

Online Research @ Cardiff

This is an Open Access document downloaded from ORCA, Cardiff University's institutional repository: <https://orca.cardiff.ac.uk/id/eprint/150113/>

This is the author's version of a work that was submitted to / accepted for publication.

Citation for final published version:

Innes, Martin ORCID: <https://orcid.org/0000-0002-8950-8147> and Dawson, Andrew 2022. Erving Goffman on misinformation and information control: the conduct of contemporary Russian information operations. *Symbolic Interaction* 45 (4) , pp. 517-540. 10.1002/symb.603 file

Publishers page: <https://doi.org/10.1002/symb.603>
< <https://doi.org/10.1002/symb.603> >

Please note:

Changes made as a result of publishing processes such as copy-editing, formatting and page numbers may not be reflected in this version. For the definitive version of this publication, please refer to the published source. You are advised to consult the publisher's version if you wish to cite this paper.

This version is being made available in accordance with publisher policies.

See

<http://orca.cf.ac.uk/policies.html> for usage policies. Copyright and moral rights for publications made available in ORCA are retained by the copyright holders.



Erving Goffman on Misinformation and Information Control: The Conduct of Contemporary Russian Information Operations

Martin Innes 

Cardiff University Security, Crime and Intelligence
Innovation Institute, Cardiff, UK

Andrew Dawson

Cardiff University Security, Crime and Intelligence
Innovation Institute, Cardiff, UK

Misinformation is a social problem increasingly, and routinely, commanding significant political and public attention. Less well known is that Erving Goffman was writing about misinformation as early as 1953 in his PhD thesis. Subsequent to which, he wrote repeatedly about the social organization and conduct of “information control” across several of his most influential publications. This article distils his ideas about these concepts to explore how they illuminate the contemporary phenomena of misinformation. To do this, empirical data are introduced from a large-scale research program exploring the causes, communication, and consequences of digital information operations and campaigns, with a particular focus upon the Internet Research Agency and Instagram. The analysis attends in particular to sequences of “revealing moves” and “concealing moves” performed by the social media account operators. The dialogue between the data and Goffman’s concepts outlines the precepts of a sociologically inflected, interactionist position on misinformation.

Keywords: Goffman, misinformation, information control, social media

Direct all correspondence to Martin Innes, Cardiff University Security, Crime and Intelligence Innovation Institute, Cardiff, Wales, UK; e-mail: innesm@cardiff.ac.uk

Symbolic Interaction, (2022), p. n/a, ISSN: 0195-6086 print/1533-8665 online.

© 2022 The Authors. *Symbolic Interaction* published by Wiley Periodicals LLC on behalf of Society for the Study of Symbolic Interaction (SSSI). DOI: 10.1002/SYMB.603

This is an open access article under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs License, which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

INTRODUCTION

The discovery that the Saint Petersburg-based Internet Research Agency ran an information operation designed to try and influence and interfere in the conduct of the 2016 U.S. Presidential election, together with a series of related subsequent revelations, has attracted significant political and public attention (Benkler et al. 2018; Francois and Lin 2021). One consequence has been a rapid growth in research illuminating how distorting and deceptive digital messaging influences and shapes a variety of issues, platforms, and situations. Indeed, how to understand and control misinforming communication has developed into one of the preeminent social and political problems of the contemporary moment.

Framed by the global health pandemic and prevalent misinformation about the causes and consequences of Covid-19 (Donovan 2020; Molter and DiResta 2020) interacting with an ambient condition of political polarization and a changing media ecosystem (Margetts et al. 2016), it seems as if almost every high-profile, politically-meaningful event now functions as a magnet for disinforming, distorting, and deceptive digital communication. Often leveraged to induce doubt, dissonance, and discord (Innes 2020), misinformation has become firmly established as a political and social challenge in its own right, but also salient because of its role in shaping political decision-making and public perceptions across multiple policy domains, including the climate emergency, public health, and democratic events.

Blended into this trajectory of development is a concern that domestic political campaign strategies progressively adopt and adapt tactics and techniques that were originally the preserve of foreign state information operations. Notable in this regard is Benkler et al.'s (2018) empirical analysis, mapping an extensive far-right media ecosystem operating to author and amplify misinforming material “at scale” around the 2016 U.S. Presidential election. Related to which, Innes et al. (2021) subsequently describe a “trickle down” process of “normalization and domestication,” whereby methods for constructing and communicating such content have become increasingly expected and accepted components of contemporary political campaigns. Some empirical support for such inferences can be gleaned from Facebook’s own “threat” analysis data, where they recently described how:

... influence operations are increasingly common tools for non-state and domestic actors. Over the past several years, we have seen new actors emerge — including commercial entities and political interest groups — running both foreign and domestic IO campaigns. (Facebook 2021:15)

Broadly speaking, longer-standing social and behavioral science work on rumor, propaganda, and conspiracy theories (Allport and Postman 1947; Douglas et al. 2017; Fine and Ellis 2010; Shibutani 1966) has shaped and guided the academic literature on digital misinformation. That said, there is a recognition that the socio-technical affordances of social media, embedded within a restructuring of the broader contemporary media ecosystem (Couldry and Hepp 2017; Gonzalez-Bailon 2017), are endowing digital misinformation a uniquely de-stabilizing and disruptive influence

upon the interactional and institutional ordering of reality (Howard 2020). In engaging with this problem, the literature grounds accounts of digital misinformation in different background disciplines, meaning they are influenced by particular conceptual predicates and methodological preferences (Salganik 2018). Cutting across these inflections however, analyses have tended to adopt one of three “master” frames.

The “political economy” frame attends to how wider structural forces and pressures, especially in the media ecosystem, have configured an information environment especially conducive to the transmission and reception of misinforming content (Woolley and Howard 2019). Underpinning this approach is a strong presence from political science, and the use of quantitative methods such as social network analysis, to configure a macro- to meso-level accounting of the organization of misinforming messaging. The leading exposition being Benkler et al.’s (2018) analysis, rehearsed above, but versions can also be found in the collection of essays convened by Bennett and Livingston (2020). An alternative iteration of this position possesses a more historical base. For example, Rid (2020) demonstrates how contemporary information operations update and reinvent established “active measures” logics and practices.

The second principal frame orientates to more “pragmatic” questions of how (mis)information operations and campaigns are organized and conducted. Its signature feature is illuminating the tactics and techniques via which misinformation is constructed and communicated. This includes how particular technological affordances and functions enable misleading material to be transmitted and have impact (Linville and Warren 2020; Margetts et al. 2016). This is inverted by a second analytic strand providing high resolution micro-level analyses of the ways that those formulating misinforming messages are able to “hack” aspects of human cognition and emotion. For example, Innes (2020) delineates eight main “techniques of disinformation,” including forms of “spoofing” and “social proofing.”

A more cultural set of influences shape the dispositions of the epistemic frame. Studies adopting this posture often mobilize concepts such as “post-truth” and “post-fact” (Fuller 2018; Malcolm 2021). Their main theme is how a profound and widespread sense of doubt and uncertainty, induced in part by the mechanics and dynamics of social media technologies, has a de-stabilizing and disruptive impact upon the institutional and interactional ordering of society, in terms of what and how we “know.” Especially rich and sophisticated illustrations can be found in the work of Pomerantsev (2019), Lynch (2016), and Kakutani (2018).

