [Internet]. [cited 25 May 2018]. Available: https://www.youtube.com/watch?v=jlPF9_1VIso
82. Gallagher B. "Modern Media Is a DoS Attack on Your Free Will." 2017; Available: http://nautil.us/issue/52/the-hive/modern-media-is-a-dos-attack-on-your-free-will
83. Marcuse H, Sozialforschung IF. "Some Social Implications of Modern Technology." Zeitschrift für Sozialforschung. 1941;9: 414–439.
84. Fowler, G. "Google wants to cure your smartphone addiction." In: Washington Post [Internet]. The Washington Post; 8 May 2018 [cited 1 Jun 2018]. Available: https://www.washingtonpost.com/news/the-switch/wp/2018/05/08/google-wants-to-cure-your-smartphone-addiction/
85. Williams P, Fitzsimons GJ, Block LG. "When Consumers Do Not Recognize "Benign" Intention Questions as Persuasion Attempts." J Consum Res. Oxford University Press; 2004;31: 540–550.
86. Lieto A, Vernero F. "Influencing the Others' Minds: An Experimental Evaluation of the Use and Efficacy of Fallacious-Reducible Arguments in Web and Mobile Technologies." PsychNology Journal. 2014;12: 87 – 105.
87. Tang L, Rosales R, Singh A, Agarwal D. "Automatic ad format selection via contextual bandits." Proceedings of the 22nd ACM international conference on Conference on information & knowledge management - CIKM '13. 2013. doi:10.1145/2505515.2514700
88. Harris T. "Making Technology Less Manipulative" [Entire Talk]. In: eCorner [Internet]. 3 Nov 2017 [cited 25 May 2018]. Available: https://ecorner.stanford.edu/video/making-technology-less-manipulative-entire-talk/
89. Williams, J. "Extracts from Stand Out of Our Light: Freedom and Persuasion in the Attention Economy - The Nine Dots" Prize. In: The Nine Dots Prize [Internet]. 30 May 2017 [cited 25 May 2018]. Available: https://ninedotsprize.org/extracts-stand-light-freedom-persuasion-attention-economy/
90. De Choudhury M, Jhaver S, Sugar B, Weber I. "Social Media Participation in an Activist Movement for Racial Equality." Proceedings of the . International AAAI Conference on Weblogs and Social Media International AAAI Conference on Weblogs and Social Media. NIH Public Access; 2016;2016: 92.
91. Siewiorek D, Smailagic A, Furukawa J, Krause A, Moraveji N, Reiger K, et al. "SenSay: a context-aware mobile phone." Seventh IEEE International Symposium on Wearable Computers, 2003 Proceedings. doi:10.1109/iswc.2003.1241422
92. Tsikerdekis M, Zeadally S. "Multiple Account Identity Deception Detection in Social Media Using Nonverbal Behavior." IEEE Trans Inf Forensics Secur. 2014;9: 1311–1321.
93. Edelman, Joe. "Is Anything Worth Maximizing? On Metrics, Society, and Selves" [Internet]. Available: http://nxhx.org/maximizing/
94. Holmen M. "A Crisis of Attention - Part II" [Internet]. [cited 2 Jun 2018]. Available: http://www.bibblio.org/blog/a-crisis-of-attention-part-ii

95. Internet Research Agency Indictment Case 1:18-cr-00032-DLF Document 1 [Internet]. 2018 Feb. Available: https://www.justice.gov/file/1035477/download
96. Lewis P. “Fiction is outperforming reality: how YouTube’s algorithm distorts truth.” In: the Guardian [Internet]. 2 Feb 2018 [cited 2 Jun 2018]. Available: http://www.theguardian.com/technology/2018/feb/02/how-youtubes-algorithm-distorts-truth
97. Harris, Sam. “Waking Up Podcast #71 — What is Technology Doing to Us? | Sam Harris.” In: Sam Harris [Internet]. [cited 25 May 2018]. Available: https://samharris.org/podcasts/what-is-technology-doing-to-us/
98. Tristan Harris on Twitter. In: Twitter [Internet]. [cited 25 May 2018]. Available: https://twitter.com/tristanharris/status/913130707978936321
99. Tufekci, Zeynep. “It’s the (Democracy-Poisoning) Golden Age of Free Speech.” In: WIRED [Internet]. WIRED; 16 Jan 2018 [cited 25 May 2018]. Available: https://www.wired.com/story/free-speech-issue-tech-turmoil-new-censorship/
100. McNamee R, Longman M, LeTourneau N, Atkins D. “How to Fix Facebook—Before It Fixes Us.” In: Washington Monthly [Internet]. 7 Jan 2018 [cited 2 Jun 2018]. Available: https://washingtonmonthly.com/magazine/january-february-march-2018/how-to-fix-facebook-before-it-fixes-us/
101. Sunstein, Cass. “How Much Is it Worth to Use Facebook? A Behavioral Perspective” Harvard Law Review. [Internet]. [cited 28 May 2018]. Available: https://blog.harvardlawreview.org/how-much-is-it-worth-to-use-facebook-a-behavioral-perspective/
102. Loder T, Van Alstyne M, Wash R. “An economic answer to unsolicited communication.” Proceedings of the 5th ACM conference on Electronic commerce - EC ’04. 2004. doi:10.1145/988772.988780
103. Edelman J. “Can software be good for us? – Human Systems.” In: Medium [Internet]. Human Systems; 17 Jan 2018 [cited 25 May 2018]. Available: https://medium.com/what-to-build/dear-zuck-fd25ecb1aa5a
104. Edelman, Joe. “Human Values” [Internet]. Available: https://medium.com/what-to-build/human-values-a-quick-primer-b01ef9617925
105. Edelman J. “What’s Next? – Human Systems” In: Medium [Internet]. Human Systems; 6 Sep 2017 [cited 2 Jun 2018]. Available: https://medium.com/what-to-build/whats-next-4b4d00bd9403
106. Davis J. “Design methods for ethical persuasive computing.” Proceedings of the 4th International Conference on Persuasive Technology - Persuasive ’09. 2009. doi:10.1145/1541948.1541957
107. “Staying free in a world of persuasive technologies” Oxford Internet Institute. [Internet]. 29 Jul 2013 [cited 25 May 2018]. Available: https://www.oii.ox.ac.uk/blog/staying-free-in-a-world-of-persuasive-technologies/
108. Neisser U. “Five kinds of self-knowledge.” Philos Psychol. 1988;1: 35–59.
109. Google. “Today we turn 10!” In: Google Lat Long [Internet]. [cited 25 May 2018]. Available: https://maps.googleblog.com/2015/02/today-we-turn-10.html
110. Introna L. “Phenomenological Approaches to Ethics and Information Technology.” 2005; Available: https://plato.stanford.edu/archives/sum2011/entries/ethics-it-phenomenology/#2
111. Chivers T. “The story of Google Maps.” In: Telegraph.co.uk [Internet]. 4 Jun 2013 [cited 25 May 2018]. Available: https://www.telegraph.co.uk/technology/google/10090014/The-story-of-Google-Maps.html
112. Perez S. “Google Maps is turning its over a billion users into editors.” In: TechCrunch [Internet]. 21 Jul

