Journal of Advertising

ISSN: (Print) (Online) Journal homepage: https://www.tandfonline.com/loi/ujoa20

Social Value: A Computational Model for Measuring Influence on Purchases and Actions for Individuals and Systems

Dmitri Williams, Euna Mehnaz Khan, Nishith Pathak & Jaideep Srivastava

To cite this article: Dmitri Williams, Euna Mehnaz Khan, Nishith Pathak & Jaideep Srivastava (2022): Social Value: A Computational Model for Measuring Influence on Purchases and Actions for Individuals and Systems, Journal of Advertising, DOI: 10.1080/00913367.2021.2002743

To link to this article: https://doi.org/10.1080/00913367.2021.2002743

Published online: 04 Jan 2022.

JOURNAL OF ADVERTISING
https://doi.org/10.1080/00913367.2021.2002743

Social Value: A Computational Model for Measuring Influence on Purchases and Actions for Individuals and Systems

Dmitri Williams[a], Euna Mehnaz Khan[b], Nishith Pathak[c] and Jaideep Srivastava[b]

[a]University of Southern California, Los Angeles, California, USA; [b]University of Minnesota, Minneapolis, Minnesota, USA; [c]Wargaming, Inc, Austin, Texas, USA

ABSTRACT
Measuring influence of one person on another has applications in advertising and marketing and across the sciences. Most approaches involve inferring influence based on speech and social media. In contrast, this paper takes existing spending data and attributes influence on to the spenders and those likely to have caused their spending. The resulting metric, Social Value, is expressed in units of behavior over time. While a person's total influence on others is called their Social Value, a person's behavior caused by someone else is called their Following Value. These metrics can also be used across an entire community, customer base, or audience, allowing an objective measure of how much spending or other behavior is social versus nonsocial. These measures in turn open up the potential to test interventions and campaigns to measure viral spread as well as overall shifts in social influence. This article presents a computational model for estimating Social Value, as well as validation of the estimation approach in a study involving players of an online game. A noncommercial open-source implementation of the computational model accompanies this paper.

Interactions with others are central to the human experience. Other people are important because we need each other to continue to survive as a species as well as for social relationships. We are well known to look to each other for social proof when deciding on an action (Cialdini 1993), even online (Marwick 2015). It is taken as an active assumption that we impact each other and that collectively these impacts are important. The study of these dynamics is the core of all social science (Backhouse and Fontaine 2014). Within advertising, it has long been taken as a given that people influence each other (Lessig and Park 1978), though measuring it has proved challenging (Berry et al. 2019). Traditional models of lifetime value typically focus on the relationship between the producer and the consumer, but more recent advances suggest also factoring in consumer–consumer factors (Kumar 2018). Nuanced models of lifetime value now include purchasing, referrals, influence, and knowledge value (Kumar et al. 2010). Yet, assessing marketing outcomes is often under-theorized and challenging to measure (Katsikeas et al. 2016). In particular, measuring human-to-human causal ties in the context of purchasing decisions has remained challenging, in part due to the complexity of human relationships and the difficulty of tying them to actual purchasing records.

Indeed, "how" and "why" other people impact us are among the most complex issues of the social sciences, and even "how much" has been challenging to quantify. The 75-year Harvard Grant and Glueck study (Curtin 2018) has demonstrated that relationships and social ties are critical for health and well-being, but such approaches only give us long-term impacts. For practitioners looking at individuals in business applications (and others), there is a need for measuring the impacts of individuals on others in specific contexts. Advertisers and marketers in particular have a need to identify the amount of influence in systems and to identify those consumers within them that might offer points of leverage out into the broader customer base. A large body of literature has

CONTACT Dmitri Williams dcwillia@usc.edu Annenberg School for Communication, University of Southern California, 734 West Adams Boulevard, Los Angeles, CA 90089, USA.
Dmitri Williams (PhD, University of Michigan) is an associate professor, Annenberg School for Communication, University of Southern California.
Euna Mehnaz Khan (BS, Bangladesh University of Engineering and Technology) is a PhD student, University of Minnesota.
Nishith Pathak (PhD, University of Minnesota) is a senior research scientist, Wargaming, Inc.
Jaideep Srivastava (PhD, University of California at Berkeley) is a professor, University of Minnesota.

accumulated examining the electronic word-of-mouth effect and what impact it may have on sales (Lee, Lee, and Feick 2006; Vrontis et al. 2021). Connecting online behaviors to sales has proved challenging, and so much of the efforts lead to purchase intent rather than purchases (Trivedi and Sama 2020). This study explores a different angle by ignoring such speech-based approaches and by focusing on everyday consumers rather than "influencers." It attempts to provide a falsifiable tool for the "how much" and "who" questions as directly applied to sales. This in turn can open up new avenues for understanding the "how" and "why."

The Social Value metric relies on purchasing behaviors rather than social media activity and so is a departure from many current approaches. To explain the approach, we first offer a brief review of the literature on computational social science, networks, and community, followed by the extant approaches to measuring influence in advertising. We then describe the Social Value concept in more detail. The concept is operationalized into a computational model and an open-source software tool, which is then validated in field settings via a large-scale test on purchasing behaviors within a large online video game. As the first empirical test of the method, this setting was chosen for validation rather than broader generalizations and uses in other verticals. The implications of the game-based purchasing behavioral results are provided, along with the potential uses of the tool for leveraging the "ripples on the pond" of real human networks through existing advertising practices. This highlights the need for experimental applications in the service of understanding what kinds of interventions will lead to social increases in spending rippling out from the right targeted consumer (Phelps et al. 2004).

Extant Approaches to Influence in Advertising and the Social Sciences

Foundational work by Friedken and colleagues (Friedkin and Johnson 1999) called for an understanding of both the process of influence as well as the statistical modeling of it. In recent years, advances in computing power and modeling have given rise to powerful techniques to map and model those relationships and processes that undergird human interaction. Network science is the study of who is connected and what the implications are (Monge and Contractor 2003). For businesses, the challenge is often both cultural and practical in that it requires a shift from thinking about individual customers to thinking about networks and viral effects. Operations focused on targeted messaging to specific people rely on databases with rows for individual customers rather than complex matrices of interrelations. A solution is to take metrics for those interrelations and put them into individual customer records. Leveraging them properly requires a shift in thinking from "how can I impact this customer?" to "how can I impact a network of customers by intervening with the right person in the right way?" Fortunately, the tools and techniques have expanded greatly.

Computational social science (Lazer et al. 2009) is the growing interdisciplinary field combining the computational capabilities of social network analysis with the advancing power of computers to map, model, and understand the complex interactions of humans. Network approaches allow the modeling of the way our relationships impact our actions. Early successes in network analysis have enabled the tracking and understanding of how we meet and mate (Christakis and Fowler 2009), how many "degrees of separation" there are between any two people on the planet (Watts 2003), and how people in some "positions" in a network have advantages and disadvantages over others (Monge and Contractor 2003). Related work on trust in networks suggests that the kinds of relationships that may be more likely to have influence dynamics may relate to homophily, measures of network closeness and the number of interactions people may have (Roy et al. 2017; Roy, Singhal, and Srivastava 2017). Many of these models draw upon work in epidemiology, tracking how an idea or emotion can spread through a network, analogous to how a disease may do the same. These are often state changes (Pathak, Banerjee, and Srivastava 2010), or modeled as contagion of a one-time behavior (Berry et al. 2019).

Because many of these approaches use viral analogues, fewer focus on the ongoing nature of a social system, with all of its messy comings, goings, events, and changes. More effort has come from commercial applications, with marketers looking to spot and leverage "influencers" in systems (Booth and Matic 2011; Li, Lai, and Chen 2011). "Viral marketing" (Petrescu and Korgaonkar 2011) is the phenomenon of consumers impacting each other through a range of social media and word-of-mouth, whether organic or engineered by a marketer. In marketing, a common assumption is that a person who has been viral once will be so again, launching an intense search for these key nodes in networks. There is now a robust literature on how marketing via "influencers" (Brown and

Hayes 2008) works. Hayes, King, and Ramirez (2016) detail how the advertising messages can spread and when they are rejected, emphasizing the importance of relationship strength, consistent with findings from Cho, Huh, and Faber (2014), Chiu et al. (2007), and Seo et al. (2018). Significant work by Berger and colleagues highlights the importance of communication channels (Berger and Iyengar 2013), message features (Berger and Milkman 2012), and specificity (Packard and Berger 2015) in determining the virality of messages.

Some, though not all, of these approaches focus on high-value influencers, operating on two assumptions: first, that a connection between people will lead to an action taking place, typically in the form of a purchase, and, second, that the number of followers a person has means they are influential (Olenski 2018; Wang and Jin 2010). Aside from people gaming the system (Narayan 2018), these assumptions have two limitations.