Given how this literature has developed at pace, there are increasing calls for determining what its theoretical precepts and anchors might be. To date, work in this area has been strongly shaped by political science and its Russian roots. The point of this article is to set out an alternative theoretical framing, grounded in the sociology of Erving Goffman; the rationale being that Goffman wrote on repeated occasions about the concept and conduct of “information control.” Moreover, in his thesis published in 1953, he deployed the term misinformation several times. The latter is not something that appears to have been acknowledged to date in the extensive

secondary literature that has grown up around his work; nor has a connection been made with the contemporary interest in this topic, as a facet of digital life.

The argument we pursue herein is that Goffman's work in this area, although not fully developed, nevertheless offers a series of invitations to develop a sophisticated and insightful understanding of how and why misinformation arises, and the influence it can exert on interactions and institutions. Goffman certainly accents and draws out qualities and features that usefully augment and supplement the more established and orthodox accounts of such issues grounded in political science. In the next section, we review Goffman's key ideas about the conduct of information control and misinformation. This, in turn, sets up a contrapuntal dialogue with some well-known contributions to the misinformation literature. Next, we briefly introduce some empirical data to evidence the value of the perspectives that can be derived from Goffman's approach. The concluding section provides a more theoretical review of the insights we distilled.

In preparing for what follows, several "terms of art" need to be defined. First, in the contemporary literature, it has become standard practice to differentiate between concepts of misinformation and disinformation, on the basis of intent (Bennett and Livingston 2020). Disinformation describes deliberately misleading messaging, whilst its conceptual cousin "misinformation" is reserved for unintentionally misleading communications. This relatively recent convention was not employed by Goffman. However, as elaborated in subsequent sections, he was clearly using misinformation to cover both forms. Reflecting this, and for the sake of conceptual clarity and consistency with Goffman's work, herein misinformation is used in this more expansive way. Such an approach has a secondary advantage, inasmuch as the empirical dynamics of the contemporary information environment are rendering the boundaries between misinformation and disinformation increasingly blurry. Thus, it can be a distraction to have to try and divine the presence of intentionality; something often unknown. Consequently, this article refers to misinformation to cover all forms of false and misleading communication and information manipulation.

The literature also makes multiple references to the concept of "information operation." This is a planned, organized, and coordinated effort to communicate false information, often taking place over a period of time and involving multiple actions. An allied notion is "active measures." Associated particularly with Soviet-era Russia, an active measure comprises a set of influencing techniques intended to induce an adversary to cause damage to their own interests. Historically, the authoring and amplification of misinformation and disinformation was an established and regularly-used component within the Russian active measures playbook (Rid 2020).

GOFFMAN ON "INFORMATION CONTROL" AND MISINFORMATION

Erving Goffman's public reputation as a significant social analyst is closely bound up with his formulation of dramaturgy. However, over the course of his career he sought

to work out a number of alternative perspectives and framings (McLuhan 2020). Attending particularly to the rituals, rhythms, and routines of co-present interactions, his recurring project was to detail, in high resolution, how our social and self-identities are situationally contingent. As we navigate different contexts and conditions, we engage in acts of social communication that simultaneously reveal and conceal fragments of our social selves. Towards the close of his life, he introduced the notion of the “interaction order” to articulate how, underpinning the vagaries and variety of appearances that can be observed, there are core patterns of interactional behavior.

That he was searching for these patterns is important for understanding his sociology. It has been suggested, in otherwise appreciative commentaries upon his work, that he had a tendency to use multiple concepts to describe similar processes and topics across different pieces of writing (Manning 2016). According to Williams (1988), Goffman was neither constant nor consistent in developing and applying concepts, preferring instead to approach matters from a variety of angles and with a range of linguistic and conceptual tools determined by the intellectual preoccupation driving his focal enquiry at a particular moment.

This tendency notwithstanding, and in what is a hitherto neglected facet of his writing, the concept of “information control” was a central concern for him. As Jaworski (2021) contends, viewed through the prism of the interaction order, “information control” is intrinsic to the stream of deceptions, counter-deceptions, and the risks of being discredited that individuals and institutions continually seek to navigate and negotiate as they engage with one another. Indeed, it features in all four of the major publications Goffman issued in the early part of his career.

Goffman’s most detailed account of the workings of information control are to be found in his book *Stigma*, where he devotes a full section to elucidating the key “techniques of information control.” His analytic focus here is on how people with *discreditable* self-identities restrict awareness of the stigmatizing information they possess to stop it inducing the onset of *discredited* social identities. Goffman delineates a series of techniques for managing such tensions, the most important being what he dubs “covering” moves. Especially insightful are his comments where he draws out how the development and implementation of particular “covers” engages an element of social collusion. In certain situations there are, he contends, a wider “team” who, through their covertly performed actions, are engaged in co-producing the cover with the principal actor, which enables the cover to function practically. To different degrees then, these collaborators are aware of the nature and status of the potentially discrediting information, but they elect to participate in its concealment.

Prior to the publication of *Stigma* however, and especially intriguingly, he highlighted similar processes of exerting control over information in his ethnography of the concept of “the total institution.” In *Asylums*, Goffman describes information control as an integral component for the functioning of the procedures and processes of formal socially-controlling institutions, as experienced by those subject to them (Goffman 1961:159 fn40). In this case, he was concerned with how bureaucratic

case records containing intimate details of the patient/inmate's behavior, beliefs, and emotions, are constructed and readily available to controllers, even when the subject of those records would prefer they were concealed. According to his analysis, one element of the skilled "art" of practising effective control is keeping these details concealed, and using the threat of their being revealed sparingly, only when particular leverage over the subject is required.

By 1969, in his extended essay "Expression Games," although Goffman did not explicitly use the concept of information control, he was still pivoting around these same issues, identifying the central theme of the piece as, "The individual's capacity to acquire, reveal and conceal information" (Goffman 1969:4).

He positions this interest in a process of interaction as follows:

Just as it can be assumed that it is in the interests of the observer to acquire information from a subject, so it is in the interests of the subject to appreciate that this is occurring and to control and manage the information the observer obtains. (Goffman 1969:10)

As he elaborates, this can be variously achieved through the performance of "control moves" (p. 12) that engage "camouflage" through "cover operations" (p. 13); "feinting" where one starts a course of action to misdirect an enemy's attention, before shifting trajectory (p. 15); and "feigning" that strategically misrepresents belief, attitudes and/or preferences (p. 16). All of these techniques involve an element of obfuscation, either through processes of concealment, or by the creation of "noise" that renders a signal of interest difficult to discern.

In developing these themes, for illustrative purposes, he attends in particular to the work of intelligence and espionage (Jaworski 2021). It is pertinent that this analysis appears within the collection of essays entitled *Strategic Interaction*, where Goffman's cross-cutting interest was in detailing how the conduct of interaction can be goal-oriented—undertaken with a deliberate and deeper purpose, and not just a series of reactions to local stimuli.

From even this relatively truncated analysis, we can infer that the subject of information control was an important and recurring interest for Goffman. This conclusion is further reinforced when we acknowledge that many of these themes were foreshadowed in *The Presentation of Self in Everyday Life* (1959). Especially in the chapter analyzing "discrepant roles," Goffman (1959:87) ties the problem of information control to the workings of what he labels "destructive information." The latter are information that groups and teams elect to restrict to their members, on the grounds that if they were to become open knowledge they would tarnish the reputations and activities of those involved. He identifies three principal forms of destructive information: "dark secrets" are things known to a team and incompatible with the impression an audience has of them; "strategic secrets" involve intentions and purposes that are kept hidden to inhibit audiences from adapting to them; and "inside secrets" are knowledge that helps demarcate a team or group member from an outsider.