2016 [cited 25 May 2018]. Available: http://social.techcrunch.com/2016/07/21/google-maps-is-turning-its-over-a-billion-users-into-editors
113. Neyfakh-Reporter L. “Do our brains pay a price for GPS?” The Boston Globe. In: BostonGlobe.com [Internet]. 18 Aug 2013 [cited 25 May 2018]. Available: https://www.bostonglobe.com/ideas/2013/08/17/our-brains-pay-price-for-gps/d2Tnvo4hiWjuybid5UhQVO/story.html
114. Heidegger M. “The Question Concerning Technology, and Other Essays.” Harper Collins; 1982.
115. Michalos AC, Borgmann A. “Technology and the Character of Contemporary Life: A Philosophical Inquiry.” Noûs. 1986;20: 573.
116. Zhang H, Zherdeva K, Ekstrom AD. “Different “routes” to a cognitive map: dissociable forms of spatial knowledge derived from route and cartographic map learning.” Mem Cognit. 2014;42: 1106–1117.
117. Dobbs D. “Are GPS apps messing with our brains?” In: Mother Jones [Internet]. 19 Nov 2016 [cited 25 May 2018]. Available: https://www.motherjones.com/media/2016/11/gps-brain-function-memory-navigation-maps-apps/
118. Ritchey T. “Swedish Morphological Society” [Internet]. [cited 25 May 2018]. Available: http://www.swemorph.com/
119. Shrager, Jeff. “SymSys 245: Interaction Analysis Glossary.” Lecture notes, Stanford University. [cited 3 June 2018].
120. Shrager, Jeff. “Cognition for readers with short attention spans.” Lecture notes, Stanford University. [cited 3 June 2018].
121. Perez, Sarah. “Apple to launch its own ‘digital health’ features in iOS 12, says report.” In: TechCrunch [Internet]. 1 June 2018 [cited 4 June 2018]. Available: https://techcrunch.com/2018/06/01/apple-to-launch-its-own-digital-health-features-in-ios-12-says-report/
122. Eyal, Nir. “Why Do Google and Apple Want You to Use Your Phone Less?” In: Nir & Far [Internet]. [cited 4 June 2018]. Available: https://www.nirandfar.com/2018/06/google-apple-less-phone-use.html
123. Maguire, Eleanor A. et. al. "Navigation-related structural change in the hippocampi of taxi drivers." PNAS April 11, 2000. 97 (8) 4398-4403; https://doi.org/10.1073/pnas.070039597
124. Brignull, Harry. “Dark Patterns.” [Internet] [cited 5 June 2018]. Available: darkpatterns.org
125. Friedman B, Kahn P, Borning A. “Value Sensitive Design: Theory and Methods.” UW CSE Technical Report 02-12-01. 12 February 2001. [Internet]. Available: http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.11.8020&rep=rep1&type=pdf