First, the approach assumes that having followers leads to an action outside of the system's sight, for example, a tweet leads to a purchase. If influencer A has 1,000 followers, they are arguably twice as influential as influencer B with 500 followers. However, we don't know whether those followers actually do something as a result of their following. The practice of hoping it works continues much as it has from the dawn of advertising, epitomized by the wry joke "Half of advertising works. We just don't know which half" (Rothenberg 1999). A robust metric of virality and influence would follow the paradigm of performance marketing, connecting the process to actual purchases. This is critical because it allows for a metric to be falsifiable (Popper 2002) and to have an accuracy measure attached to it.

Second, the approach focuses on large-scale networks of influencers and followers, that is, finding the right large-scale celebrity. However, as we demonstrate below, there is immense influence at the small-scale level, especially among close contacts. Celebrities may have sway, but they may not be as powerful as our family, friends, coworkers, and neighbors in their aggregated power (Bakshy et al. 2011).

The Social Value approach uses an epidemiological framework based in network science, plus a reliance on behavioral data to address these two shortcomings. Going forward, we highlight the distinction between speech and behavior: knowing what someone did or caused—in contrast to what their speech might have caused—allows for falsifiability and accuracy tests, and for advertisers, a return on investment (ROI). Subsequently, this means that while we acknowledge the foundational work of speech-based research on influence, the approach here focuses entirely on resulting behaviors, that is, spending, playing, viewing, etc. How and why people influence each other is not a component to the approach here, though it can be added after the fact for both insights and to add value to potential interventions. The approach outlined here treats interactions as a black box, focusing instead on the results of those interactions and tying them to verifiable actions.

Social Value

Social Value of a person is the collective behavioral impact in a network that can be attributed to the person, that is, the extra behavior of others that is most likely due to the presence of the person in the network.

The premise is simple: the presence of other people causes us to do more (or less, since impact can be negative as well) of something. Their absence leads us to do less (or more) of the same thing. In other words, some amount of our behavior is attributable to other people. The questions are then "which people?" "which behaviors" and "how much behavior?" Because we already have the total amount of behavior, answering these questions becomes an attribution challenge: what portion of an individual's behavior should be attributed to others, and vice versa? This leads into the logic for how to test and validate any solution: if we think that a person has a certain amount of influence on another's amount of behaviors, then the removal of that person should lead to an equal reduction of that behavior. This could be either positive or negative. A person who depresses others' behavior could be said to have negative Social Value and their removal would lead to an increase in others' behavior.

It follows that if we can measure the Social Value of one person in a network, we can do it for everyone. That level of analysis allows us to see the aggregate total of all personal influence on behavior across the community. In other words, it allows us to see what percentage of the community's behavior is driven purely by interpersonal influence and what percentage must be driven by something else. We call that remainder Nonsocial Value. Some communities may have more Social Value in them, perhaps when the members have stronger affective ties or when they are more interdependent. This does not mean that a community or system is better or worse; it does, however, mean that more of its behaviors are caused by social

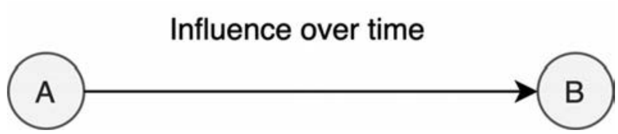

Figure 1. Influence's basic unit is a dyad, that is, two people, who interact over time.

interactions than others. Why this might be is left for future researchers presumably employing experiments or natural experiments to reverse engineer Social Value levels across a range of communities and contexts.

The purpose of the approach is to eventually allow the modeling of social influence on any observable, quantifiable behavior. "Behavior" is generic and could mean anything measurable, as the algorithm itself doesn't know what a unit of behavior means. For example, this approach can be used to measure influence on time spent, sessions, monetary spending, or anything else where the behavior is trackable and capturable with a time stamp. Behaviors such as voting or sexual encounters may be impractical or unethical to track and capture.

As noted above, this method does not explain how influence works operationally. We know mathematically how much there is and who wields it, but we don't know whether it's happening through discussion, imitation, persuasion, force, or one of many other possible mechanisms. For the present analysis, the discussion of what's inside the black box of the influence process is out of scope, but it is an obvious and intriguing area for future research.

We start with examining the interaction of two people, then between a person and many others, and then among everyone. The most basic (or atomic) unit of analysis is the dyad, that is, two connected people. Take person A and person B, who we know to be connected (see Figure 1). If we want to understand how much impact, if any, is person A is having on person B's behavior, the issue of causality and its three requirements immediately arises: correlation, time order, and ruling out other explanations (Mill 1884). Clearly causality is difficult to establish from data alone and is a widely researched area (Pearl 2000), with an explosion of new research in the recent past (Pearl and Mackenzie 2018). However, in this research our claims are not as strong as causality as per Popper's (2002) definition, where a single counterexample can refute the existence of causality. Rather, our claims are more akin to weaker constructs like probabilistic causation (Eells 1991), where counterexamples could be observed with some (low) probability, or even the psychological concept of attribution (Schwarz 2006), where a causality is perceived to exist. Of course, since we actually test with data, the attribution is converted into probabilistic causality.

Our approach to measuring probabilistic causality is based on a retrospective analysis of behavior, which is the core of the Social Value approach. First, we must have access to data on A and B; for each there has to be some observable, recordable action and an accompanying time stamp for it. These kinds of data are increasingly available.

If we have a series of observations where A and B take some action, we can impose Mill's three criteria. First, when A does something, does B also do it? This gives us correlation. Second, does A's action precede B's? This gives us time order. Finally, and most challengingly, can we rule out other explanations? Perhaps B was going to take the action anyway or was as impacted by some other cause as much as A was, but always acted later than A, for example, they both drive home but one gets off of work earlier. To rule out such explanations, we must have the equivalent of an experiment. There must be cases where we are observing B's actions in which A is present and in which A is not present.

If we have enough of these, we are effectively watching a field experiment of A and B out in the world. We call this approach inclusion/exclusion—the over-time observation of behavioral chains when the focal person is, and is not, present. The result is that we can estimate with reasonable accuracy how much more behavior B engages in because of A's presence and thus the Social Value attributable to A. Again, without engaging in any interventions or testing, we do not know how or why this occurs, but we know with some degree of certainty that it does and that it is due to A. As in any experiment, when all else is equal, the difference between test and control must be due to the stimulus. Here, A's presence is attributable as the stimulus.

As an example, let us say that A and B are residents in a neighborhood who sometimes meet at a local park. Some days A comes to the park on her own and stays for an hour. On days when A and B both come, A and B stay at the park for an hour. However, on days when A does not come, B stays for only 30 minutes. In this case, it can be reasonably argued that A's impact on B is 30 minutes of park time. This approach allows us to take B's park-going behavior and to break it down into two mutually exclusive pieces. There is the portion of behavior that A caused, that is, the extra park time of 30 minutes.

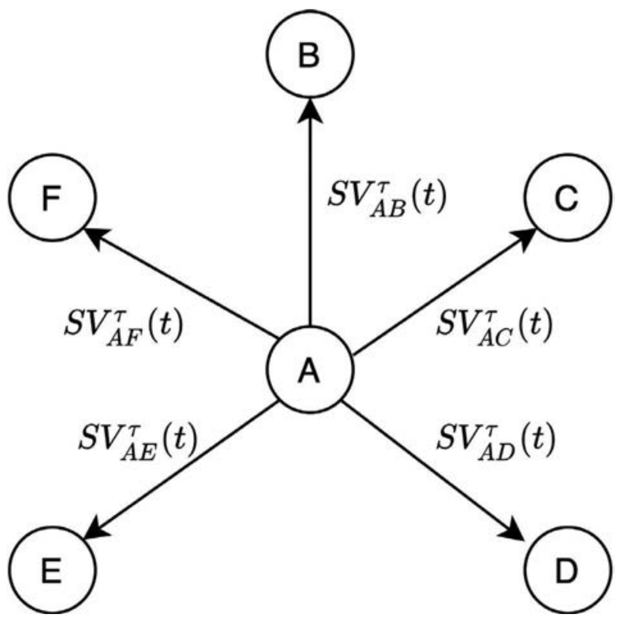

Figure 2. Larger networks show that expanding from a dyad will reveal that a person has some total amount of Social Value that can be aggregated from all of their relationships. Person A's Social Value is the sum of their influence on each of the five others here.

