Changing Societies & Personalities, 2022
Vol. 6, No. 4, pp. 750–763
https://doi.org/10.15826/csp.2022.6.4.201

ARTICLE

Belief Formation in the Social Context: A Bayesian Decision Account

Francesco Rigoli
City, University of London, UK

ABSTRACT

A recent computational model argues that people's beliefs arise from two basic sources: a motivation to be accurate plus other forms of motivation such as pursuing economic interests. Yet, this previous proposal has focused exclusively on the individual level of analysis, neglecting the question of how social contexts shape belief formation. The present paper addresses this by examining how group dynamics underpin the formation of beliefs. The argument is that different social groups vary in their beliefs because, at least to some degree, each group acquires different life experiences and is motivated by different incentives. At the same time, the proposal is that each group can influence other groups' beliefs in three ways: by expressing opinions to other groups, by affecting the experiences acquired by other groups, and by setting incentives for other groups. This picture suggests that, within a community, three types of groups can often be identified: (a) the dominant groups, defined as those particularly capable of affecting other groups' experiences and incentives; (b) the intellectual groups, regarded as those whose opinions are particularly influential; and (c) the subaltern groups, encompassing people with minimal power to affect beliefs of others. By examining influence dynamics among these groups, the paper investigates how consensus or disagreement emerge within a community. Altogether, the paper offers insight on the interaction between social dynamics and psychological mechanisms that contribute to shape people's beliefs.

KEYWORDS

Bayesian decision, belief, influence, persuasion, dominant group, intellectual

Received 31 July 2022
Accepted 16 December 2022
Published online 30 December 2022

francesco.rigoli@city.ac.uk

Introduction

Beliefs can be defined as perceptions about whether a state of affair is true or false (e.g., "I believe that it is raining"). Research on the nature of beliefs is at the crucible of different social science disciplines, including psychology, neuroscience, sociology, and political science (Converse, 1964; Harris et al., 2008; Hutto & Ravenscroft, 2021; Jackson & Pettit, 1990; James, 1889). This research has raised several fundamental questions: What are the mechanisms underlaying the formation of beliefs? Why do some people hold some beliefs and other people embrace other beliefs? What are the social dynamics responsible for the formation of beliefs? To address these questions, two separate research perspectives have been developed. Focusing on the mental processes at play, one perspective (common in psychology and neuroscience) examines the individual level (Harris et al., 2008; James, 1889; Shermer, 2012), while the other perspective (common in sociology and political science) analyses social dynamics (Converse, 1964; Spillman, 2020; Swidler & Arditi, 1994). Individual and social approaches have each offered insight about the nature of beliefs. Yet, although individual approaches provide a fine-grained description of the psychological mechanisms, they typically disregard how these mechanisms arise in social contexts. Conversely, social perspectives explore social dynamics in detail, but often remain vague about the underlying psychological mechanisms. Accounts integrating both approaches are rare (Boudon, 1989; Goldthorpe, 1998; Opp, 2012); as a consequence, the interaction between the individual and social levels of analysis remains poorly understood.

This paper aims to integrate the individual and social levels of analysis by developing a broad theoretical framework describing how beliefs emerge within the social environment. The argument builds upon a recent computational model about the psychological processes underpinning belief formation (Rigoli, 2021a, 2021b, 2022; Rigoli et al., 2021). While this previous model focuses exclusively on the individual level of analysis, the goal here is to broaden the perspective and embed the model within a general framework where the psychological processes shape, and are simultaneously shaped by, social dynamics. Let us start by introducing the above-mentioned model in the next section.