Given the latter book was a reworking of his PhD thesis, it is perhaps unsurprising that processes of information control feature strongly there too. Intriguingly however, information control is more prominent and receives a more fulsome treatment in his PhD than was provided in the published book. As such, it provides more substantive clues about how Goffman understood information control and the influences upon his thinking in this regard. Of particular note, is how, in initially working out this idea, he quotes passages respectively from Simmel, Whyte, and Park (Goffman 1953:81–82) in a rapid sequence. This is important positioning work, inasmuch as he clearly perceives the conduct of information control as fitting within the traditions of Chicago School Sociology.

Within his PhD he distinguishes between acts of “expression” and “impression” in the process of communication (Goffman 1953:73). As documented in the preceding passages, the former concept was one he returned to at several points during his career and relates to information “sent” or conveyed by the communicator. Significantly, for our interests herein, he expends considerable effort to differentiate between information that is consciously communicated, and the array of signals that a messenger unwittingly “gives off.” Both are vital to the resulting “impression,” or how an audience makes sense of what is sent and received. There are then clear analogies here between Goffman’s thinking and the role of the notions of “transmission” and “reception” in information theory. But, Goffman endows them with a slightly more interpretative twist, in terms of how meaning is ascribed to the materials being communicated; the meanings that such information takes on shape how and why people react to them in particular ways. Critically then, in Goffman’s formulation, the “receiver” of transmitted misinformation is not to be understood as a passive recipient, for they actively participate in constructing the meanings and reactions any message will elicit. Moreover, through their direct and indirect interactions with those engaged in acts of transmission, they influence the content of what is “transmitted.”

From the perspective of this article, of especial consequence is that in his PhD, Goffman makes several explicit references to the role of “misinformation.” He introduces it as part of a linked triptych of concepts, the other two being “inadequate information” and “unserious information,” that collectively account for how individuals “exert control over information about self that others are able to acquire” (Goffman 1953:74).

From this standpoint, he teases out two modalities of misinformation. Deceit involves “linguistic signs” or the content of what is said, and occurs relatively rarely under normal circumstances. The second, “feigning” is associated with “expressive signs” of emotion or conduct and is a more common focus for misinformation. As he elaborates in a footnote:

In the case of feigning or dissimulation, the sender appreciates that his expressions are “false” and “misinformative;” they are employed, in fact, precisely in order to throw the observer “off the scent.” (Goffman 1953:75)

Here then, Goffman directly addresses the issue of intent that was rehearsed in the opening discussion of definitions of misinformation. But significantly, he goes on to assert that there are circumstances when the actor can be persuaded by, and come to believe in, the fabrication they are performing. Thus, misinformation not only influences the external audience, but also possesses recursive properties, whereby the content of what is expressed “folds back” onto the communicator’s own situational understanding. Indeed, it is this sense of nuance and subtlety that may afford Goffman’s theoretical contributions particular contemporary value and resonance. He is not suggesting that an interaction is defined by the presence or otherwise of misinforming communication, but rather that distortions and deceptions, of varying intensities, are folded into the conduct of the negotiated order of many such encounters, co-existing with other more or less accurate and truthful material. As he elaborates:

... the recipient may find that part of the message is “clear,” that is, it can be taken at face value, and that another part of the message is “coded,” that is, it is a distortion of some kind and must be decoded before providing truthful information. (Goffman 1953:80)

What is innovative about Goffman’s account of misinformation and information control is how he is not casting it as anomalous or pathological, but rather *intrinsic* to the conduct of co-present interactions and the maintenance of normal appearances:

The transmission of misinformation and inadequate information appears to be a very general practice ... In Dixon, the practice of conveying misinformation or inadequate information seems well developed. (Goffman 1953:76–7)

He goes on to describe, in almost forensic detail, how such processes are used to regulate and suppress emotions and behaviors that arise out of the frustrations of life. In an important passage he concludes that there are conventions and moral norms to the deployment of misinforming messaging. These involve a, “... screen of distortions, evasions, omissions etc. ... to be analyzed and translated so as to reveal the real information that is hidden by it” (Goffman 1953:80).

As discussed in the next section, the use of omissions and distortions as well as outright fakery, are crucial to understanding the social organization of the work of Russian misinformation operatives.

APPLYING GOFFMAN’S IDEAS TO DIGITAL MISINFORMATION

To understand how and why misinforming tactics and techniques are deployed increasingly routinely within political campaigns and across other significant events, it is helpful to return to “first principles” about how such messaging is designed and delivered. In this respect, there is much to be learned from analyzing the logics and practices of those engaged in crafting such forms of communication. Especially insightful in this regard is Ladislav Bittman’s (1985) memoir. Bittman

served, between 1964 and 1966, as deputy chief of the Czechoslovak intelligence service's disinformation department. His book can be usefully supplemented and augmented by the additional historical materials reported by Grant (2020) and Rid (2020). Then, as a second step, these "principles" can be further elaborated by bringing them alongside the literature on the construction and communication of digital misinformation, briefly reviewed under the rubric of the three frames in the opening passages of this article.

Bittman (1985) provides a fascinating account of several historic active measures campaigns, and the design processes and thinking that informed them. For instance, amongst other things, he talks about the importance of "surfacing and orchestrating." The former referring to how disinforming content is presented so that it possesses a patina of plausibility for the intended target, but also deniability should the deception be discovered. Similarly, Grant (2020) describes the craft skills of manufacturing disinformation as knowing how to mask a falsehood in amongst truths, and how influence is achieved through selective omission, as well as forgery.

There is a clear touchpoint here with Goffman's information control, and the idea that this involves both revealing and concealing certain material. Across the series of empirically informed studies reviewed in the previous section, we saw how he repeatedly documented the ways some "discrediting" information is edited and restricted, where other signifiers are purposively projected into the public realm. Indeed, a strength of his analysis is how it draws attention to the blurry lines between communicating false and true information. For his framing recognizes how, as interactions unfold, some of what is expressed is more and less accurate, rather than contending it is either all fake, or all correct.

What Bittman calls "orchestrating," today tends to be labeled "co-ordination." It captures how persuasion is more likely if content is transmitted via multiple communication channels and social media accounts. These are techniques that clearly resonate strongly with how contemporary information operations and campaigns are conducted. The connection that can be read across to Goffman, is the role of "team" dynamics in dramaturgical performances, and his recognition that there is frequently a collective dimension to misleading expressions and impressions. For example, his analyses of "covering" and "passing" go to significant lengths to articulate how this necessarily engages the active participation of more people than just the stigmatized individual (Goffman 1963).

At the core of Bittman's account are three design principles that shape the overall effectiveness of a (dis)information operation and/or campaign. The first of these is summarized as follows:

The overall purpose is not only to deceive but to cause damage to the target. The victim of disinformation must be led to inflict harm upon himself, directly or indirectly ... (action or inaction). (Bittman 1985:56)

This clarifies that messaging content is not necessarily harmful on its own account, but works to exacerbate and amplify extant fissures and fractures in the target. In

seeking this effect, the author of the disinforming message has to think about the processes of reception and how these can be influenced and manipulated. It is exemplified by the ways that the Internet Research Agency's influence operation, around the 2016 Presidential election, focused so much upon racial tensions and identity politics (Benkler et al. 2018; Jamieson 2018). This is an important point in terms of understanding the sophistication and "art" of crafting disinforming communication, that often gets overlooked. The intent is not simply to use brute force by transmitting malignant information at the target, but rather to introduce it in ways that it corrodes and infects from within; degrading trust and inducing discord far more subtly.