We call this B's "Following Value." There is also the portion of behavior that B was going to do anyway, an effective baseline of behavior that is unrelated to B's connection to A. We could think about this as the value of the park itself. We call this B's "Nonsocial Value." Nonsocial Activity is akin to the marketing concept of loyalty in that it is behavior driven by the relationship between the person/consumer and the product/place. We can look at A's behavior and apply the same labels. In this example, A has no Following Value, because she is there for an hour whether B comes or not. Therefore, her Nonsocial Activity is an hour. Additionally, A's impact on B can be measured, which we know to be 30 minutes. This is A's influence just on B, that is, her Social Value on B. It is important to note that influence going out and influence coming in always even out. There is a basic conservation theorem at work. After all, we are not creating new behavior, just attributing the behavior that is already there. Also, this is a hypothetical and clean case. There may be other factors at play, which we haven't factored in and which will lead to error in estimating the attribution. For example, maybe it rained one day, limiting A's or B's access to the park for reasons unrelated to the park or their relationship. Because there are likely other social factors at play that we don't capture, the measure is inherently conservative, either meeting or undercounting the ground truth.

The dyad is the most basic unit, but of course people often interact with more than one other person, and the concepts of Social Value and Following Value are meant to capture a person's whole set of impacts across all of their interactions. In network language, they are properties of a node, not an edge. If a person interacts with 20 others and has varying amounts of influence on each, their Social Value is the sum of all of those. Within their network, it's how influential they are overall.

This is illustrated in Figure 2. We still have person A and her friend B, but now we also have friends C through F, and so we can look at A's interactions and impacts on all six of them. If A has her influence of 30 minutes with B, no impact on C, D, and E, but 15 minutes of influence with F, then in this network, her Social Value is 45 minutes. Note that we could use the term Social Value at the dyad level, but by default when we use it we are referring to a person's aggregate influence on all others. We suspect that although the dyad is measurable and important, most users of this approach will focus on an individual's overall impact.

Further, from prior social networks research, we know that there are effects that can go out more steps than just this first one (Monge and Contractor 2003), though we expect them to diminish as they radiate out (Newman 2003), just as a stone's ripple on a pond fades as it moves away from the initial point of impact. The inclusion/exclusion logic still applies, but now is carried out to the larger network. What are the behaviors of all of the people in the network when A is around versus when she isn't? When the network is larger than the simple one in Figure 2, this total is still her Social Value. And, the same conservation theorem applies as it did with A and B. A's total Social Value must add up to all of the Following Activity she caused in the rest of the network.

Human relationships are of course far more complex than the simple examples presented here. There are more people impacting others than just A. Perhaps B has some Social Value as well. And B may also influence A. The impacts do not have to be one-sided between people. It's straightforward to imagine that people come to the park to see each other and will both come for more time than they would alone. Perhaps a great deal of the overall activity is due to people's interest in seeing others rather than the value of the park on its own. This is why we speculate that overall community levels of Social Value will be higher in systems with affective ties and

Table 1. Terms and notations.

	Terms	Notations
	Network	$<V, E>$
	Person	$\upsilon_i,\ \upsilon_j,\ etc.$
Social influence of i on j at time t		$I_{ij}(t)$
Estimate of $I_{ij}(t)$		$I\hat{}ij(t)$
Time interval		$\tau = (t,\ t + \tau\ )$
Social Value of i with j at interval τ beginning at time k: $SV_{ij}{}^{\tau}(k)$		$SV_{ij}^{\tau}(k) = \int_k^{k+\tau} \left(I_{ij}(t) * spendRate(t)\right)dt$
Estimate of $SV_{ij}{}^{\tau}(t)$		$S\hat{}V_{ij}^{\tau}(t) = \sum_{w_k \in [t,t+\tau]} \hat{I}_{ij}(w_k)\mathrm{spendRate}(w_k)dw_k$
Social Value of i at interval τ beginning at time t		$SV_i^{\tau}(t) = \sum_{j\in N(i)} SV_{ij}^{\tau}(\mathrm{t})$
Estimate of $SV_i{}^{\tau}(t)$		$S\hat{}V_i^{\tau}(t) = \sum_{j\in N(i)} S\hat{}V_{ij}^{\tau}(t)$

interdependence. We introduce these terms using some notations below, with more detail in Table 1.

- Nonsocial Value (A,c,t): It is akin to the loyalty of a customer A to an organization's product or category c at time t. It can vary with time. This is defined as the Nonsocial Value of A.
- Influence (A,B,c,t): It is the influence of customer A on customer B on a specific category c at time t. It can also change over time. A's influence on B causes change in the Nonsocial Value of B. The amount by which B's Nonsocial Value with product c at time t changes by the presence of A is called the influence of A on B about product c at time t.
- Influence $(A,B,{*},t) = \sum_c Influence(A,B,c,t)$: The Influence of A on B across all categories,
- at time t. The "*" symbol indicates that influence over all the categories has been aggregated while other factors are still fixed (A, B, and t).
- Influence $(A,{*},c,t) = \sum_{b\in N(A)} Influence(A,b,c,t)$: Influence of A on all neighbors in category c at time t. Note that $N(A)$ is the set of all neighbors of A.
- Influence $({*},B,c,t) = \sum_{a\in N(B)} Influence(a,B,c,t)$: Influence of all its neighbors on B, for category c at time t. We call this Following Value of B.
- Influence $(A,{*},{*},t) = \sum_c \sum_{b\in N(A)} Influence(A,b,c,t)$: Total influence of A on the network, at time t. This is defined as the Social Value of A.

We can roll up these component parts into other meaningful metrics. Someone's unaffected behavior (Nonsocial Value) and their impact on others (Social Value) are together their impact on the system, which we call "Network Power." Their own total amount of activity, which is what we have traditionally always observed, comprises their unaffected behavior (Nonsocial Value) and their behaviors driven by others (Following Value). We call this already known value "Personal Activity," and again note that this is an attribution exercise within already observed values. The final metric is someone's overall presence in the system and is made of their unaffected behavior (Nonsocial Value), their influence on others (Social Value), and other's influence on them (Following Value). This is their "Total Value," and it is how much activity would disappear from the system if this person were to leave.

Individuals whose Social Values are very high are called Social Whales. Their presence causes disproportionately high behavior by others. And, their absence causes behaviors of others to drop significantly. This last use of this approach comes at the aggregate level and so is no longer about A, B, or any other individual and is instead about the community overall. At this level, we examine a table of all of the individuals, as the example in Table 2 shows.

Here we can observe two new insights. First, we can see the distribution of Social Value across the entire group. Is it evenly spread out or concentrated among a few individuals? Second, we can see how much social activity there is overall. In this example there are only two people who cause others to engage in more activity (A and F). By adding all of the columns' values we see the totals for Nonsocial Activity, Following Activity, and Social Value. The latter two must be the same since the total influence coming out of people must equal to total influence affecting people. Again, this is an attribution exercise within a zero-sum space. We are not creating extra behavior in the analysis. Whichever of the social totals we use (Following Activity or Social Value), this is the measure of activity that is purely attributable to social causes. This is the answer to the question "How important are other people?" In this example, 150 of the 450 units of behavior are socially driven, so we can say that this group's behavior is one-third social and two-thirds nonsocial.

The implications of this value are twofold. First, it provides an objective measure by which we can compare two groups. If group 1 is 20% social and group 2

Table 2. An example of a small community with leaders and followers and totals for social and nonsocial activity.

Person	Nonsocial Value (what a person does entirely on their own)	Following Value (what person does because of influence from neighbors)	Social Value (what neighbors do based on the influence of this person)	Total Value
A	60	20	60	140
B	30	30	0	60
C	45	40	0	85
D	55	10	0	65
E	70	50	0	120
F	50	0	90	140
Totals	310	150	150	610

Note. The sum of Following Value and Social Value over the entire network will be identical.

is 50% social, it prompts the question of why. Second, we will naturally wonder whether we can change that percentage, which would require experiments, or quasi-experiments (Cook and Campbell 1979), likely involving different advertising messages. If we can observe a group and then effect some change in it, we can see whether there is a resulting change in the overall totals or in the individual measures. For example, if we see that overall activity increases, that tells us the intervention had an impact. Yet now we can see into the change matrix to determine how that impact took place, which gives us insight into why and how it changed. Imagine that a community has 1,000 units of behavior, and then after some change it has 1,500 units of behavior. If those 500 extra units came as a result of increases in Nonsocial Activity versus Social Value increases, we gain insight into the process. We would know whether the intervention led to increased individual activity or the mechanism led to increased social activity or some mix of both. If we have the ability to examine two similar communities, we can have higher validity if we see a change in a test community with an intervention versus a control community without it.

As noted above, the algorithm treats the social process like a black box, but we can then take these results to open the box and consider those social processes. Different interventions and messages will shift different metrics. For a practitioner with an understanding of the consumer base and the context, this opens the possibilities to both measure objective impacts as well as infer how and why they occurred. In turn, this enables learning about the types of interventions that work best in various applications.