Psychological Processes

The number of accounts examining the psychological mechanisms underpinning belief formation is enormous. Although an exhaustive overview of these accounts is beyond the scope of this paper, I propose to group these in two broad families. On the one hand, *explanation theories* presuppose that beliefs emerge from an attempt to describe reality accurately. Examples of these include the philosophical proposals of Plato, Aristotle, and Kant (Edwards, 1967), contemporary cognitive models claiming that Bayesian inference underpins mental processes (Friston, 2005; Knill & Pouget, 2004; Oaksford & Chater, 2007), and rational-choice theory (Boudon, 1989; Goldthorpe, 1998; Opp, 2012). Many of these accounts acknowledge the limits and biases of the human mind (often attributed to the restricted computational

capacity of the brain), but nonetheless maintain that, despite these mental constrains, beliefs ultimately reflect humans' effort to explain reality as accurately as possible (Gigerenzer & Selten, 2001). Contrary to explanation theories, *motivation theories* argue that beliefs arise from motives other than accuracy seeking. Examples are the philosophies of Marx, Nietzsche, and Freud (Edwards, 1967), as well as contemporary perspectives in social psychology such as cognitive dissonance theory (Festinger, 1962; Harmon-Jones & Mills, 2019) and social identity theory (Abrams & Hogg, 1990; Tajfel, 1982). For example, Marx famously argued that, far from being grounded upon a disinterested analysis of reality, the beliefs advocated by the dominant class conceal its economic interests (Marx & Engels, 1970).

Does empirical evidence lend more support to epistemic or to motivation theories? The picture emerging from empirical research suggests that each family of theories captures some, but not all, aspects of reality (Bartels, 2016; Jost & Banaji, 1994; Jost et al., 2009; Kunda, 1990). Thus, the emerging picture is one where epistemic and non-epistemic motives coexist and both contribute to the formation of beliefs. This means that a sound theory of belief formation should acknowledge both epistemic and non-epistemic drives. However, how these two kinds of drives interact has been rarely explored. A recent computational model developed by the author of the present paper has offered a possible answer to this question (Rigoli, 2021a, 2021b, 2022; Rigoli et al., 2021). The model relies on a Bayesian decision framework, and hence it is referred here to as the Bayesian Decision Model of Beliefs (BDMB). The BDMB focuses on the mechanisms through which individuals arbitrate among alternative hypotheses for explaining aspects of life and reality. To illustrate how the model works, consider the example of landowners operating within the slave system in the USA South before the civil war erupted in that country (Rigoli, 2021b). Imagine that these individuals formed their opinions about races by arbitrating between two hypotheses, one—a racist hypothesis—claiming that Blacks are genetically more violent and less intelligent than Whites, and the other—an antiracist hypothesis—claiming that Blacks are not different from Whites in any respect. The first hypothesis implies that, by providing slaves with a regulated existence under the landowner's paternalistic eye, the slave system benefits landowners and slaves alike. Conversely, the second hypothesis implies that the slave system is detrimental for slaves, inasmuch as they are exploited at the advantage of Whites. According to the BDMB, three factors establish which hypothesis will be endorsed by a landowner. The first factor is represented by prior beliefs, namely, by relevant assumptions already available before reasoning. Prior beliefs capture various aspects such as general views about the world or about society. For example, one prior belief might be the conviction that human races are genetically different; the alternative view might be that all races are characterized by an equal genetic endowment. A landowner who entertains the former prior belief, the argument goes, will be more likely to accept the racist hypothesis.

The second factor envisaged by the BDMB as being central to belief formation is represented by novel evidence, which can be experienced in two ways: directly, when evidence is conveyed by one's own perception, or socially, when social sources such as another person or the media provide indirect information (Rigoli, 2021b). In our example, spending time working with Blacks is an instance of direct evidence,

whereas listening to the opinions expressed by other landowners is an example of social evidence.