A second related principle that can be distilled from Bittman's (1985:49) writings is how, to be persuasive, "every disinformation message must at least partially correspond to reality or generally accepted views."

This accounts for how and why information operations and campaigns so often gravitate around ambiguous and contested "wedge issues." The false material introduced into the public conversation is regularly targeted towards those points where value judgments intersect with uncertain, or unstable, knowledge. Evidence from social psychology demonstrates people are more likely to believe rumors and conspiracies, if they resonate in some way with their underlying social values and perceptual frames (Sharot 2018). In her historical account of Soviet disinformation methodologies, Grant (2020) develops this insight, identifying two key pathways: (1) persuading people to believe messages that are not actually true; (2) encouraging them to think things are false, when they are in fact correct.

The resonance here with Goffman's work is the extent to which different situations induce particular conventions and conditions that shape what is perceived by others as plausible, credible, and acceptable. Critically, what we can take from his work is how the influence of a particular message is structured by more than its content. It also depends upon the qualities of the messenger, the prevalent context, and the underlying receptivity of the audience to believe it, in terms of appealing to their established beliefs, interests or concerns.

Bittman's third key principle has been rather neglected by researchers of contemporary foreign state information operations. He writes:

they [Soviets] sincerely believe in the cumulative effect of active measures. Even though a single operation may not visibly change public perceptions of an issue, a series of related operations will eventually bring the desired changes. (1985:86)

Thinking in such terms, is not how actors engaged in analyzing and exposing digital information operations have typically approached such matters. Instead, most empirically-guided reports have been bounded either by the ability to link different sets of accounts together, or because they have been authored by a particular state, at a particular time. Far less effort has attended to the cumulative and aggregated societal consequences of a steady stream of misinformation. Where such issues have been picked up, for example by more epistemically-framed enquiries into the conditions of "post-truth" social orders, the resulting renditions tend to sweeping

generalization and abstraction, agglomerating a host of issues, as opposed to separating out the information effects of false and misleading communicative actions performed by defined groups of actors.

This is not unproblematic given how Facebook “de-platformed,” or took action against, 1.3 billion fake accounts in the last quarter of 2020 (Facebook 2021). This involved removing over 100 networks for coordinated inauthentic behavior engaged in manipulating public opinion (Rosen 2021). Relatedly, of the roughly 2.8 billion monthly active users in the last quarter of 2020, Facebook estimates 5% were fake accounts (Tankovska 2021). Of course, such data encompass “policing” multiple online harms, and not just misinformation control. But even as proxy indicators, they do suggest a significant problem. Within this tumult of activity, perhaps because of the sheer brazenness of their activity, the Russian state and the St Petersburg based Internet Research Agency in particular, have acquired a particular status as the totemic misinformation actor.

The Internet Research Agency

For understandable reasons, most analyses of the Internet Research Agency (IRA) have centered their activities in and around the 2016 U.S. Presidential election (Benkler et al. 2018; Jamieson 2018). The notoriety of what they did during this period notwithstanding, it appears that the agency’s assets were first mobilized “at scale” around 2012. We obtained detailed empirical insights into these and later activities by analyzing the datasets released by the major social media platforms following the U.S. Senate Intelligence Committee enquiry into the events of 2016. In the passages below, this includes material from the so called “Five Thirty-Eight” dataset, a corpus of 2,973,371 tweets from 2848 Twitter handles that were attributed to the Internet Research Agency by Twitter in their submissions to the U.S. Congress enquiry in November 2017 and June 2018. In terms of our processing and interpretation of these data, the focus was upon analyzing “content” and “behaviors.” The former layer attending to the messages sent and the account personas projected. The behavioral analysis used more technical signals derived from the account metadata, to explore indicators of inauthenticity, and patterns of co-ordination and orchestration (see Figure 2). The insights derived from these methods show that in 2012 the IRA accounts had a domestic propaganda function designed to help shore up the popularity of President Putin, given concerns about the “Snow Revolution”¹ that occurred between December 2011 and July 2013. In the middle of 2013, with the scale and intensity of the domestic protests waning, the IRA started to create large numbers of English language Twitter accounts — over 400 in August alone. This is suggestive of a strategic shift in interest in deploying the Agency’s assets towards foreign affairs.

However, at the end of 2013, protests broke out in Ukraine. Events escalated quickly with: (1) the appearance of the so-called “little green men” (soldiers devoid of any insignia) in Crimea; (2) its annexation in an illegal referendum; (3) the

shooting down of flight MH17; (4) and the war in the Donbas, all occurring within 6 months. Mining datasets released by Twitter and attributed to IRA operators, suggests a rapid pivot in activity during this period, with over 800 new Russian language social media accounts set up inside 6 months. However, as the geopolitical tensions subsided and attention shifted to other matters, the IRA returned to developing its international focus and reach, investing in creating many of the accounts that would target the United States in 2016.

Several “whistle-blower” accounts from workers at the IRA during this period,² paint a picture of the organizational routines of an agency providing an out-sourced capacity and capability for the Kremlin to try and influence public and political conversations around a range of issues and events of interest (Linville and Warren 2020). With respect to the latter, broad tasks were issued and distributed amongst the multiple specialist departments comprising the agency that employed over 1000 people (DiResta et al. 2018). Some departments had a specific geographic focus, where others were responsible for meme production, or posting comments in response to particular media stories. From what we know, departments and individual workers were given clear performance targets. Most of the workforce was relatively young, attracted by a decent salary rather than a particular ideological interest in the work, and there was a reasonable turnover in staff (Dawson and Innes 2019; Pomerantsev 2019).

In terms of developing an avowedly Goffman-inflected understanding of how this division of labor and these organizational routines were integrated into the production of political misinformation around the 2016 election, it is useful to engage his “conceal-reveal” dyad; pivotal to his formulation of the conduct of information control. As discussed in previous sections, Goffman’s point was that performing information control always involves a double movement, between those aspects of identity, behavior, and emotion to be suppressed, and those that should be projected and promoted. This is clearly evident in the behaviors of some of the accounts the IRA ran on Twitter. Although a comparatively large number of Twitter accounts have been attributed to the IRA — around 2800 — it appears that relatively few of them achieved significant audience reach and thus “influence-ability.” One who did was @SouthLoneStar (Figure 1):

Texas Lone Star’s biography and “presentation of digital self” clearly exemplifies the conceal-reveal connection that Goffman accents. The real identity of the operator is masked by the development of a cover identity, as a right-wing, Trump-supporting, male, from the Southern United States. In representing this persona, the operator has clearly sought to harness what Goffman (1961) refers to as “identity kits,” or clothing and other artifacts that function as expressions of social belonging. In this case, there is the use of the Stetson hat in the thumbnail image and the inclusion of the flag in the banner-head, as well as some clearly provocative hashtags. Juxtaposed in this way, it is a combination of signifiers clearly intended to project a particular social identity to help the account attract target audiences.