Appendix A: Nudget Intervention Text

Nudget Location	Nudget Text	Character Count
Near a video in the feed	**Unauthorized Autoplay:** For Facebook, your time is $$$, so it auto-plays all videos without consent to grab your attention.	123
Near an advertisement post in the feed	**Clever Ad Placement:** Facebook shapes ads like your friends' posts, making it hard to ignore & selling your attention to the highest bidder.	139
Below the "Home" button	**Glued Reminders:** Facebook affixes these buttons here like a billboard reminding you to keep clicking for more and more…	119
Below the "Notifications" button	**Quantified FOMO:** Facebook shows a number to tell you exactly how many "important" things you're missing out on. What % is actually important?	142
Below the "Notifications" button	**Foot-in-the-door:** Seemingly harmless, Facebook has a tiny request: check your notifs! But we're bad at forecasting time, and 5 mins later, you'll probably still be here.	171
Below a post with comments	**Endless Conversation:** Facebook always enables comments, so endless chatter arouses your curiosity even if it's not worth your time.	127
Below a post with comments	**Comments:** Comments = social approval. Without them, there's no incentive to post content & help Facebook earn even more $$$.	124
Below the "Event" section of the right scrollable feed	**Social Proof:** With events, Facebook lists your friends to tap into your tendency to accept "social proof" as validity. These friends might not even go!	151
Below the "Like" button of a post	**Lucrative Likes**: Likes = self-validation. Facebook designed likes so you'd click to see friends' names. It's their one-click ticket to buying more of your time.	160
To the left of the side feed	**Silly Side Feed**: With this scrollable feed, Facebook creates an artificial sense of "breaking news" depicting a permanent flurry of activity.	141
Near the top of the feed	**Bottomless Bowl**: All-you-can-eat buffets increase food intake. Similarly, your Facebook feed never ends, so they keep you scrolling.	133
Near the top of the feed	**Mystery Algorithm:** Even thousands of FB employees don't really know how the feed algorithm shaping your thoughts everyday works.	128
Near the top of the feed	**Slot Machine Feed**: Facebook's algorithm makes useful posts appear intermittently so we keep refreshing, never sure when we'll receive that dopamine-activating prize.	165
Near the left menu list	**Missing Options:** The more we rely on Facebook, the more we limit our sense of the menu to what is shown here. What about options for human connection NOT on the menu?	166

Near the "Pokes" feature on the left menu (pictured above)	**Reciprocity**: Facebook created pokes to leverage our vulnerability to social reciprocity...whereas the initiator's poke was evoked by a web of manipulation in the first place.	174
Near the "Messenger" icon on the top menu bar	**Intentional Interruption:** Facebook knows an instantly delivered message is more likely to elicit a response.	108
Above the messenger chatbox (bottom right of screen)	**Receipts:** Facebook always tells the sender when you saw their message, increasing social pressure to respond.	109
Below the rightmost arrow for the dropdown on the top menubar	**Don't Leave:** Compare how easy it is to see your feed vs. deactivate your account in settings. Facebook purposefully makes it hard for you to exit.	146
Near the top of the feed	**Digital Trails:** Facebook uses your every click to understand your preferences in layout and content. Mega-personalization keeps you clicking.	141
Above a post that begins "...was tagged in a photo"	**Tagging = Social Approval**. Facebook makes you think a tag is an organic action - whereas the tagger is often prompted by a targeted notif.	138
Near a post with a photo	**Self-Representation:** Facebook makes it way too easy to conflate a perfect self-curated photo with real life.	108
Below the "Notifications" button on the top menubar	**Loud Red:** Facebook knows that red is the optimal color to get you to click. Red biologically draws you in. It's a race to the bottom of the brainstem.	150

Appendix B: Nudget Transfer Assessment

To measure transfer learning in our Nudget study, we included the following 3 images of LinkedIn screens, each divided into 3 regions (A, B, C). We left 6 blank, numbered spaces below each screen and requested participants to identify and write out all the persuasive design techniques they think LinkedIn used in that region. We did not blur the images so as to be as realistic as possible, and instead merely redacted last names.

Screen 1:

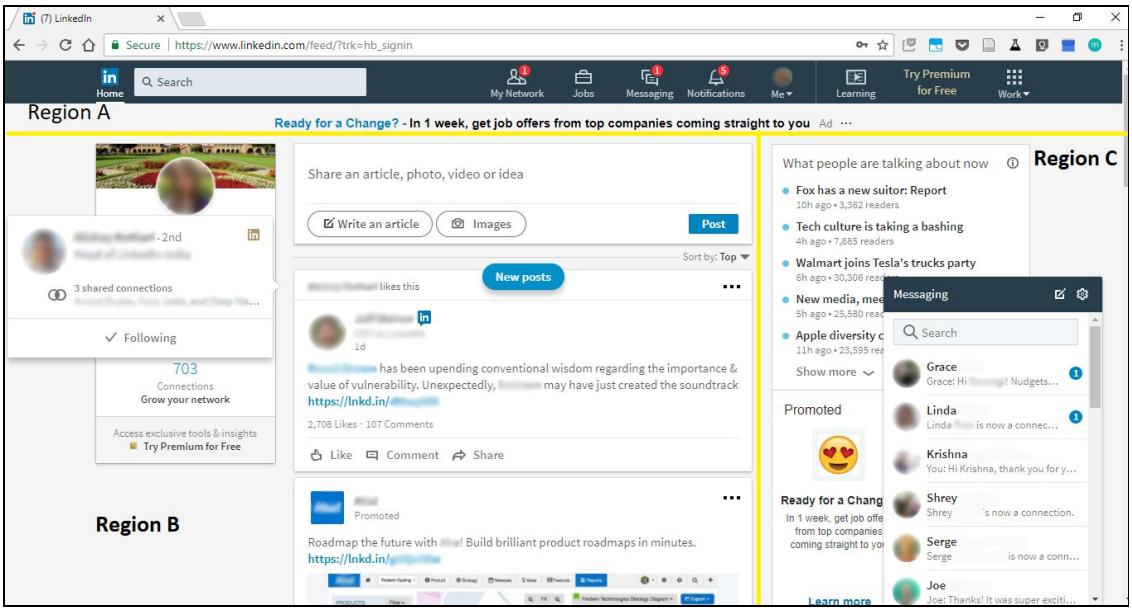

Screen 2:

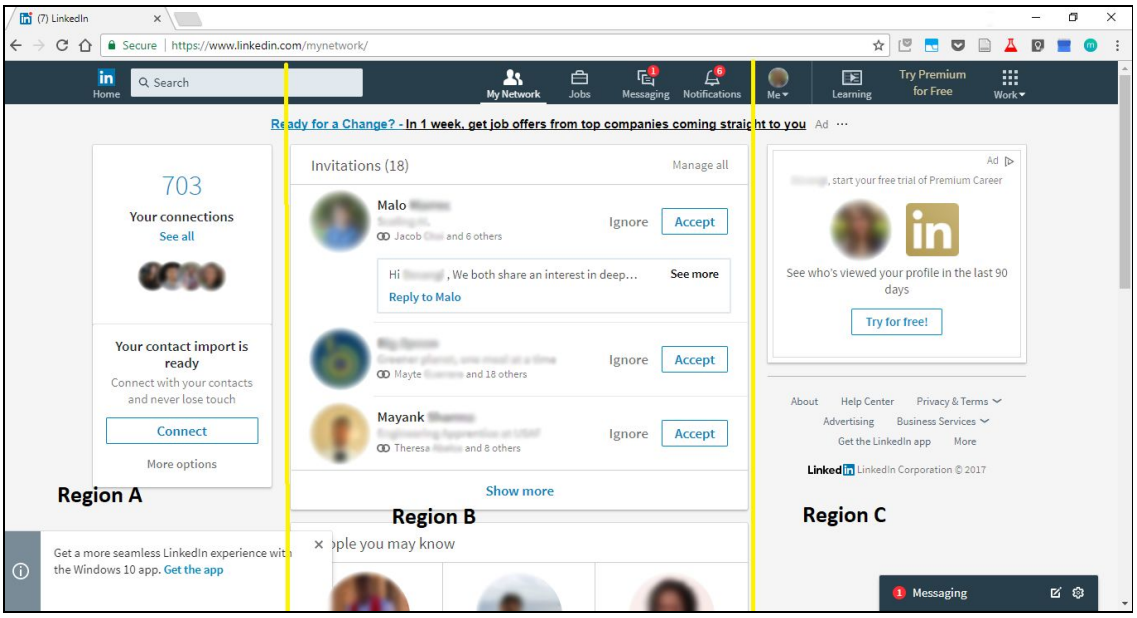

Screen 3:

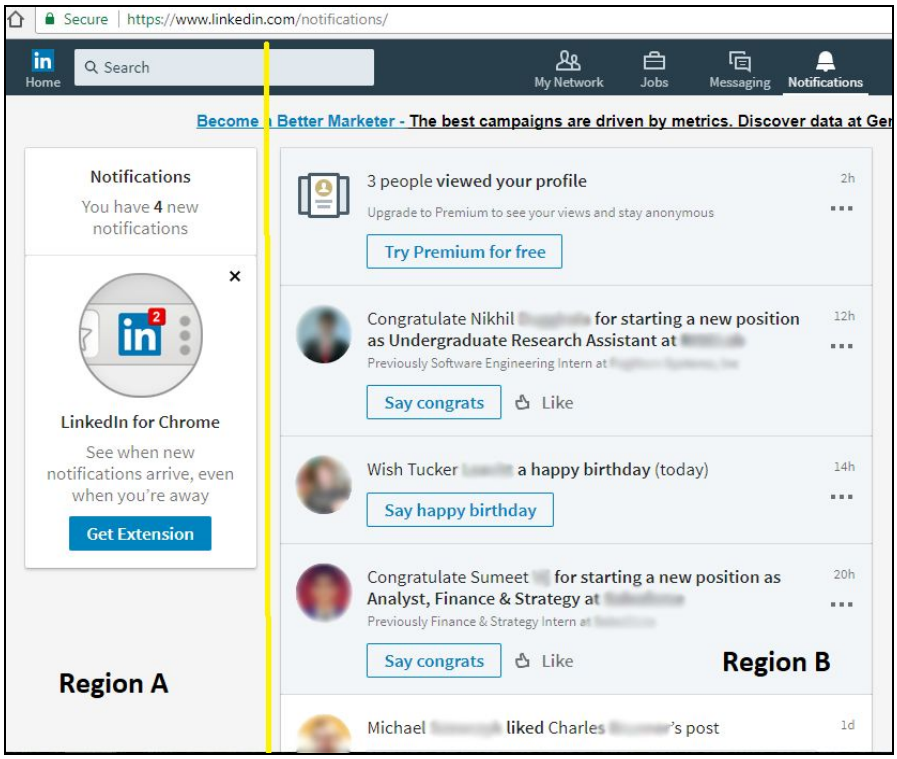
Secure | https://www.linkedin.com/notifications/
Home
Search
My Network
Jobs
Messaging
Notifications
Become a Better Marketer - The best campaigns are driven by metrics. Discover data at Ger
Notifications
You have 4 new notifications
LinkedIn for Chrome
See when new notifications arrive, even when you're away
Get Extension
3 people viewed your profile
2h
Upgrade to Premium to see your views and stay anonymous
Try Premium for free
Congratulate Nikhil for starting a new position as Undergraduate Research Assistant at
12h
Previously Software Engineering Intern at
Say congrats
Like
Wish Tucker a happy birthday (today)
14h
Say happy birthday
Congratulate Sumeet for starting a new position as Analyst, Finance & Strategy at
20h
Previously Finance & Strategy Intern at
Say congrats
Like
Region A
Region B
Michael liked Charles 's post
1d

Appendix C: LinkedIn Persuasive Design Techniques

Below, we list all the persuasive design techniques identified by our participants in the pre-transfer assessments. To constitute a unique persuasive design technique, both of the following factors must be different from all other techniques: 1. the vehicle (i.e., a UI element or group of elements), and 2. a reason why that element creates persuasion. The second factor took many forms; for instance, participants described their personal interactions with an element on the page, the underlying social strategy at play, network effects, visual theory, and hypotheses about how users interact with the element in a way that aligns with the website's bottom line goals, to name a few.