The algorithm is designed to work with a variety of scenarios. Social influence can be observed in many varied settings where the key performance indicators (KPIs) of interest as well as the factors influencing them depend on the environment and objective of the practitioner. However, certain common elements can be identified and an approach is developed based upon these common elements. First, it is required to have a scenario consisting of multiple users interacting and participating in some activity of interest. Typically, this activity of interest will be consumer behavior such as buying items or engaging with a service. Two important components required for Social Value computation are data on user activities (timestamped so that temporal ordering can be observed) as well as user interactions (or some proxy of interactions between the users, such as location). For example, data on users buying items from a website as well as data on users inviting friends to use the website data can be used as inputs to the Social Value computation algorithm. The approach is designed to estimate what percentage of users' activities can be attributed to other users they interact with in the environment.

The core idea is to develop a machine learning model that treats some KPI, measuring the intensity of the activity of interest, as a target variable and learns a function that predicts this KPI based upon various factors describing the environment and the users' behaviors. Note that factors measuring frequency and intensity of interactions between users are also included in this model. Once such a function is obtained, the absence of a users' interactions is simulated by zeroing out the frequency and intensity of interactions and estimating the KPI based upon on just the nonsocial factors. For a given user u_a, the difference in the KPI values when social factors are present and absent is considered the aggregate social impact of all other users in the system. Generally, this can be narrowed down to the set of users that user u_a has interacted with and, based upon the intensity of interactions, the aggregate value can be further broken down into individual attributions for each user interacting with user u_a. The Appendix provides a more detailed and formal description of this approach.

Validation and Testing

The algorithm used here has been coded into a Python package that is open-sourced and is freely available for noncommercial purposes. It, and the accompanying data set used below, are downloadable at https://github.com/eunakhan/social-value.

Measuring Social Value directly is difficult as there is no observable ground truth. Instead, we measure Social Value by determining its impact on observable behaviors and whether the results are consistent with their expected measures. Because the behavior of interest is spending, the study focuses on the impact of Social Value on spending in an online game setting. These were available and complete data and are not intended to generalize to all possible settings. An advantage of gaming data is that they allow perfect and complete measurement of all transactions as well as all interactions to easily construct a social graph. However, any purchasing data set can be converted into a social graph by assuming a perfect graph in which all consumers are connected (see Limitations and Future Directions).

The approach for validation is the following: First, all active users in a given time period are considered, then the spending history of these users' is collected and is used to compute the expected spending for each user in the subsequent time period. These projected spending values are compared to the actual spending values and the error is observed. The next step is to compute Social Value impact on spending for each user and then adjust the predictions using these values. It is expected to observe a systematic decrease in the error values for a significant portion of the population. Thus, the core idea behind the validation process is to see whether accounting for the impact of social behavior on user spending will explain an appreciable portion of the error observed when predicting user spending values.

Consider A and B in Figure 2. Suppose a state-of-the-art method, based on past spending behavior and profile of B, but not on B's interaction with A, predicts the future spending amount of B, over some interval τ, which starts at time t, as $Spend_B^\tau(t)$. Besides, we also have the information that B's friend A has churned (stopped using the system) at time t, that is, was present at time $t-1$ but not at time t. The Social Value algorithm estimates that $SV_{AB}^\tau(t)$ is the spending amount of B for period τ, attributable to A's presence, and hence should contribute toward A's Social Value. Based on this intuition, when A churns, B's spending amount should decrease by $SV_{AB}{}^\tau\ (t)$. We can adjust user B's spending prediction, for period τ, as $Spend^\tau{}_B(t)$ - $SV_{AB}^\tau(t)$, and this value should be closer to the actual value than the predicted $Spend^\tau{}_B(t)$, that is, explicitly accounting for user A's departure should reduce the error in spending prediction for user B. Moreover, the phenomenon should be consistent for significant population of users, particularly if they are in the higher percentiles of social active users. Applying this adjustment should systematically reduce error and is used to validate Social Value computation. The full process is outlined below:

1. Build a model that predicts the spending amount for users. Use data up to time t and predict users' spending amount in interval $\tau = (t, t + \tau)$.
2. Based upon the network at time t, find all users (U) whose neighbors have churned (i.e., neighbors are absent from the system) in this interval τ and consider the pairwise Social Values from each of these users to their churned neighbors.
3. Predict spending amount at τ for each user $u \in U$, using the model from step 1. Denote this estimate as $Spend_u^\tau(t)$.
4. For each user $u \in U$, subtract sum of pairwise Social Values of all churned neighbors on user u, from the spending amount estimate of the previous step, to get the Social value adjusted spending amount: $AdjSpend_u^\tau(t) = Spend_u^\tau(t) - \sum_{y \in N(u)\,\&\,y\ \text{churned in}\ \tau} SV_{yu}^\tau(t)$.
5. Compare $Spend_u^\tau(t)$ and $AdjSpend_u^\tau(t)$ with the actual spending amount. It is expected that $AdjSpend_u^\tau(t)$ will be more accurate than $Spend_u^\tau(t)$.

Validation and Social Value Computation Experiments on Online Game Data

With the assistance of a game publisher, we gained access to detailed play records over time, for a large-scale online game, *World of Tanks*,[1] to validate the Social Value model. Notably, we did not rely on self-reports of behavior, entirely avoiding experimenter effects by using unobtrusively collected data (Webb et al. 1966) logged by the game's operator. All data were collected, anonymized, and used in accordance with an institutional review board–approved protocol. Additionally, the supplying company complied with General Data Protection Regulation laws for further privacy safeguards. No personally identifiable information was used, and the universal player IDs were subjected to a one-way encrypted hash so that no analysis could be connected to a player.

To add context within *Tanks*, players participate in online team versus team battles of tanks. Each user selects a tank of their choice and two teams consisting of 7 to 15 tanks on either side are pitched against each other on a virtual battlefield. Users can play solo

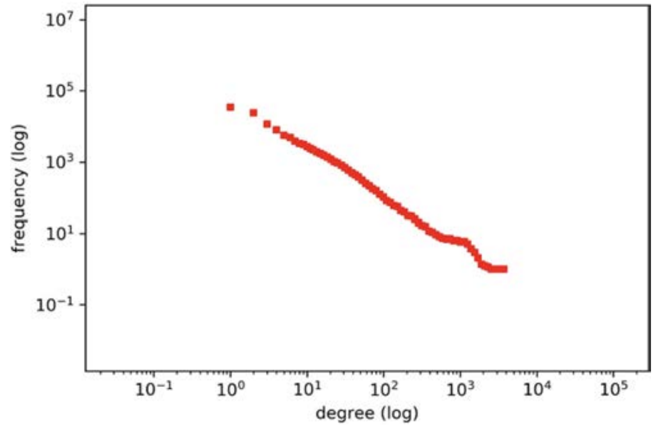

Figure 3. Degree distribution (log scale).

or in groups with their friends. In case of solo play, teammates are other similarly skilled users selected randomly. The other, more social, option is to invite friends to form a small group or even an entire team. Each time a player actively chooses to play with another, it is represented as an edge in the network. The edge weight between two users is the log of total time they played together. The log transformation was used to better deal with the exponential nature of edge weight distribution and is consistent with Zipf's Law, which is found in a range of social science data (Wikipedia 2021). Players spend real money on virtual items and experiences within the game. These range from new tanks to cosmetic alterations, more powerful ammunition, and accelerated rewards from play. The in-game currency is called gold, and the results below are given in gold units, which are generally converted at a rate of 5 gold to $1. One important question is whether spending from an online game are a proxy for other forms of spending. No one form of spending would suffice for all commercial behaviors, as they range across sectors and levels of activity. That is a large-scope issue worthy of significant future analysis. For now, the goal is validation, reporting on real people, in a real commercial system, making real purchases.

For our study, we considered a three-week period ranging from March 4, 2019, to March 25, 2019. This period was chosen to minimize the impact of New Year's as well as Summer season special events, which usually happen before and after, respectively, the month of March. The sampling window was thus on the conservative side of activity. All active users, that is, a user with at least one tank battle in the period, were included in the data set. There were 183,754 players and for all of them recency, frequency, and tenure-based metrics were compiled, along with information on how much they played, with whom, on each day. Figure 3 shows the degree distribution of the network where the nodes are the players and an edge exists between a pair of players when they had played together.