FIGURE 1. Screenshot of @SouthLoneStar's Twitter Persona

In terms of puncturing such covers, one of the key behaviors that social media companies attend to are signals of co-ordination. Transposed into Goffman's terms, this would be at least partly analogous to the accent he places upon teams and the collaborative nature of performances engaged in information control activities. In an effort to apply this to understanding how coordination and collaboration were a feature of the work of the IRA, the researchers developed a novel data analysis method dubbed "synchronicity analysis." This works by looking for temporal patterns in the posting activity of different accounts, to see if their messages occur at around the same time. Figure 2 provides a visualization of the results from implementing this approach using a sample of 10 IRA linked accounts.

Time is plotted on the X-axis and the dots depict where posting activity by each of the 10 accounts occurred. The larger dots depict an increased volume of messaging activity. What can be immediately observed is how there are common "pulses" of activity across multiple accounts. Using this method, it was possible to show that approximately 70% of the accounts in the dataset of 2800 released by Twitter were coordinating with at least one other IRA account. Pragmatically, this probably reflects the bureaucratic organization of the work, with individual operators responsible for running multiple accounts. But having "teams" of accounts is also part of amplifying and boosting the visibility of the content being promoted. It enables an operator to flood an information zone with a particular message.

It is worth noting that it has been repeatedly documented that, although the majority of the Twitter accounts attributed to the IRA had right-wing social identities, not all of them did. There were accounts "spoofing" more liberal-left and black rights activist personas (Benkler et al. 2018; Dawson and Innes 2019). Potentially, this might suggest the overarching objective was to induce discord, disputes, and political polarization, as opposed to advancing a specific ideological position.

Far more is known about the IRA's activities on Twitter and Facebook than Instagram, reflecting some of the methodological predilections and skillsets of the




FIGURE 2. Synchronicity Analysis of 10 IRA Twitter Accounts

research community. However, as we elaborate below, it is instructive to focus upon their use of Instagram on the grounds that we can observe clear traces of some of the techniques of deception and misdirection that Goffman illuminated. Instagram further warrants attention because of its popularity with younger social media users. At around the time when the IRA’s American activities were discovered, Instagram was being used by 71% of American 18 to 24 year olds compared to Twitter’s 45%,³ with younger social media users especially appreciative of its visual and image based communication format.⁴

A report commissioned for the U.S. Senate Select Committee on Intelligence inquiry estimates there were 187 million engagements with IRA content on Instagram, distributed across some 20 million users. Moreover, as their activities on Twitter and Facebook were exposed, the IRA operators shifted more attention to Instagram. They produced a prolific volume of memes and image based content, re-purposing other popular memes, as well as appending large numbers of hashtags to build visibility. Twelve IRA attributed Instagram accounts achieved over 100,000 followers, and 40% of their accounts had more than 10,000 (DiResta et al. 2018:26). On Instagram, the fake social identities were targeting the black community far more, sending messages seeking to suppress voting behavior, and to encourage secessionist and insurrectionist ideas. Based upon dates of the first posts of their accounts, it appears that the IRA’s campaign started on Twitter, then moved to Instagram and only subsequently developed a significant Facebook presence. Especially intriguing is how some of the Instagram accounts were used to sell merchandise, such as T-shirts featuring politicized slogans (DiResta et al. 2018:31). The analysis shows that the IRA accounts were embedded within a wider network of

similarly-presenting authentic accounts, sharing overlapping content, and the study authors concluded that the social media companies probably had not identified and attributed all of the IRA creations operating across their “surfaces” (DiResta et al. 2018:45–46).

On November 5th 2018, a day before the U.S. midterm elections, Facebook released a statement announcing they had removed 30 Facebook accounts and 85 Instagram accounts (later revised to 99 accounts) for engaging in “coordinated inauthentic behavior.” This followed a tip from U.S. law enforcement agencies suggesting that these accounts were connected to foreign governments. According to Facebook, around 1,250,000 people followed at least one of the deleted Instagram accounts, giving an average of 12,000 followers per account. Because Facebook removed most of these accounts before the account names were made public, it was not possible to archive the data. However, thanks to the large numbers of 3rd party websites offering “Insta metrics,” researchers were able to archive an average of the last 12 posts made by each of the accounts, and to profile their behaviors and messages. Consistent with what was described by DiResta et al. (2018), this led to the identification of an additional 19 accounts, sharing similar and overlapping content, and displaying common patterns of behavior. Similar to the approach applied to the Twitter materials, analysis of the Instagram data pivoted around a blend of the content and behavioral layers of the accounts and their messages.

In order to pull through more of these resonances, it is insightful to develop a qualitative case study of one Instagram account displaying similar traits to those associated with the larger corpus of known IRA accounts. *Kasarov_eli* first came to attention on the 18th February, 2019, because of an anomalous pattern of behavior. Scrolling through its content history showed that most of the time it was posting memes and images of Japanese anime cartoon type figures, but these were occasionally interspersed with highly-politicized content. For example, the account reposted statements from the official Instagram account for the Presidency of the Syrian Arab Republic and Ramzan Kadyrov (the Head of the Chechen Republic and currently sanctioned under the Magnitsky Act for involvement in repression, torture and murder) – not something typical for your average anime fan! It also posted explicit disinformation messages attacking the work of the White Helmets in Syria, a key focus of the Kremlin around this time (Starbird et al. 2018). There were also periodic posts involving repeated tagging of pro-Russian words, such as #VladimirPutin. These were sometimes buried within lists of as many as 30 hashtags, likely to “game” the platform algorithms in such a way as to maximize the reach of the account’s messages.

This pattern of behavior is coherent with some of Goffman’s concepts mentioned earlier. Specifically, it is an example of “feinting and feigning” as a way of camouflaging the account’s real purposes, both from its audience of followers, but also the platform regulators. This is an important point in that the operators of the accounts are seeking to construct and maintain a cover or camouflage that works for two different audiences. First, they need to deceive other social media users who constitute

their potential followers. But also, they need to disguise their activities such that they avoid the detection algorithms used by the platform owners to scan for signals of inauthentic and coordinated behavior, and increasingly harmful content, that contravenes platform community standards and use policies.

An especially intriguing instance of such “covering moves,” was a sequence where the *Kasarov_eli* account posted a series of memes featuring Natalia Poklonskaya. Poklonskaya was the Senior Prosecutor of the Prosecutor General of Ukraine under Victor Yanukovich, and became famous when she was appointed the Prosecutor of (the autonomous/occupied) Crimea. A video of a press conference she gave to mainstream media became an online viral hit because of her attractiveness, with jokes circulating like “Steals Crimea, and your heart.”