Screen #	#	Persuasive Vehicle	Method of Persuasion
1A	1	Notification badges on the horizontal toolbar for "notifications," "messages," and "network"	Makes you want to click and see new notifications (arouses curiosity)
1A	2	Red color of notification badges on the horizontal toolbar	Stands out/catches your attention/indicates urgency in order to redirect your clicks to other people's or companies pages.
1A	3	Number on the notification badges on the horizontal toolbar	Makes it feel like a to-do list and makes you want to get the numbers to 0 (arouses our "base desire for having order instead of chaos"
1A	4	Intermittent variable notifications	The delivery schedule of notifications is varied and intermittent, which keeps it changing and thus interesting.
1A	5	Textual ad at the top "Ready for a change…"	Tries to get you to click on that page by appearing organic and relevant.
1A	6	"Try Premium for free" at the top right of the horizontal toolbar	Uses its prime accessible location near the first things you see at the top to get you to pay more money,
1A	7	"Try Premium for free" at the top right of the horizontal toolbar	Places this action on equal footing with the other calls to action, making it seem like a normal action that most people would take.
1A	8	Gold color of "Try Premium for free"	Evokes the concept of royalty and makes this feature seem even more special.
1A	9	The name "Premium" for the monetized version	Makes the monetized version seem like a special offer.
1A	10	Auto-billing after your free trial has expired	The "free" offer will automatically turn into a payment unless you actively cancel, which makes it harder to remember.
1A	11	Misnomer of the "Work" feature on the horizontal toolbar	Makes you think it's about jobs for you but is actually a collection of LinkedIn products they want you to use.
1A	12	"Learning" feature	Removes any guilt of time spent on that feature due to use of a positive, productive word
1A	13	"My Network" element on the horizontal toolbar	Encourages you to click by arousing your curiosity about what others are up to.
1A	14	"Messaging" element on the horizontal toolbar	Makes it appear like all messages are organic conversation in order to persuade you to check the sponsored ones as well.
1A	15	Ordering of elements on the horizontal toolbar	Placing "Notifications" next to your photo (which naturally attracts your attention first) makes you more likely to click on Notifications rather than "My Network" (Fitt's Law!) which funnels your prescribed clicks.

1A	16	Number of elements on the horizontal toolbar	Breaking notifications into so many categories increases your number of clicks
1A	17	Simplicity of icons	Makes them “friendly” looking at not “intimidating”
1A	18	Placing “My Network” first in reading order (left-->right) on the horizontal toolbar	Foregrounds the rewarding nature of connecting with someone, an important theme throughout the site (Maslow’s hierarchy of needs [70]).
1A	19	Including a photo of yourself on the horizontal toolbar	Generates a feeling of self-worth or self-identity which frames your reference of mind and feels like it gives you agency when interacting with the rest of the site
1A	20	Including a photo of yourself on the horizontal toolbar	Makes you want to check out your own profile - “like mounting a mirror next to your desk.”
1A	21	“Me” section	Foregrounds the idea of personal identity and representation whereas it is primarily a way to get more data about you to drive personalization.
1A	22	“Jobs” button on the horizontal toolbar	Mismatch between expectations and reality: this button draws you in by making you think the feature’s main purpose is to add value to your job search whereas the site’s priorities are split between you and the companies that advertise on it.
1A	23	“Jobs” button associated email campaign	The site sends daily emails saying it found jobs for you to pull you back to the jobs feature.
1A	24	Placement of the text ad “Ready for a change…”	Placing this ad about jobs near the “Jobs” icon makes it an outlet to satisfy your use case of searching for job, making it more likely to be clicked.
1A	25	Placement of the text ad “Ready for a change…”	Placing this at the top means you are forced to read it of you look at things lower on the site, even if you might not want to read the ad.
1A	26	No ‘X’ button to remove ads	Forces you to read ads, reducing your agency
1A	27	Placement of “Home” button	As the leftmost button, provides an easy way to refresh the page to load new content to the feed, which is an unending cycle
1A	28	Overall color scheme	LinkedIn uses colors that are known to be “trustworthy” to gain your trust.
1A	29	Different color of horizontal toolbar	Makes it stand out as the prominent elements
1A	30	Showing your number of notifications in the tab on Google Chrome	Attracts attention towards the LinkedIn tab by making you wonder what kind of notifications await you.
1A	31	Showing a little red dot on the LinkedIn icon in the tab on Google Chrome	Attracts attention towards the LinkedIn tab by making you wonder what kind of notifications await you.
1B	32	“New posts” pill	Prompts loading more updates, which increases time spent on site
1B	33	Interaction-based expansion: re-expanding of comment box when you click on an article and return to feed	Suggests that you should write a comment on your thoughts
1B	34	Likes and comments	As a content creator, these keep you coming back to the site to see who has responded to your post
1B	35	Comments	Keeps you reading below a post even when you’re done reading the post itself.
1B	36	Interaction-based recommendations: after an article click, you see similar articles or after following someone you see similar people to follow	Suggests that you should do similar actions