Social Value scores were computed for all users in this set. The time played together (on a log scale) by two users in the interval was used as the weight of the edge connecting them. The features used (taken as an average from the most recent three weeks) were number of sessions, number of inactive days, frequency, and number of unique neighbors in the co-play network. For each user the feature set also included weighted averages of the number of sessions, number of inactive days, and frequency of that user's neighbors. The weights for this average were taken from the edge-weights in co-play network. These edge-weights corresponded to edges measuring co-play between the user and his/her neighbors. These data were used to compute the Social Value influence between pairs of users. At the core of Social Value estimation is a machine learning model that estimates the effect of the factors described above on the variable of interest (in our case, gold spending). Details on the methodology can be found in the Appendix. A random forest regression model was used inside the Social Value algorithm and it had an R^2 value of 92.92% and an accuracy of 82.64%, where the number of trees used was 100 (Figure 4).

Social Value was computed for all active users in the data set. However, an appreciable percentage of these users were new with very low activity, that is, one to a handful of days at most. There are very few data on such users to make meaningful predictions and conclusions since they are not yet participating/engaged in the system, even though technically they are part of it. This is a common phenomenon in free-to-play games where the barrier to entry is very low but progression in the game demands significant time

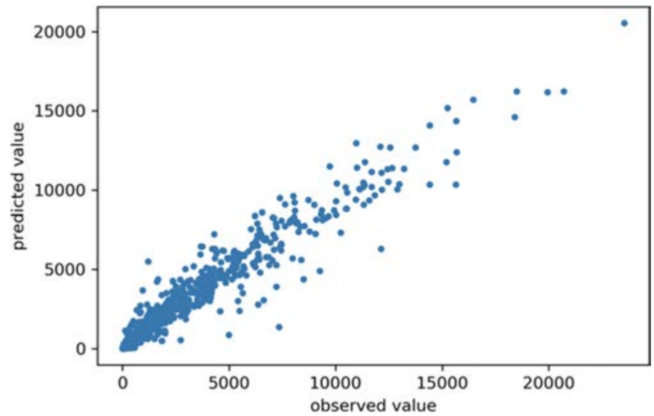

Figure 4. Predicted value vs. observed value.

and monetary investment from the player. Moreover, users with low Social Value are very likely to have a negligible impact on their neighborhood, and it is difficult to get any measurable effect of their social behaviors' impact. Almost all these users are those with zero to very few relationships with other users. Therefore, for a less noisy validation, the following users were filtered out:

i. Users who did not have sufficient history. The number of days played in a 13-week history, for each user, was compiled and the 25th, 50th, and 75th percentiles corresponded to gaming sessions with more than 2, 6, and 30 days, respectively. In order to capture the more engaged section of the population, the top quartile, that is, users who had gaming sessions of at least 30 days over the prior 13 weeks, were considered for the validation set.
ii. Next, we eliminated users who had zero to very little play time with others. Figure 5 shows how much they had played in teams versus their total playtime during the three months from January 1 to April 1, 2019. We computed the ratio of time played in teams to total time played for each player and found that 25th, 50th, and 75th percentiles for this amount were 0.018, 0.06, and 0.3. In other words, 75% of the players had this value less than 0.3. Once again, we pick the top quartile and selected users with their ratio of playtime with friends versus total playtime being greater than 0.3. After selecting the most engaged and social players, we ended up with 14,587 nodes and 30,976 edges.

We built a time series model that predicted the amount of gold spent given the history of spending data for users. Because the playing behavior in this kind of environment iterates in a weekly fashion, we used week-level granularity. Three weeks of data were used to generate predictions for users' gold spending in the subsequent fourth week. To predict the amount of gold spent for week four, we use a vector autoregression (VAR) model. VAR models generalize the univariate autoregression model to multiple parallel time series, for example, multivariate time series. They can predict the values for multiple variables at time t from their corresponding previous values, that is, t-1, t-2, etc. The order of our VAR model was 3. The relative absolute error (RAE) for the model was 69.65% when all the training data were used as testing data and 71.17% when 1,964 players were held out during training and used for testing. The difference is not large, which indicates that the model did not overfit. The higher error rates in this phase were not a matter of concern because the goal here is not to build a predictive model to efficiently predict future spending behavior but to see how the use of Social Value improves the predictions. Next, we identified the 769 users who were active in the three-week history but churned in the subsequent week, that is, they had zero sessions then. We identified their neighbors (set U where $|U| = 1{,}964$) and observed their predicted gold-spending using the VAR model.

Next, the pairwise Social Value scores were used to adjust these predictions from the VAR model. For each user u in set U, we subtracted the sum of pairwise Social Values that the churned neighbors had on this user, from the predicted amount of gold spent. The intuition behind this approach is that the amount of gold spent by u should decrease by the Social Value amount from its churned neighbors.

For users with significant Social Values, in their absence a drop in their neighbors' spending amount is expected. In order to illustrate this, we partitioned the users into 10 equal-width bins based on their churned neighbors' Social Value scores. Table 3 shows that bin 10 has 14 players whose churned neighbors had on

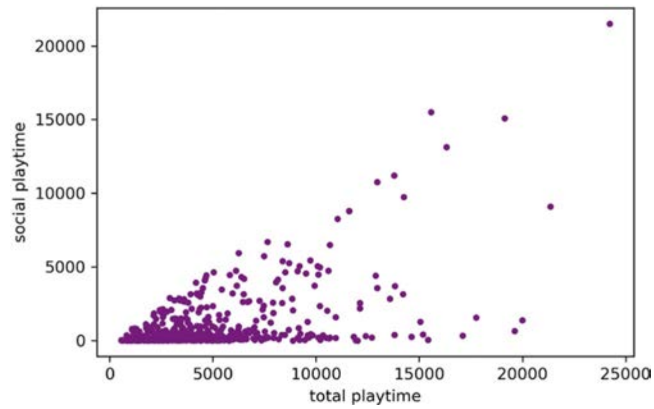

Figure 5. Social playtime against total playtime.

Table 3. Distribution of different output attributes.

Metric	Minimum	Maximum	Standard deviation	Mean	Total
Social Value	0	222.550	1.266	0.125	43195.949
Nonsocial Value	0	631.375	5.428	0.594	205229.185
Following Value	0	392.924	1.962	0.125	43195.949
Network Power	0	631.427	5.652	0.719	248425.135
Personal Activity	0	688.180	6.450	0.719	248425.135
Total Value	0	720.362	6.682	0.844	291621.084

average 2159.852 units (in gold) of Social Value on them. As we move across the bins, the Social Value drastically decreases. So, we focused on bin 10 (and bin 9) for the validation. Table 4 lists the gold spent predictions for players in bin 10, without and with adjustment using Social Value scores. The adjusted predictions have a much lower relative absolute error rate with an improvement of 29.06% on average. Bin 9 has 74 players and the Social Value adjustment lowers the error rate by 45.89%.

Apart from validation, we also present some aggregate statistics on Social Value based metrics on the entire population of active users (before filtering out users for validation). Various measures discussed earlier like Following Value, Nonsocial Value, Network Power, Personal Activity, and Total Value are shown in Table 3. This experiment was conducted for 183,754 users and it was estimated that the percentage of Social Value out of total spending (i.e., total Social Value/(total Social Value + total Nonsocial Value)) is 17.39%. This means that 17.39% of spending in this set of users can be attributed to social impact during the three-week period. The distribution of Social Value for these users is plotted on a log-log scale in Figure 6, which demonstrates a power law, which is quite consistent with Zipf's Law (Wikipedia 2021).

Figure 6 indicates that few people had a high Social Value score. These are the people who have higher influence on their neighbors and their presence or absence can cause changes in their neighbors' behavior. During validation, we focus on these users and observe changes in their neighbors' behavior, amount of gold spent being the behavior of interest in our experiment, after these influential players churn or leave the game. These neighbors correspond to users in bins 9 and 10 in Table 4.

Table 5 lists the predictions while predicting the amount of gold spent (i) without adjusting for Social Value and (ii) with adjusting for Social Value for the players in bin 10. Without being able to factor in the influencer causing spending, the predictions are systematically wrong because the influencers did not spend in the tested periods. On average, the error is 64.39% before and 35.33% after using Social Value, which is a 29.06% improvement. For bin 9 with 74 players, the improvement is 45.89%.

Discussion

The central premise above was that the presence or absence of people in a system will have measurable effects on others' behaviors. Friends cause behavior, and thus looking at all behavior becomes an attribution exercise to determine which portion is social and which portion is nonsocial. This was validated by looking at cases where the person was present and those where they were not present to determine the impact on others' behaviors. In the data here, the Social Value spending distribution followed power laws (Pareto 1896), with a minority of individuals wielding a majority of the influence in a system. There are high-level implications for advertising practice: If the total amount of influence in a system is X%, it means that some cents on the dollar are attributable to social forces among consumers, not to the product or its messaging. Practitioners can therefore view that same X% as the amount of new signal they

Table 4. Social Value in each bin.