The third critical factor for belief formation proposed by the BDMB is represented by the utility, in terms of reward or punishment, expected if any hypothesis is accepted or rejected (Rigoli, 2021b). In our example, the landowner assesses the utility expected to occur (a) if the racist hypothesis is true and is accepted (and slavery is supported), (b) if the antiracist hypothesis is true and is accepted (and slavery is opposed), (c) if the racist hypothesis is false but is accepted (and slavery is supported), (d) if the antiracist hypothesis is false but is accepted (and slavery is opposed). To understand the influence of expected utility, compare two landowners, one who owns slaves and the other who does not. For the first landowner, accepting the racist hypothesis (and supporting racism) is obviously conducive to higher utility, in terms of economic incentives, compared to accepting the antiracist hypothesis (and fighting racism); the opposite is true of the second landowner. Therefore, the BDMB implies that the first landowner will be more likely to believe in the racist hypothesis compared to the second landowner[1].

In summary, according to the BDMB, three factors contribute to belief formation: prior beliefs, novel evidence (either direct or social), and expected utility. Note that, because of the influence of expected utility, a hypothesis might be selected because it is the costliest to reject even though it is not the best supported by evidence and by prior beliefs (i.e., even though it is not the most accurate). However, prior beliefs and novel evidence remain fundamental because a hypothesis will be less likely to be accepted if it is poorly supported by them. In this way, the BDMB integrates epistemic and non-epistemic motives, the former captured by the role assigned to prior beliefs and to novel evidence, the latter embodied by the expected utility component. The BDMB assumes that people are largely unaware of the factors shaping their beliefs, and simply perceive these beliefs as being true: in our example, a landowner might be staunchly convinced about the validity of the racist hypothesis, without realizing that this conviction does not stem from a disinterested analysis of reality.

In a nutshell, by integrating both epistemic and non-epistemic factors, the BDMB offers a systematic description of the various psychological processes responsible for belief formation. The implications of the BDMB have been examined in several domains, including the domain of political reasoning (Rigoli, 2021b), the domain of religious beliefs (Rigoli, 2021a), the domain of conspiracy theories (Rigoli, 2022), and the domain of delusional beliefs (Rigoli et al., 2021). Yet, all these domains pertain to the individual level of analysis. It remains to be assessed whether the BDMB can offer any insight about how the individual and social levels of analysis interact. In other words, can the BDMB elucidate how beliefs emerge during social interactions? The aim of the present paper is to employ the BDMB to explore how social dynamics shape belief formation.

[1] Although, in this example, expected consequences reflect economic interest, more generally the BDMB adopts a broad definition of expected consequences, encompassing any sort of value such as fostering social bonds, exerting power, realising bodily pleasures, or avoiding other people's suffering (e.g., Schwartz, 1992).

Social Context

Beliefs are not the result of an individual mind interacting with the environment, but of multiple minds interacting among themselves and with the environment. Along this line, scholars have argued that, in an evolutionary perspective, humans have developed a remarkable ability to share beliefs that has enabled them to cooperate effectively (Adolphs, 2009; Balliet & van Lange, 2013; Castelfranchi & Falcone, 2010; Dunbar, 2003). Following this argument, evolution is deemed to be responsible for humans' predisposition to influence other people's opinions and, at the same time, to be influenced by other people's opinions (Friedkin & Johnsen, 1990; Turner, 1991). Employing the BDMB as a framework, here I propose an explanation of how beliefs emerge when individuals interact with one another and with their environment. The argument starts by asking two apparently trivial, but in fact fundamental, questions. First, why do people vary in their beliefs? Second, how can a person influence other people's beliefs? This section addresses the first question, while the second question will be explored later.

Following the BDMB, an obvious reason why beliefs vary among people is that everyday activity, and thus the everyday experience one gains, varies greatly among people. Consider the European Medieval society as an example. This was arranged in three orders: the nobility, the clergy, and all other people who were collectively grouped in the third estate (Duby, 1982). The third estate, the clergy, and the nobility were engaged in working, praying, and fighting, respectively. Obviously, these different tasks imply that very different everyday experiences were acquired by the different orders: for example, a life spent cultivating the fields had little in common with a life spent praying or with a life spent fighting. Adopting the BDMB as a framework, different everyday experience implies different direct evidence—one of the key factors which, as explained above, shape beliefs according to the model. Thus, within the BDMB, differences in direct evidence among groups are one reason explaining why people vary in their beliefs.