1B	37	Personalization of all content in the feed	You are only shown content first that you are likely to find interesting according to the site's algorithm
1B	38	No choice of what you want to see in the feed - LinkedIn chooses all the content	Allows them to display what they think is best for their success metrics (my clicks)
1B	39	"Akshay likes this": text indicating your friend likes a post	Increases likelihood of post engagement via social proof: "If my friend X liked this, maybe I'll like it too."
1B	40	"Akshay likes this": text indicating your friend likes a post	Gives you a false sense of your friends' digital presence on the site
1B	41	Existence of ads as part of the news feed	Gets you to click on ads
1B	42	Enticing text at the beginning of the Aha! advertisement	Makes you want to keep reading (arouses curiosity)
1B	43	Unverified content in news articles	Not fact-checking the content in news articles allows for sensationalized content in feeds, which gets more clicks.
1B	44	Calling advertisements like "Aha!" promotions	Makes it appear like ads are special deals.
1B	45	De-emphasis of the word "promoted" on a promoted post	Makes ads look like other content in the feed.
1B	46	Showing profile pictures along with posts	Makes the posts seem more personal and also indicates that content is personalized for you.
1B	47	Having the second post cut off halfway	Encourages us to scroll down to see the rest of it below the fold.
1B	48	Including blue links in posts	Encourages us to click on people and things mentioned in posts
1B	49	Bottomless bowl: feed is an infinite scroll list	Keeps you endlessly scrolling to the bottom.
1B	50	Adoption of "feed" structure from Facebook	Borrows mindless scrolling behavior from Facebook by borrowing the feed element
1B	51	Variable reinforcement temporal delay upon reload of feed	Builds anticipation for the content that's loading
1B	52	Default "sort by: top" ordering posts on the feed by popularity not recency	Makes popular articles even more popular (snowball effect)
1B	53	Box at the top of feed providing ways to "write an article," etc.	Makes content generation (which otherwise requires commitment and time) very accessible, and increases the pool of content available to be shown in feeds.
1B	54	Presence of "jobs recommended for you" in the news feed	Brings personalized job recommendations outside of the "Jobs" page to the feed; increased personalization means more relevance and thus aroused interest.
1B	55	"Like" button	A one-click indication of content relevance which allows LinkedIn to improve personalization in the future
1B	56	Left column stays fixed	Allows affixed cues to continue to cue you as your scroll
1B	57	Promoting famous people	Uses your friends' names as bait to get more followers to already famous people
1B	58	Tooltip with a summary profile when you hover over someone's name	Gets you to easily click on their profile to learn more
1B	59	"3 shared connections" in the tooltip that appears when hovering over someone's name	Makes you want to know who those people are by clicking on their profiles as well.
1B	60	Brown "in" badge for Premium users	Displayed like a status symbol which makes it compelling to want and respect, nudging you towards a Premium membership.
1B	61	Displaying your number of connections	Makes this a metric that you want to maximize

1B	62	Displaying a message encouraging you to "access exclusive tools"	Makes these tools (such as finding out how many people viewed your profile) a metric of interest that you want to maximize
1B	63	Displaying "Try Premium for Free" message	Makes it easy for you to maximize insights by clicking and subscribing to the paid version.
1B	64	Placement of "Grow your network" message	Placing "grow your network" near your number of connections makes it seem like growing this number is as easy as the click of a button.
1B	65	"Grow your network" message	Makes you feel like you're not doing enough and could do more through LinkedIn
1C	66	New message notification	Cues user to click on the message
1C	67	Highlighting unread messages	Prompts you to read them first and check them off your list
1C	68	Chatbox element (as opposed to chat being a different tab)	Allows chatbox to be ubiquitously present, and unable to be fully rid of. This allows message responses without leaving engagement on a particular page.
1C	69	Chatbox element	Makes it appear like all messages are organic conversation in order to persuade you to check the sponsored ones from recruiters as well.
1C	70	Chatbox element	Drives conversion to the site, since you now have a concrete investment in a communication pipeline through the site.
1C	71	Size of chatbox	Doesn't completely cover the articles and ads, which lets you be distracted to click on those even as you write a message.
1C	72	Size of chatbox	Makes messages prominently displayed, which conveys a sense of urgency to reply
1C	73	Message preview size	Not showing the full message arouses curiosity about the rest of the message and encourages users to click on the message.
1C	74	"Online" status of messaging contacts	Makes it easier to start a conversation with a contact
1C	75	Automatic popup of chat box when a contact is online and you visit their profile	Makes it easier to start a conversation with a contact
1C	76	Automatic new message when you make a connection	Prompts you to start sending messages to each new contact
1C	77	The text "What people are talking about now" at the top of the trending articles	Uses social proof to show that you can "stay in the loop" with others by focusing your attention on this section.
1C	78	The text "What people are talking about now" at the top of the trending articles	Making it seem like this content is driven purely by others hides the fact that LinkedIn actually curates it
1C	79	Trending topics list	Uses social proof (what others are interested in) to convince you to spend longer on the site
1C	80	Number of readers listed below each article in the trending topics	Greater number of people increases your likelihood of clicking (bandwagon effect)
1C	81	Blue dot next to each trending topic	Similar to notification badges, makes you want to click through all the headlines to get rid of the blue dot.
1C	82	Timestamp on each news item	Entices you to make sure you are caught up with news over a time period
1C	83	"Show more" link below the news items	Encourages us to interact with the box by expanding it downward, and keeps us hooked to the new content since we purportedly chose to see more content using our own agency
1C	84	Tracking of your clicks on the trending topics	Leads you to see content that you are more likely to click on