Bins	Numplayers	Average Social Value in bin	Total Social Value in bin
10	14	2159.852	30237.928
9	74	269.354	19932.219
8	180	35.772	6439.078
7	272	5.406	1470.450
6	327	0.711	232.657
5	239	0.108	26.005
4	102	0.014	1.468
3	31	0.001	0.060
2	5	0.0003	0.002
1	935	1.283e-07	0.0001

can now see and potentially leverage. X becomes a sort of "dark matter" equivalent in advertising practice: there and exerting force, but not previously measurable. The validation of the algorithm affords practitioners with a tool to identify the amount of social influence in a system and to target those who are likely to spread it or be susceptible to it. In parallel, it gives researchers a broad tool for identifying communities with high and low levels of social activity. By determining the influence in a system, the tool enables follow-on work into interventions within that system.

Potential Applications

We propose three initial use cases where this signal might add value: in targeted messaging, in user acquisition optimization, and in inventory management. In each case, existing practices can be used to harness the new signal.

Interventions through an influential person in a network are the most direct application. By activating such a person, their network neighbors should be more likely to follow suit. The right form of activations will likely vary greatly by context. We speculate three basic forms that merit testing: a giveaway, in which the influential consumer receives a good at a reduced price or for free; a pay-forward intervention, in which the influential consumer receives an offer not for themselves, but to give to others around them; and a sharing intervention, in which an influential consumer is given an offer that only works when redeemed in tandem with another person. While the first intervention is fundamental marketing simply using the new targeting signal, the latter two are more socially oriented in nature and are more likely to activate the existing organic relationship that the algorithm has detected.

Social Value measurement allows a new means to evaluate the success or failure of such interventions to spread to others. By definition, the Social Value of a consumer is others' spending, so changes in it are a measure of that ripple effect. As an example, Kumar and Rajan (2012) suggested that "social coupons" (e.g., Groupon) may have poor ROI due to cannibalization of the existing consumer base. The Social Value approach allows a more targeted approach and evaluation not by blanketing a social group or area, but by casting a stone to hit the maximum ripple. Moreover, the approach is testable. Having offered a coupon or any intervention, the algorithm allows the business to see the impact on the person at the epicenter and then on those around them. Rather than estimating word-of-mouth, the measurement allows the business to see it in purchasing behaviors and calculate an exacting ROI versus the opportunity cost of the intervention.

User acquisition in online performance marketing is the evaluation of a funnel and its several steps to ultimately determine the ROI of marketing spend by comparing the cost of acquisition to the lifetime value (LTV) of the consumer (Hoban and Bucklin 2015). The Social Value "dark matter" amount is the optimization potential of these funnels because it is a proportional correction to LTV. It indicates that users are worth more or less than previously thought (still netting to 0). If different advertising messages or platforms are more or less likely to bring in consumers who increase others' spending, the ROI of those messages or platform expenses can be adjusted accordingly. This allows arbitrage for the users of the algorithm compared to those without it; a company knows that ad platform A is +5% more valuable than otherwise known because it generates an LTV from its delivered consumers that is 5% higher than was previously observable. Ad platform B is underdelivering in that its delivered consumers are 5% less valuable than was previously thought. In both cases, the company can bid for ad services with information asymmetry. Also, separate from the ad platform, the measure allows qualitative insights on the creative used. If, for example, an ad featuring dogs brings in X% net extra spending by delivering influential consumers compared to an add featuring cats, it is a valuable signal for future creative decisions.

Last, Social Value has an inverse application with products themselves. Although individual SKUs have no influence or influenceability, their purchasers do. Consider a pair of jeans bought by many customers but whose purchasers have no Social Value. Compare these jeans to another pair bought by fewer customers but whose purchasers have strong influence. Which should the business order more of, stock, and

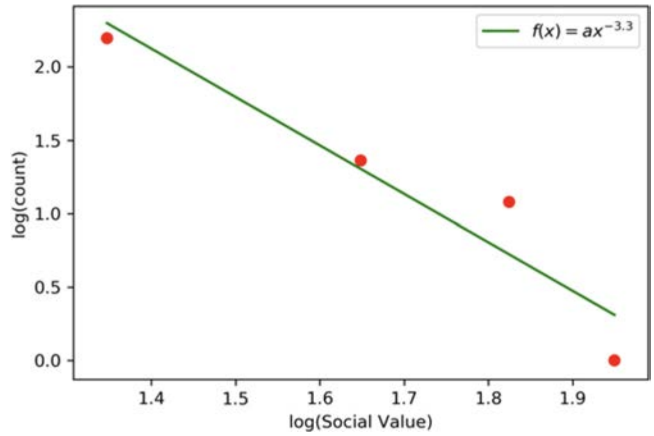

Figure 6. Social Value demonstrating power law.

Table 5. Spend predictions with and without using Social Value (SV) for the topmost bin.

	Estimated spend (gold)		
Players	Without SV	With SV	Observed spend (gold)
1	667.871	0	0
2	819.125	0	0
3	167.879	0	0
4	1115.043	299.065	6700.0
5	807.837808	0	0
6	196.432	0	1000

advertise? Without the Social Value signal, the business will erroneously choose the first pair, which is currently popular. With the signal, the business will consider the second pair, which is likely to *become* popular.

A Note on Organic Applications

Experimental designs are possible in any social system where we can collect objective, complete data on human behaviors in ethical ways. The ability to measure these values at one point in time can be extended to repeated measures, and the algorithm can be rerun for different time periods. If we observe that the percentage of social activity at time 1 is X% and then at time 2 is Y%, we can ask why. If we can effect some experimental change in a system compared to a control, we can know why.

Understanding this process could as also be used for positive social change. Research has shown that peer pressure can induce cooperation in a network (Mani, Rahwan, and Pentland 2013). This can take other forms and become more efficient by starting with the most influential people in that network. Crucially, it is unlikely to work in a typical top-down corporate sense where the hope of starting with an "influencer" who is seeded with some product or intellectual property will somehow simply "go viral." If an intervention in a network is to spread through these real affective ties, it may be most efficient if it does so because the people want to retain and strengthen their ties. Consider the pay-forward and sharing applications above, in which both activate existing ties and demand a positive social act. The most successful interventions may be those in which the tie is strengthened on purpose; this bears future testing. Regardless, successful interventions are most likely to be organic in the sense that they fit within the relationships rather than coming from an unconnected outside source like the advertiser. For example, the intervention might give people a benefit if they take an action together or feel more connected. If this bears out, it would align the incentives of marketers with civil society—the most profitable marketing uses of social influence will be those in which friendships and communities are strengthened, increasing the likelihood of true organic influence to propagate.

Limitations and Future Directions

As with any research, there are limitations as well as future directions to consider. We do not assume that the players in one online game are representative of other online games or of other consumer verticals. Moreover, we have selected only the active players, filtering out those who did not play as much or interact with others. Each game ecosystem will have different such distributions of activity to consider. In that context, the findings here are only an initial data point. We assume that the distributions of social activity and of Social Value itself will vary from context to context. In turn, this will allow us to explore why the variation occurs, as many have (Lazer et al. 2009; Muchnik et al. 2013). As with many others, we have an essential confound: if some systems have more or less skews in their power laws, is that due to the system driving behavior or to the self-selection of different kinds of people into those systems? These data were collected from field settings where consumers had opted into a range of experiences, and it is not clear whether they knew the social implications of the various settings. Without a controlled experimental design in which we force large numbers of users into these conditions, we cannot know with certainty. It is likely, though, that the amount of Social Value in a system will vary and be based on the nature of that system.

For example, staying with games for a moment, some games have mechanics that force interdependence among players and tend to generate longer-term associations. This is a direct indicator that the features and social architectures—the literal code of the system (Lessig 1999)—is related to the social experience. More social players may select more social experiences. More social architectures may cause players to have more, and more substantive, social experiences. Either, or both, could be at work. Without a controlled test, we cannot say with certainty that this is causal, but a comparison of a more socially incentivized game with one like in our initial test would suggest that there will be more Social Value. Preliminary, non-peer-reviewed commercial tests using the algorithm have been done in gaming, e-commerce, and retail environments (Williams 2018). In these tests, the algorithm did not rely on populations with known ties, as was the case for the test here with gaming data. Instead, those tests started with the assumption that all users were connected, resulting in a super-dense network graph that was computationally expensive. Iterations of that method suggested that these graphs could be pared down with heuristics, which allowed for quick and less costly run times as some cost to the overall accuracy. Those tests suggested a rough baseline rate of 30% to 40% of behavior being driven by others rather than by products and services, across a range of commercial activities and game types (Williams 2018). Significant future testing is needed to explore these preliminary findings.