Reflecting the account’s cover as an anime fan account, *Kasarov_eli* picked up on the wider interest in Poklonskaya amongst segments of the audience on Instagram and posted several drawings of Natalia in an anime style, as illustrated in Figure 3.⁵

This sequence of messages was designed to engage the audience and was then followed up with blatant propaganda, as captured in the fourth image above. It is also illustrative of Goffman’s point about the role of the active, participatory audience. In this instance, Poklonskaya’s fans had reinterpreted her original media press conference, using it as source material for their artistic acts of creativity, attributing to her an “influencer” status. This manifested in the anime drawings that the likely Russian operator of the *Kasarov_eli* account sought to amplify and make use of as part of their own messaging behaviors. Thus, the focus on Poklonskaya was notable because it represented a moment where the account was able to blur and blend its overt adopted persona as an anime fan account, with its covert political purpose – propagating messages aligned with Kremlin interests. In so doing, it adopted a “soft power” style of operation, presenting an attractive and positive image of Russia designed to appeal to some in its core audience, rather than just attacking its perceived adversaries.⁶

Interestingly, *Kasarov_eli* was not the only account utilizing this kind of cultural awareness. Looking across other Instagram accounts that have been formally attributed to the IRA, several “spoofed” personas suggest that they were celebrity



FIGURE 3. Pictures Posted by the *Kasarov_eli* Account Based on Google Cache

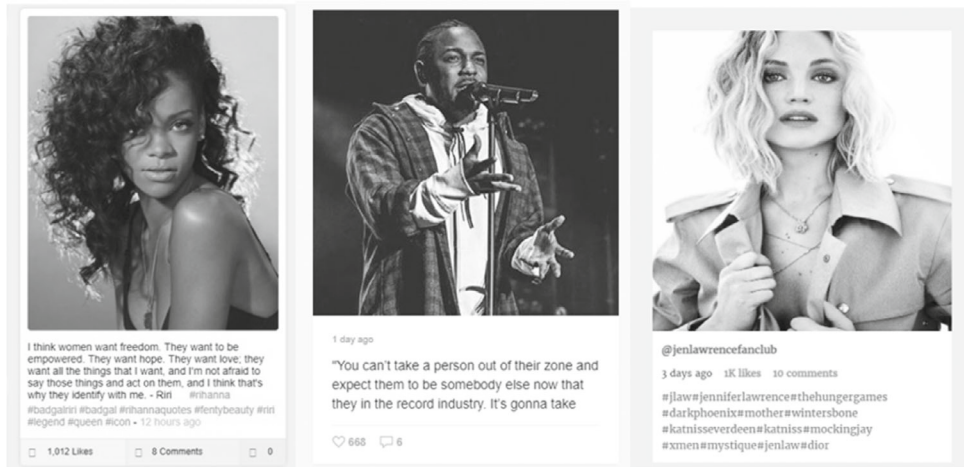


FIGURE 4. Examples of Celebrities to Attract and Influence Social Media Users. Sources: https://www.instagram.com/Riri_one_love (80,000 followers); https://www.instagram.com/Kendrick_dna (86,000 followers); <https://www.instagram.com/Jenlawrencefanclub> (34,300 followers)

fan accounts. This was their misdirection or “feint” in Goffman’s words, in that by feigning interest in a particular celebrity they could build an audience for their political messaging that would be slipped into the media stream at certain intervals. Some illustrative examples of IRA linked accounts engaging in this pattern of behavior are reproduced in Figure 4.

These celebrity-associated accounts were especially striking because of the relatively high numbers of followers they attracted. Again, this conveys a degree of sophistication possessed by some operators behind these accounts. They were sufficiently aware of internet culture and the conventions and affordances of Instagram specifically, that they had worked out a method to harness the “celebritification” of culture to enhance the reach of their activities.

Detailed analysis of the content over time, posted by the IRA attributed accounts and by *Kasarov_eli*, revealed a roughly 80:20 ratio in terms of their pro-Russian political content. That is, around 80% of *Kasarov_eli*’s messages were anime and/or largely unremarkable, apparently designed to preserve the cover and camouflage its real purposes. 20% of its posts focused upon topics and propaganda themes of interest to the Russian state. This represents a significant finding for understanding the social organization of (mis)information operations and campaigns, inasmuch as the accounts were not simply pumping out misleading and false political information; this type of content was blended into other activity that was more consistent for the “normal appearances” of the account in question.⁷

Further investigation of the *Kasarov_eli* Instagram persona unveiled an additional intriguing feature of how it had been operated. Tracking back across the account’s

history, using third-party Instagram metrics websites, revealed that this was only the latest in a series of identities, all inflected with Russian geopolitical themes. Previously it had also operated under the aliases of: *kasarov_red*; *spooky_kasarov_eli*; *god_hates_kebab* (kebab being internet slang for Muslims in some communities), *united_russia_republics*, and *stop_russophobes*. That this was possible reflects how Instagram made it comparatively easy to change account identities on the platform, but did not make such changes apparent to followers. Looking across these names, one interpretation is that they are indicators of a process of innovation and discovery, as the account operators were finding out that developing and adopting a cover to camouflage their political activity was more effective than having this overtly signaled in the account identities. The limitations with the publicly available data provided by Facebook means that this cannot be investigated further and so must remain a somewhat speculative inference. But it is a feature warranting further examination in future.

MISINFORMATION AND THE UNIT OF ANALYSIS PROBLEM

An interesting feature of Goffman's treatment of information control is how he introduces it across a series of differently-focused empirical studies. As we elaborate in the sections above, each time, the shifts in the framing of the study enables him to leverage new insights into how the control over the information that is simultaneously concealed and revealed is enacted, subject to a range of situational contingencies. Similar considerations pertain to the contemporary study of digital misinformation and specifically its "unit of analysis" problem.

Fundamentally, this problem relates to how researchers are analyzing different issues and episodes from a variety of vantage points, and the conceptual and research design choices they make alter the insights and findings generated. But, they are generally not attending to how these differing situations alter what is revealed and concealed in terms of the perspectives they are developing on the construction and consequences of misinformation. Put another way, in talking and writing about misinformation there is a tendency to assume that findings from one study can be easily extrapolated to a different situation and setting, even when this involves different actors and issues. However, a rather different perspective is gleaned if one starts with the view that misinformation is a general purpose "catch-all" concept, that captures a range of problematic and troublesome behaviors and content.

For example, most available empirical data that can be confidently attributed to particular misinformation actors comes from research conducted by the social media platforms' own internal investigative teams. This gives a good understanding of activities on that platform, but is typically less insightful about any cross-platform dimensions. Contrastingly, studies conducted by civil society organizations are often focused upon particular democratic events, such as elections, themes such as vaccine hesitancy, or climate emergency denial, shaped by their funders' interests.

A rather different set of influences have impacted academic research in this domain. Foremost amongst which is the pragmatic matter of the availability of datasets. Consequently, a lot of this work has focused upon Twitter and Facebook, neglecting what is happening on other platforms. From Goffman's multiple studies of information control, we can distil how the selection of different situations and analytic lenses affords alternative perspectives on an issue, illuminating particular features, whilst obscuring others. Applied to the study of misinformation this is not ineluctably problematic. It becomes so because of how the choices and selections made with regards to differing units of analysis are down-played and glossed over. In many cases, it is misleading to study a particular type of misinformation situation and then generalize the findings to all other types. Indeed, the irony of this tendency is that it yields misinformation about misinformation.

It is for this reason that we need to pursue the quest for theory of the type that Goffman was outlining. His particular conceptualizations of misinformation and information control foreground key patterns in the organization and conduct of distorting and deceptive social communication. Transposing them from their original forms describing features of co-present "analogue" encounters, to the digital realm of physically and (sometimes) temporally distanced interaction, necessarily requires incorporating a modest degree of interpretative adaptation (something in keeping with Goffman's own conceptual/methodological practices). This notwithstanding, although formulated for a very different age, Goffman's concepts continue to both apprehend and articulate a number of core features of how and why misleading messages are transmitted and received.

CONCLUSION

Information control, understood as a blend of tactical and strategic efforts to both reveal and conceal information about the self, was a recurring interest for Erving Goffman, as he sought to illuminate the workings of co-present interaction. It was a concept he deployed across several of his key writings, especially in the first half of his career. Misinformation is a notion he only appears to have used in his PhD study, subsequently replacing it with other terms of art, such as "destructive information," albeit describing analogous processes. Nevertheless, it is striking that this reference was being made in 1953, thus constituting a significant early statement in the intellectual genealogy of what has developed into a major topic of interest across the social, political and human sciences.