1C	85	Default status of messaging tab as closed	Lets you see articles and ads and allow their persuasion to attract you first so you click on them
1C	86	Right column stays fixed	Allows affixed cues to continue to cue you as your scroll
1C	87	Incentivizing text like "<USERNAME>, are you ready for a job"	Directly appeals to your need of finding a job (Maslow's hierarchy of needs [70])
1C	88	Personalized ads matching your photo with a company logo	Makes you wonder what it would be like to have that company on your profile (which is an artificially constructed representation of self in the first place)
1C	89	Promoted ad	Encourages you to click on content that will make the site money.
1C	90	Use of huge emoji in promoted ad	Distracts you from other useful information to encourage you to click on the promoted content.
1C	91	Blue "learn more" text below promoted ad	The color blue makes it enticing to click on the promotion.
1C	92	"In 1 week, get job offer…"	Enticing (and likely false) promises in advertisements encourage clicking.
2A	93	Showing your number of notifications in the tab on Google Chrome	Attracts attention towards the LinkedIn tab by making you wonder about the significance of the number
2A	94	Toast prompting user to download Windows app	Opens the door to moore notifications even while not on LinkedIn, allowing the site to pull you away even if you're doing something else and spend more time on LinkedIn
2A	95	Text in toast prompting user to download Windows app	Use of words like "seamless" makes it seem like users will have a better experience if they get the app.
2A	96	Placement of toast prompting user to download Windows app	Layout breaks the grid of LinkedIn since the toast crosses over into the next region, which gets us to notice it.
2A	97	Size of "X" button on toast prompting user to download Windows app	The relatively smaller size of the "X" button makes it harder to close the toast.
2A	98	Gray bar & info icon on toast prompting user to download Windows app	Makes it seem like the toast needs your attention, and the icon makes it look especially important
2A	99	Prompt to "See all" connections	Proposes and normalizes a seemingly ridiculous action of scrolling through hundreds of people you've already connected with
2A	100	"Your contact import is ready" text	Making it seem like a difficult, helpful task is already done hides the fact that this is LinkedIn's way of getting more information and furthering its network effects.
2A	101	"More options" below the connect button	Makes you feel another click invested in the process if you click on it (a process whose purpose is to get more data).
2A	102	Textual framing "never lose touch"	Makes user think site is suggesting an action in the user's best interest whereas adding more contacts is beneficial to keeping you on the site (more content for feed, more possible people to message and interact with).
2A	103	Number of connections displayed	Foregrounding this key metric reinforces your commitment to the site at a user, showing how much investment you've already made (sunk cost)
2A	104	Font size of number of connections displayed	Makes this seem like a measure of how effectively you are using LinkedIn.
2A	105	Thumbnail photos of your connections	Makes you want to click on their profiles and want to add even more people
2A	106	Location of thumbnail photos of connections adjacent to "see all"	Tiered effect: human faces attract you to the photos, which makes you more likely to click on "See all"

2A	107	Repeated use of the word “connect”	Makes you want to connect with people (need for human connection - Maslow’s hierarchy of needs)
2A	108	Repeated use of the word “your” (“your connections,” “your contact(s)”	Makes the product feel more personal, improving your relationship with it.
2A	109	White box on gray background	Color contrast focuses attention more efficiently on white box
2A	110	Consistent stylization (blue text and same font) for all calls to action “See all,” “Get the app,” “Connect,”	Makes it easier to spot the calls to action
2B	111	Text ad at the top of the screen	This ad persists throughout screens, making it a permanent fixture that appears more often with a greater chance of being clicked.
2B	112	Personalized invitation snippet: “See more” prompt	Snippet arouses curiosity and “See More” gives users a way to satisfy that curiosity.
2B	113	Showing a message from people below their invitation	Gives the invitation a more personal context and incentivizes you to connect with them by making you feel like you must reply.
2B	114	Showing the positions of people who sent you invitations below their names	Serves as a hook to make you want to learn more about them.
2B	115	Showing profile photos	Makes everyone seem more human, which draws upon the norms and pressures of human interaction.
2B	116	Showing mutual connections you share with people who sent you invitations	Serves as a hook to make you want to learn more about them because of your shared network (social proof)
2B	117	Showing mutual connections you share with people who sent you invitations	Helps you develop more of a personal bond with a potential connection.
2B	118	Personalized suggestions of who to connect with (“People you may know”)	Creates a goal that you then embrace; makes the norm of behavior to want to increase the number of connections you have by exploring the network of people in your network.
2B	119	Personalized suggestions of who to connect with (“People you may know”)	Increases the likelihood of you making connections with more people, which accelerates LinkedIn’s network effects.
2B	120	Bottomless bowl: list of people you know is an infinite scroll list	Increases the number of new connections you might makes which is linked to how much time you spend on the site.
2B	121	“Accept” button blue & highlighted vs. “Ignore” button grayed out	Makes you more likely to notice and click on the accept button, which is the desired action that increases LinkedIn’s network effect, which leads to more time invested in the site
2B	122	Loaded word “ignore”	Makes it harder to choose this action, which has a negative connotation.
2B	123	Loaded word “invitation”	Makes it more tempting to open a connection request by calling it an invitation rather than a reviews.
2B	124	“Manage all” feature	Allows you to accept all invitations instantaneously, which accelerates LinkedIn’s network effects.
2B	125	Showing the number of invitations (in parenthesis)	This large number in parentheses looks unappealing and makes us want to decrease the number.
2B	126	Red color of notification badges	Stands out/catches your attention/indicates urgency in order to redirect your clicks to other people’s or companies pages.
2B	127	“Show more” button	By not showing all the relevant details, encourages you to click/scroll further to truly find out more information, given that you’ve come this far. “You can keep going with it forever”
2B	128	Blue color of “Show more” button	Makes it more enticing to click on it (blue stands out compared to the rest of the color palette)