The takeaway point is broader than an analysis of online gaming. If we can find the amount of Social Value in one system, we can do it in all systems. This will enable us to learn which systems lead to higher or lower levels of social activity. This has implications for advertisers seeking to use the organic connections within a group to spread messages or to incentivize behaviors. Rather than seed a message with social media influencers, advertisers can explore targeting particular users to engineer ripples on the pond. Using the tool here, as well as extending it through an open-source effort, will allow us to learn what works and what doesn't.

In addition to pondering how the phenomenon can be leveraged, we have inadvertently created a bevy of fundamental questions to address in future work. Can Social Value be thought of like Customer Lifetime Value (Kumar 2018), which has an accruing, cumulative function? It is based on a social dynamic, rather than solely on an individual's traits, but the same patterns may apply, with increasing values over time. We have merely looked at a three-week period. What happens in the next three, the next, and so on? And for those who drop out, is there a decay function? Do losses in relationships create vacuums that are in some cases replaced, and if so, how? How will the density of actions and interactions vary across systems, and will that suggest different units of time for different applications? For example, the necessary number of events needed to reach accurate values would be unrealistic in, say, car-buying data, while three weeks may be longer than needed in a denser event space such as consuming music. These basic concepts suggest that issues of density, velocity, and context are fertile ground for exploration.

Note

1. https://na.wargaming.net/en

Acknowledgment

We are indebted to Wargaming for providing access to the anonymized data and in particular to Eugene Kislyi and Jeremy Ballenger for their help and feedback.

Code and Data

The analysis code and data set are available at https://github.com/eunakhan/social-value, with a non-commercial and commercial license.

Conflict of Interest

The first author has worked as a consultant for Wargaming. The third author works there full time. No work for this research was compensated. The other authors declare that they have no competing financial interests.

ORCID

Dmitri Williams http://orcid.org/0000-0001-7995-4429
Euna Mehnaz Khan http://orcid.org/0000-0002-0030-7841
Nishith Pathak http://orcid.org/0000-0002-9873-5157
Jaideep Srivastava http://orcid.org/0000-0001-9385-7545

References

Backhouse, R. E., and P. Fontaine, eds. 2014. *A Historiography of the Modern Social Sciences*. Cambridge: Cambridge University Press.

Bakshy, E., J. Hofman, W. Mason, and D. Watts. 2011. "Everyone's an Influencer: Quantifying Influence on Twitter." Paper presented at the Proceedings of the Fourth ACM International Conference on Web Search and Data Mining, Hong Kong, China.

Berger, J., and R. Iyengar. 2013. "Communication Channels and Word of Mouth: How the Medium Shapes the Message." *Journal of Consumer Research* 40 (3):567–79. doi:10.1086/671345

Berger, J., and K. L. Milkman. 2012. "What Makes Online Content Viral?" *Journal of Marketing Research* 49 (2): 192–205. doi:10.1509/jmr.10.0353

Berry, G., C. Cameron, P. Park, and M. Macy. 2019. "The Opacity Problem in Social Contagion." *Social Networks* 56:93–101. doi:10.1016/j.socnet.2018.09.001

Booth, N., and J. A. Matic. 2011. "Mapping and Leveraging Influencers in Social Media to Shape Corporate Brand Perceptions." *Corporate Communications: An International Journal* 16 (3):184–91. doi:10.1108/13563281111156853

Brown, D., and N. Hayes. 2008. *Influencer Marketing: Who Really Influences Your Customers?* New York: Taylor & Francis.

Chiu, H. C., Y. C. Hsieh, Y. H. Kao, and M. Lee. 2007. "The Determinants of Email Receivers' Disseminating Behaviors on the Internet." *Journal of Advertising Research* 47 (4):524–34. doi:10.2501/S0021849907070547

Cho, S., J. Huh, and R. J. Faber. 2014. "The Influence of Sender Trust and Advertiser Trust on Multistage Effects of Viral Advertising." *Journal of Advertising* 43 (1): 100–14. doi:10.1080/00913367.2013.811707

Christakis, N., and J. Fowler. 2009. *Connected: The Surprising Power of Our Social Networks and How They Shape Our Lives*. New York: Little, Brown and Company.

Cialdini, R. B. 1993. *Influence: The Psychology of Persuasion*. New York: Quill.

Cook, T. D., and D. T. Campbell. 1979. *Quasi-Experimentation: Design & Analysis Issues for Field Settings*. Boston: Houghton Mifflin Company.

Curtin, M. 2018. "This 75-Year Harvard Study Found the 1 Secret to Leading a Fulfilling Life." *Inc.* https://www.inc.com/melanie-curtin/want-a-life-of-fulfillment-a-75-year-harvard-study-says-to-prioritize-this-one-t.html

Eells, E. 1991. *Probabilistic Causality*. Cambridge, UK: Cambridge University Press.

Friedkin, N., and E. Johnson. 1999. "Social Influence Networks and Opinion Change." *Advances in Group Processes* 16:1–29.

Hayes, J., K. W. King, and A. Ramirez. 2016. "Brands, Friends, & Viral Advertising: A Social Exchange Perspective on the Ad Referral Processes." *Journal of Interactive Marketing* 36: 31–45. doi:10.1016/j.intmar.2016.04.001

Hoban, P. R., and R. E. Bucklin. 2015. "Effects of Internet Display Advertising in the Purchase Funnel: Model-Based Insights from a Randomized Field Experiment." *Journal of Marketing Research* 52 (3):375–93. doi:10.1509/jmr.13.0277

Katsikeas, C. S., N. A. Morgan, L. C. Leonidou, and G. T. M. Hult. 2016. "Assessing Performance Outcomes in Marketing." *Journal of Marketing* 80 (2):1–20. doi:10.1509/jm.15.0287

Kumar, V. 2018. "A Theory of Customer Valuation: Concepts, Metrics, Strategy, and Implementation." *Journal of Marketing* 82 (1):1–19. doi:10.1509/jm.17.0208

Kumar, V., L. Aksoy, B. Donkers, R. Venkatesan, T. Wiesel, and S. Tillmanns. 2010. "Undervalued or Overvalued Customers: Capturing Total Customer Engagement Value." *Journal of Service Research* 13 (3):297–310. doi:10.1177/1094670510375602

Kumar, V., and B. Rajan. 2012. "Social Coupons as a Marketing Strategy: A Multifaceted Perspective." *Journal of the Academy of Marketing Science* 40 (1):120–36. doi:10.1007/s11747-011-0283-0

Lazer, D., A. Pentland, L. Adamic, S. Aral, A.-L. Barabasi, D. Brewer, N. Christakis, N. Contractor, J. Fowler, M. Gutmann, et al. 2009. "Computational Social Science." *Science* 323 (5915):721–3.

Lee, J., J. Lee, and L. Feick. 2006. "Incorporating Word-of-Mouth Effects in Estimating Customer Lifetime Value." *Journal of Database Marketing & Customer Strategy Management* 14 (1):29–39. doi:10.1057/palgrave.dbm.3250033

Lessig, L. 1999. *Code and Other Laws of Cyberspace*. New York: Basic Books.

Lessig, V. P., and C. W. Park. 1978. "Promotional Perspectives of Reference Group Influence: Advertising Implications." *Journal of Advertising* 7 (2):41–7. doi:10.1080/00913367.1978.10673220

Li, Y.-M., C.-Y. Lai, and C.-W. Chen. 2011. "Discovering Influencers for Marketing in the Blogosphere." *Information Sciences* 181 (23):5143–57. doi:10.1016/j.ins.2011.07.023

Mani, A., I. Rahwan, and A. Pentland. 2013. "Inducing Peer Pressure to Promote Cooperation." *Scientific Reports* 3 (1):1–9. https://www.nature.com/articles/srep01735. doi:10.1038/srep01735

Marwick, A. 2015. *Status Update: Celebrity, Publicity, and Branding in the Social Media Age*. New Haven, CT: Yale University Press.

Mill, J. S. 1884. *A System of Logic, Ratiocinative and Inductive: Being a Connected View of the Principles of Evidence and the Methods of Scientific Investigation*. New York: Harper.

Monge, P. S., and N. S. Contractor. 2003. *Theories of Communication Networks*. Oxford: Oxford University Press.