Framed in this way, the insights and analysis reported herein, are of considerable interest to the secondary literature that has grown up around the scholarship of Erving Goffman. For whilst there has been a growing appreciation for how Goffman's work can help understand the social dynamics of mass communication (Pinch 2010; Winkin and Leeds-Hurwitz 2013) and the social construction of knowledge (Raab 2019), the present analysis suggests it has especial relevance to

understanding a stream of micro-deceptions and mis-directions in the contemporary media ecosystem.

Crucial in this regard are his ideas on “feigning” and “feinting.” For, as we observed in relation to the activities of the Internet Research Agency, these describe perfectly how they built their digital disguises to exploit targeted audience interests. Their social media accounts were not just continually and exclusively transmitting misinformation, their operation was more refined than that. This included deploying a series of core “revealing moves” and “concealing moves.” Misinformation can be crafted either through acts of omission or addition. That is, by editing or concealing specific details, the meaning of a narrative can be turned. Likewise, insertion of new information into a message as it is repeated and distributed across multiple social media surfaces, is another way authors of misinformation perform their work. Importantly, this is often more subtle and nuanced in terms of its potential to influence, than when messages are outright fabrications. This resonates especially with Bittman’s (1985) assertion that a misinforming message tends to be more persuasive if it contains a “grain of truth,” or semblance of reality.

The empirical analysis further highlighted the role played by “teams” of fake accounts synchronizing and coordinating their activities, as a mechanism for boosting the visibility of particular messages. This does not always work, but when it does, it enables an actor to “flood the zone” with a misleading message dominating the information space. In this regard there is frequently a division of labor between accounts authoring misinformation messages, and those specializing in their amplification. The latter often including bot networks, which have become the focus of much political and public discussion around the misinformation problem. Positioned in the kind of framework we are developing here however, it is clear they comprise only one component in a more complex misinformation ecosystem.

This “social” component of communication highlights a further dimension that can be distilled from Goffman’s work. To fully understand what is being communicated, we need to attend to the roles of both misinformation “transmitters” and “receivers.” Unlike with some other theoretical positions, a Goffman inspired framing casts the latter as active participants in how and why misinforming and disinforming content either does, or does not, become influential in shaping the ordering of reality. A comprehensive theory of misinformation and its effects has to integrate the transmission and reception of such content, and attend to how they interact, as they mutually and recursively anticipate and react to the other in configuring their communicative actions, if we are to fully understand the meanings that particular distortions and deceptions are ascribed. Indeed, this “collapsing” of the content “producer” and “consumer” roles is arguably the key innovation of social media technologies, and renders these participatory interactions between senders and receivers of particular analytic interest.

Taking Goffman’s work as a base suggests the potential for developing a theoretical posture that prioritizes *how* misinformation is authored, amplified, interpreted, and understood. The significance of this is that, to date, the more avowedly

theoretical work that has been conducted has drawn its conceptual base from the predicates of political science. That is, being concerned principally with explicating how issues of political economy structure and condition the information environment within which misinformation is circulated. A Goffman-inspired framework affords an alternative perspective that attends more closely, and in higher resolution detail, to the pragmatics of constructing and communicating misinforming messages.

Transposed into the digital realm, Goffman's scholarship, affords the following insights:

- Social media and other media platforms function akin to a situation, setting out a series of opportunities and conventions for organizing the unfolding of communicative actions and interactions between users. Different platforms are built around different socio-technical affordances, and consequently possess different rituals and routines of communication.
- An author of a misinforming message often has to do more than fabricate false content; they have to build the credibility and plausibility of the message and the messenger account. For example, influential spreaders of Covid-19 misinformation may play on their professional medical qualifications, or highlight how the material concerned comes from a "scientific source."
- Given Goffman's interest in social identity, it is telling how IRA trolls constructed their digital social identities using fairly stereotypical identity signifiers. We can infer that these covers were designed to possess a semblance to the digital identities of ordinary social media users, on the grounds that such personas were found to be more persuasive than when political themes were made too overt.
- Drawing back from the individual accounts to consider how the IRA operation functioned as a team performance, we can observe how the campaign was not pursuing a single narrative or political line. Rather, particular IRA managed accounts adopted conflicting positions from one another around contentious political issues, reflecting how their strategic aim was to induce doubt, discord and distrust.

There are limits to what can plausibly and credibly be said about the social organization and conduct of misinformation communication based upon a case study of a single entity like the Internet Research Agency. Not least because of how the burgeoning literature on misinformation is continually being updated with analyses documenting an increasing range of state and non-state actors participating in misinformation campaigns. This notwithstanding, what a Goffman-inflected approach offers is a set of multi-dimensional, contextually sensitive, conceptual instruments, especially suited to understanding the practical tactics and techniques of information control engaged in the manufacture of misinformation.

Inspired by Goffman, analyses should attend to the unfolding point and counterpoint, between what is revealed and what concealed, to understand how misinformation authors and amplifiers engage in fabricating their communications. Promoting and projecting particular forms of information to misdirect audience attention in one

direction, whilst simultaneously using this to conceal other more discreditable material, is integral to the tradecraft of the practical “dark arts” of designing and delivering effective misinformation campaigns.

ACKNOWLEDGMENTS

The authors would like to thank the members of the OSCAR team at Cardiff, whose wider work has informed and provided a background to the arguments put forward in this article. Elements of the empirical research reported herein draw upon work investigating the causes and consequences of disinformation, originally funded by UK Government. Thanks also to the editor and reviewers of the journal who made thoughtful and helpful suggestions about how to improve the article.

NOTES

1. https://en.wikipedia.org/wiki/2011%E2%80%932013_Russian_protests.
2. See for example Adrian Chen (2015) “The Agency” <https://www.nytimes.com/2015/06/07/magazine/the-agency.html> (accessed 22/01/22).
3. <http://www.pewinternet.org/2018/03/01/social-media-use-in-2018/>.
4. <https://www.recode.net/2018/10/9/17938356/facebook-instagram-future-revenue-growth-kevin-systrom>.
5. Anime is a particular style of Japanese animated cartoon drawing that has spawned a very large and devoted online subculture.
6. Joseph Nye defines “soft power” as a mode of influence that operates through cultural capital and attraction.
7. “Normal appearances” was the title of another of Goffman’s essays.

REFERENCES

- Allport, Gordon and Leo Postman. 1947. *The Psychology of Rumor*. New York: Henry Holt & Co.
- Benkler, Yochai, Robert Faris, and Hal Roberts. 2018. *Network Propaganda: Manipulation, Disinformation, and Radicalization in American Politics*. New York: Oxford University Press.
- Bennett, W. L. and Steven Livingston. 2020. *The Disinformation Age*. Cambridge: Cambridge University Press.
- Bittman, Ladislav. 1985. *The KGB and Soviet Disinformation: An Insider's View*. Washington, DC: Pergamon-Brassey's.
- Couldry, Nick and Andreas Hepp. 2017. *The Mediated Construction of Reality*. Basingstoke: Palgrave Macmillan.
- Dawson, Andrew and Martin Innes. 2019. “How Russia’s Internet Research Agency Built its Disinformation Campaign.” *The Political Quarterly* 90(2):245–256. <https://doi.org/10.1111/1467-923X.12690>.
- DiResta, Renee, Shaffer, Kris, Ruppel, Becky, Sullivan, David, Matney, Robert, Fox, Ryan, Albright, Jonathan, and Ben Johnson. 2018. *The Tactics and Tropes of the Internet Research Agency*. New Knowledge. Accessed May 06, 2022 (<https://digitalcommons.unl.edu/senatedocs/2/>).