2B	129	“In 1 week, get job offers from top companies coming straight to you” text in the ad	Enticing, vivid language in advertisements shown on LinkedIn prompts clicks
2C	130	Red “new message” badge	Makes you want to click to reopen the chatbox even when it’s closed
2C	131	Custom personalized ad in the right column	Directly forces you to envision yourself in a context that LinkedIn wants.
2C	132	Size of ad in the right column	Makes ad most likely item to clicked whereas useful things like “Help Center” are tiny links (Fitt’s Law)
2C	133	Layout of the “About,” “Help Center” links	This layout is formless and unattractive, which serves to add more emphasis to what we should be looking at: the big square ad.
2C	134	Showing your photo next to company logo	Your own face attracts your eye (hook),
2C	135	Showing your photo next to company logo	Makes you think of what it would be like to work there and have that company on your profile (another place on the site you look at your photo).
2C	136	“See who’s viewed your profile”	Makes you curious (in a narcissistic way) - a manufactured desire that can be mitigated by the solution presented: purchasing LinkedIn Premium. Ironically this rarely improves your professional life.
2C	137	Use of your name “<NAME>, start your free trial”	Makes the monetization request more personal.
2C	138	Promoting LinkedIn ads	LinkedIn ads ensure that you will stay on the site.
2C	139	Repeating the “Try for free” call to action	The recurring nature of this call to action (repeated left to right across the screen) makes it seem more urgent.
2C	140	Echoed stylization of “try for free” button	The “try for free button” uses the blue and white stylization of the “accept” and “connect” buttons, which have already been associated with positive actions.
2C	141	Default status of messaging tab as closed	Lets you see articles and ads and allow their persuasion to attract you first so you click on them
2C	142	Highlighting the “golden” color associated with the premium membership	Brings to mind the “status” that the site has associated the color with
3A	143	Extension download message	Opens to the door to more & faster notifications even if not onsite, which creates more opportunities to pull you back to the platform if you’re online doing something else.
3A	144	Including the number of notifications (“4” written and “2” overlayed on the app icon)	Gets you to look at the ad
3A	145	Glimpse of what the Chrome extension looks like	Helps you imagine what it would look like, which makes it easier to make a decision to get it.
3A	146	“Home” button	Provides an easy exit to more of the feed, which is architected solely by LinkedIn
3A	147	Blue fill of “Get the extension” button	The use of blue is enticing, and filling the buttons with blue makes them the most visually attractive
3A	148	Depiction of notifications	Makes it seem like they are comprehensive reminders whereas LinkedIn selectively chooses notifications
3A	149	“Even when you’re away” text	Makes you feel like you should be thinking about LinkedIn even when you’re not using it.
3A	150	Horizontal text ad crosses over into this vertical region	Guides your eye across the screen from left to right.

3B	151	Text ad at the top of the screen	This ad persists throughout screens, making it a permanent fixture that appears more often with a greater chance of being clicked.
3B	152	Selection of notifications shown	Showing notifications related to connections you already care about increases the likelihood of you checking notifications.
3B	153	Suggested reply buttons (“Say congrats,” “Say happy birthday”)	Lowers the barrier (time you need to spend) to reply
3B	154	Suggested reply buttons (“Say congrats,” “Say happy birthday”)	Makes it seem like a lower effort interaction whereas these buttons actually send a message to the other person
3B	155	Suggested reply buttons (“Say congrats,” “Say happy birthday”)	Makes you feel bad if you don’t follow the “Wish person X” instructions
3B	156	Bolded suggested reply buttons	Entices you to click on the more prominent option
3B	157	Blue color of suggested reply buttons (“Say congrats,” “Say happy birthday”)	Conditions us to associate the color “blue” with “good”
3B	158	Existence of birthday and work milestone reminders	Makes it more likely for you to outsource your memory of birthdays or work milestones to this feature
3B	159	Receiving birthday and milestone wishes	Makes these seem like organic messages to the receiver whereas they were prompted by targeted notifications
3B	160	Text of suggested replies	Makes it seem like you’re building or strengthening your connections.
3B	161	Recording the replies you choose to send	Allows LinkedIn to personalize your feed to rank these connections’ posts higher both in the feed and in these notifications.
3B	162	“Like” button	Allowing you to “like” a post from within your notifications makes them more interactive and gives you a call-to-action for the notifications about posts.
3B	163	Cue to sign up for a premium membership	Gives you a way to satisfy your desire for more features
3B	164	Placement of cue to sign up for a premium membership at the top of the notifications	Draws your attention to your need to know what others think of you first & also what makes the site more money first
3B	165	“3 people viewed your profile” before the Premium call to action	Hooks you by feeding into your curiosity about how others perceive you, and taunts you with your lack of further information
3B	166	Text giving a taste of a premium membership: free trial, stay anonymous	Whets your appetite for the greater number of features to engage with
3B	167	Premium membership itself	Makes you feel like you should make the most of your membership and use the platform more.
3B	168	“X liked Y’s post” notification	Normalizes “liking” behavior and encourages you to do the same.
3B	169	White color of previous notifications	The softer shade of notifications below the fold makes our eyes more comfortable looking at the notifications, encouraging us to spend more time in this section at the end.
3B	170	White color of previous notifications	Contrast with the blue background color of newer notifications, which makes you want to take action to make these more uniform with the others.
3B	171	Including the timestamp of a notification	Conveys the urgency of reacting by counting the number of hours elapsed.