Muchnik, L., S. Pei, L. C. Parra, S. D. S. Reis, J. S. Andrade, Jr., S. Havlin, and H. A. Makse. 2013. "Origins of Power-Law Degree Distribution in the Heterogeneity of Human Activity in Social Networks." *Scientific Reports* 3: 1783. doi:10.1038/srep01783

Narayan, L. 2018. February 1. "It's Time for a Reboot of Influencer Marketing." *Adweek*. http://www.adweek.com/digital/lux-narayan-unmetric-guest-post-reboot-influencer-marketing/

Newman, M. E. J. 2003. "Ego-Centered Networks and the Ripple Effect." *Social Networks* 25 (1):83–95. doi:10.1016/S0378-8733(02)00039-4

Olenski, S. 2018, January 17. "50 Ways Brands Can Increase Their Instgram Followers." *Forbes*. https://www.forbes.

com/sites/steveolenski/2018/01/17/50-ways-brands-can-increase-their-instagram-followers/#3b9ab79c779d

Packard, G. M., and J. Berger. 2015. "Following the Blind: How Expertise and Endorsement Style Impact Word of Mouth Persuasion." Available at SSRN 2615023.

Pareto, V. 1896. *Cours D'economie Politique*. Lausanne: F. Rouge.

Pathak, N., A. Banerjee, and J. Srivastava. 2010, December 13. "A Generalized Linear Threshold Model for Multiple Cascades." Paper presented at the IEEE International Conference on Data Mining Sydney, Australia.

Pearl, J. A. 2000. *Causality: Models, Reasoning, and Inference*. Cambridge, UK: Cambridge University Press.

Pearl, J. A., and D. Mackenzie. 2018. *The Book of Why: The New Science of Cause and Effect*. New York: Basic Books.

Petrescu, M., and P. Korgaonkar. 2011. "Viral Advertising: Definitional Review and Synthesis." *Journal of Internet Commerce* 10 (3):208–26. doi:10.1080/15332861.2011.596007

Phelps, J. E., R. Lewis, L. Mobilio, D. Perry, and N. Raman. 2004. "Viral Marketing or Electronic Word-of-Mouth Advertising: Examining Consumer Responses and Motivations to Pass along Email." *Journal of Advertising Research* 44 (4):333–48. doi:10.1017/S0021849904040371

Popper, K. R. 2002. *The Logic of Scientific Discovery*. London: Routledge.

Rothenberg, R. 1999. *Ad Age Advertising Century: The Greatest Icon*. https://adage.com/article/special-report-the-advertising-century/ad-age-advertising-century-greatest-icon/140149

Roy, A., J. Huh, A. Pfeuffer, and J. Srivastava. 2017. "Development of Trust Scores in Social Media (TSM) Algorithm and Application to Advertising Practice and Research." *Journal of Advertising* 46 (2):269–82. doi:10.1080/00913367.2017.1297272

Roy, H., A. Singhal, and J. Srivastava. 2017. "Formation and Reciprocation of Dyadic Trust." *ACM Transactions on Internet Technology* 17 (2):1–24. doi:10.1145/2996184

Schwarz, N. 2006. "Attitude Research: Between Ockham's Razor and the Fundamental Attribution Error." *Journal of Consumer Research* 33 (1):19–21. doi:10.1086/504124

Seo, Y., X. Li, Y. K. Choi, and S. Yoon. 2018. "Narrative Transportation and Paratextual Features of Social Media in Viral Advertising." *Journal of Advertising* 47 (1):83–95. doi:10.1080/00913367.2017.1405752

Trivedi, J., and R. Sama. 2020. "The Effect of Influencer Marketing on Consumers' Brand Admiration and Online Purchase Intentions: An Emerging Market Perspective." *Journal of Internet Commerce* 19 (1):103–24. doi:10.1080/15332861.2019.1700741

Vrontis, D., A. Makrides, M. Christofi, and A. Thrassou. 2021. "Social Media Influencer Marketing: A Systematic Review, Integrative Framework and Future Research Agenda." *International Journal of Consumer Studies* 45 (4): 1–28. 10.1111/ijcs.12647.

Wang, R., and Y. Jin. 2010. "An Empirical Study on the Relationship between the Followers' Number and Influence of Microblogging." Paper presented at the E-Business and E-Government (ICEE), 2010 International Conference on E-Business and E-Government, Guangzhou, China.

Watts, D. 2003. *Six Degrees: The Science of a Connected Age*. New York: W.W. Norton & Company, Inc.

Webb, E., D. Campbell, R. Schwartz, and L. Sechrest. 1966. *Unobtrusive Measures: Non-Reactive Research in the Social Sciences*. Chicago: Rand McNally and Company.

Wikipedia. 2021. "Zipf's Law." Accessed July 11, 2021. https://en.wikipedia.org/wiki/Zipf%27s_law.

Williams, D. 2018. "For Better or Worse: Game Structure and Mechanics Driving Social Interactions and Isolation." In *Video Game Influences on Aggression, Cognition, and Attention*, edited by C. Ferguson, 173–83. New York: Springer.

Appendix

Computing Social Value

Consider a directed network of users, G(V,E), where each user $u_i \in V$ and each edge $e(u_i,u_j) \in E$ is a positive real number representing the intensity of some given relation from user u_i to u_j. Examples of such relationships can be the number of likes on social media posts, messages sent, invitations sent, and time spent playing together. For each user we can have a behavior of interest; for example, in the field of advertising, the behavior of interest is user response. In most businesses, it is often a user-related KPI measuring some form of engagement. For example, the behavior of interest can be a simple binary indicator of whether a given user clicked on an ad. It can also be a KPI measure such as time to return, session length, or number of sessions. Assume for a given user u_i, value $y_i \in Y$, is the measure of users' behavior of interest. The neighbors of any given user u_i are defined as the set of users $n_j \in N(u_i)$, such that $e(n_j,u_i) \in E$. The objective is to quantify and compute the impact of every given user u_i's neighbor $n_j \in N(u_i)$ on u_i's behavior of interest y_i. Specifically, this is an attribution problem as a portion of observed y_i is "credited" to u_i's neighbors. For example, if a user typically plays one hour of an online game by himself but his playtime goes up to one and a half hours if he has friends joining in his party, then one can argue that the half an hour increase in playtime can be attributed to these friends.

It should be recognized that the above situation arises in a variety of scenarios and it would be difficult/impossible to develop a specific technique that will work for all cases. Therefore, a more general approach/framework for Social Value computation is presented. The proposed approach can be applied in all cases as long as sufficient data on user behavior and the environment are available. Extending the formal scenario presented earlier, assume the presence of data $d_i \in D$ which comprehensively describe user u_i's behavior including his/her interactions with the environment as well as other users. The following methodology can be followed to compute Social Value:

1. Develop machine learning model M (random forest used in the experiment), which learns the function, $f(d_i) = y_i,\ \forall\ u_i \in V$.
2. Modify data d_i to $d_i^{''}$, by omitting instances of user u_i's interactions with other users from d_i.

3. Use model M to estimate $y_i^{''} = f(d_i^{''})$, i.e., the expected behavior of u_i if neighbors are excluded and social behavior is essentially "zeroed out."
4. Difference $\Delta_i = (y_i - y_i^{''})$ can be attributed to the combined social effect of u_i's neighborhood. This quantity Δ_i is defined as the Following Value of ui, Inf(ui), and is a measure of how much ui's neighbors influence his/her behavior.
5. Proportional to relationship intensity, $e(n_j,u_i) \in E$, the Following Value for u_i is distributed among each of the neighbors and pairwise Following Value from n_j to u_i, PairInf(n_j,u_i), is defined as, $Inf(u_i) * e(n_j, u_i) / \sum_{n \in N(u_i)} e(n, u_i)$. The idea is that neighbors with stronger relationships get more credit. Thus, a portion of u_i's behavior y_i, is attributed to each of his neighbors.
6. The pairwise influence a given user u_i exerts over all his neighbors can be aggregated to represent the total influence u_i has over his neighborhood. This quantity is defined as Social Value: $SV(u_i) = \sum_{n \in N(u_i)} PairInf(u_i, n)$.

The approach presented above requires the practitioner to use machine learning methods for computing model M which, as long as sufficient data are available, can be developed and fine-tuned to perform optimally based upon the underlying domain and other contextual information. The core idea is to develop a model of user behavior and then use this model M to simulate the absence of any given user's neighbors and estimating that user's expected behavior. The difference in the expected behaviors, when a given user's neighbors are present and absent, is the value attributed to the neighbors' influence on that user. This is then distributed among the neighbors and aggregated over the "influencees" to compute the respective Social Values for the influencers (see Figure 2).

As shown in Figure 2, for a pair of nodes (*i,j*), SV_{ij} represents the amount of activity done by *j* which is attributable to the influence of *i*. Thus, this activity by *j* should be considered as part of the Social Value of *i*. Thus, in Figure 2, SVA = SVAB + SVAC + SVAD + SVAE + SVAF

As discussed above, computing Social Value (SV) is an attribution problem. Consider SV_{AB} in Figure 2. The actual activity is done by *B*, whether of money, time, or any other behavioral metric of interest. However, it is attributed to *A*, who is the cause for this activity by *B*.