- Donovan, Joan. 2020. "Concrete Recommendations for Cutting through Misinformation during the COVID-19 Pandemic." *American Journal of Public Health* 110.s3(October):S286–S287. <https://doi.org/10.2105/AJPH.2020.305922>.
- Douglas, Karen M., Robbie Sutton, and Aleksandra Cichocka. 2017. "The Psychology of Conspiracy Theories." *Current Directions in Psychological Science* 26(6):538–542.
- Fine, Gary and Bill Ellis. 2010. *The Global Grapevine: Why Rumors of Terrorism, Immigration and Trade Matter*. New York: Oxford University Press.
- Francois, Camille and Herb Lin. 2021. "The Strategic Surprise of Russian Information Operations on Social Media in 2016 in the United States: Mapping a Blind Spot." *Journal of Cyber Policy* 6(1):9–30. <https://doi.org/10.1080/23738871.2021.1950196>.
- Fuller, Steve. 2018. *Post-Truth: Knowledge as a Power Game*. London: Anthem Press.
- Goffman, Erving. 1953. *Communication Conduct in an Island Community*. Chicago: University of Chicago.
- Goffman, Erving. 1959. *The Presentation of Self in Everyday Life*. London: Penguin.
- Goffman, Erving. 1961. *Asylums: Essays on the Social Situation of Mental Patients and Other Inmates*. New York: Anchor Books.
- Goffman, Erving. 1963. *Stigma: Notes on the Management of Spoiled Identity*. London: Penguin.
- Goffman, Erving. 1969. *Strategic Interaction*. Oxford: Basil Blackwell.
- Gonzalez-Bailon, Sandra. 2017. *Decoding the Social World*. Cambridge, MA: MIT Press.
- Grant, Natalie. 2020. *Disinformation: Soviet Political Warfare 1917–1991*. Washington D.C.: Leopold Press.
- Howard, Philip. 2020. *Lie Machines: How to Save Democracy From Troll Armies, Deceitful Robots, Junk News Operatives and Political Operatives*. New Haven: Yale University Press.
- Innes, Martin. 2020. "Techniques of Disinformation: Constructing and Communicating 'Soft Facts' after Terrorism." *British Journal of Sociology* 71(2):284–299.
- Innes, Martin, Helen Innes, Colin Roberts, Darren Harmston, and Daniel Grinnell. 2021. "The Normalization and Domestication of Disinformation: Far-Right and Russian State Information, Influence and Interference Operations in Europe." *Journal of Cyber Policy* 6(1):31–49. <https://doi.org/10.1080/23738871.2021.1937252>.
- Jamieson, Kathleen. 2018. *Cyberwar: How Russian Hackers and Trolls Helped Elect a President*. New York: Oxford University Press.
- Jaworski, Gary. 2021. "Goffman's Interest in Spies and Espionage: The University of Chicago Context." *Symbolic Interaction* 44(2):392–411. <https://doi.org/10.1002/symb.498>.
- Kakutani, Michiko. 2018. *The Death of Truth*. London: Harper Collins.
- Linville, Darren and Patrick Warren. 2020. "Troll Factories: Manufacturing Specialized Disinformation on Twitter." *Political Communication* 37(4):447–467.
- Lynch, Michael. 2016. *The Internet of us: Knowing More and Understanding Less in the Age of Big Data*. New York: W. Norton.
- Malcolm, Dominic. 2021. "Post-Truth Society: An Elisian Sociological Analysis of Knowledge in the 21st Century." *Sociology* 55(6):1063–1079. <https://doi.org/10.1177/0038038521994039>.
- Manning, Philip. 2016. "Goffman and Empirical Research." *Symbolic Interaction* 39(1):143–152. <https://doi.org/10.1002/symb.220>.
- Margetts, Helen, Peter John, Scott Hale, and Taha Yassera. 2016. *Political Turbulence: How Social Media Shape Collective Action*. Princeton: Princeton University Press.
- McLuhan, Arthur. 2020. "Adopting a Cloak of Incompetence: Impression Management Techniques for Feigning Lesser Selves." *Sociological Theory* 38(2):122–141.
- Molter, Vanessa and Renee DiResta. 2020. "Pandemics and Propaganda: How Chinese State Media Creates and Propagates CCP Coronavirus Narratives." *Harvard Misinformation Review* https://misinforeview.hks.harvard.edu/wp-content/uploads/2020/06/Ipeditis_FORMATTED_PandemicsandPropaganda_HKSReview.pdf.

- Pinch, Trevor. 2010. "The Invisible Technologies of Goffman's Sociology: From the Merry-Go-Round to the Internet." *Technology and Culture* 51(2):409–424.
- Pomerantsev, Peter. 2019. *This Is NOT Propaganda: Adventures in the War against Reality*. London: PublicAffairs.
- Raab, Jurgen. 2019. *Erving Goffman: From the Perspective of the New Sociology of Knowledge*. London: Routledge.
- Rid, Thomas. 2020. *Active Measures: The Secret History of Disinformation and Political Warfare*. New York: Farrar, Strauss and Giroux.
- Rosen, Guy. 2021, March 22. How we're tackling misinformation across our apps. *Facebook*. <https://about.fb.com/news/2021/03/how-were-tackling-misinformation-across-our-apps/>.
- Salganik, Marc. 2018. *Bit by Bit: Social Research in the Digital Age*. Princeton, NJ: Princeton University Press.
- Sharot, Talia. 2018. *The Influential Mind*. London: Little Brown.
- Shibutani, Tamatso. 1966. *Improvised News: A Sociological Study of Rumor*. New York: Bobbs-Merrill.
- Starbird, Kate, Arif Ahmer, Tom Wilson, Katherine Van Koevering, Katya Yefimova, and Daniel Scarnecchia. 2018. "Ecosystem or Echo-System? Exploring Content Sharing across Alternative Media Domains," *Presented at 12th International AAAI Conference on Web and Social Media (ICWSM 2018)*, Stanford, CA (pp. 365–374).
- Tankovska, H. (2021, February 2). Number of monthly active Facebook users worldwide as of 4th quarter 2020. Statista. <https://www.statista.com/statistics/264810/number-of-monthly-active-facebook-users-worldwide>.
- Williams, Robin. 1988. "Understanding Goffman's Methods." in *Erving Goffman: Exploring the Interaction Order*, edited by Paul Drew and Andrew Wooten. Cambridge: Polity Press.
- Winkin, Yves and Wendy Leeds-Hurwitz. 2013. *Erving Goffman: A Critical Introduction to Media and Communication Theory*. New York: Peter Lang.
- Woolley, Samuel and Philip Howard. 2019. *Computational Propaganda: Political Parties, Politicians and Political Manipulation on Social Media*. Oxford: Oxford University Press.

ABOUT THE CONTRIBUTOR(S)

Martin Innes is a Professor in the School of Social Sciences and Director of the Security, Crime and Intelligence Innovation Institute at Cardiff University.

Andrew Dawson trained as a Computer Scientist before joining the Security, Crime and Intelligence Institute at Cardiff where he specializes in the analysis of disinformation.