However, according to the BDMB, differences in direct evidence are not the whole story: expected utility is also another critical factor. Remember that the BDMB posits that, beyond any role played by prior beliefs and by novel evidence, beliefs are shaped by expected utility. Regarding the latter, let us assume that any individual of a community is motivated by two types of values (Schwartz, 1992): *community values*, aiming at the well-being of the whole community, and *self-interest values*, aiming at one's own well-being or at the well-being of one's own specific group. Distinguishing community versus self-interest values is useful to understand why beliefs vary among people: although community values are typically shared by everyone, self-interest values are specific for each group. Consider again the example of the nobility, the clergy, and the third estate in the Middle Ages (Duby, 1982). It is safe to assume that all orders shared certain community values such as fostering agricultural productivity and protecting the community from invaders. Yet, it is also reasonable to imagine that self-interest values varied widely among the orders. For instance, the tolls destined to the Church were probably evaluated very differently by the nobles and by the clergy. Given divergent

self-interest values, the BDMB implies that each group embraces beliefs that are not shared by other groups. In the example of Medieval society, it is reasonable to imagine that, more than the nobility and the third estate, the clergy was convinced that the tolls destined to the Church were necessary for ensuring God's grace.

In summary, following the BDMB, two factors are proposed to explain why beliefs vary among people: differences in direct evidence and differences in self-interest values. Yet, as argued above, humans do not live in isolation, but they live in social environments. This means that each person's beliefs are constantly under the influence exerted by other people (Friedkin & Johnsen, 1990; Turner, 1991). Can the BDMB offer any insight on how such social influence works? The next section addresses this question.

Social Influence

We can employ the BDMB to interpret how a group can influence beliefs embraced by other groups. Following the BDMB, the proposal is that beliefs can be influenced in three ways (Fig. 1): (a) by novel information experienced via one's own perception (direct evidence), (b) by opinions expressed by others (social evidence), and (c) by changes in expected utility (note that prior beliefs cannot be influenced directly, but via direct or social evidence). Social influence can target any of these aspects. Moreover, it is reasonable to assume that groups vary in their ability to target each aspect. Some groups appear to be particularly capable to shape the direct evidence experienced by other groups. In the Middle Ages, an example of this is the nobility, with its power to determine, among other things, serfs' labour schedule. Moreover, there are groups who are especially influential upon the utility experienced by other groups. Again, Medieval nobles are an example of this, as evident from their power to request tolls from serfs. As the example of the Medieval nobility suggests, the groups capable of affecting direct evidence have often also great power to impact upon utility. These can be referred to as the *dominant groups* (Gramsci, 1948–1951/2011). However, groups with power regarding direct evidence and utility are often not as powerful in terms of affecting social evidence. This is because, in many complex societies, some groups emerge who, despite being unable to influence direct evidence and utility, have specialized in acquiring knowledge. The opinions expressed by these groups are therefore typically judged as highly trustworthy, endowing these groups with a special power to affect social evidence. In the Middle Ages, an example of this is the clergy, a group with poor influence upon direct evidence and upon utility but, at the same time, a highly powerful group in terms of affecting social evidence. Groups in this category, encompassing roles such as priests, academics, artists, and journalists, can be referred to as the *intellectual groups* (Gramsci, 1948–1951/2011).

In summary, the BDMB offers an interpretation about the nature of social influence processes. The model identifies three potential targets of influence: direct evidence, social evidence, and utility. Groups can be classified based on their power to influence each target. This classification highlights three broad categories of groups: (a) the dominant group, influential upon direct evidence and utility, (b) the

intellectual group, influential upon social evidence, and (c) the subaltern group, characterised by a minimal ability to influence the beliefs of other groups. Armed with this model of social influence, the next section enquires about how influence dynamics can lead to consensus or to disagreement within a community.

Figure 1

Role Played by Social Influence According to the BDMB

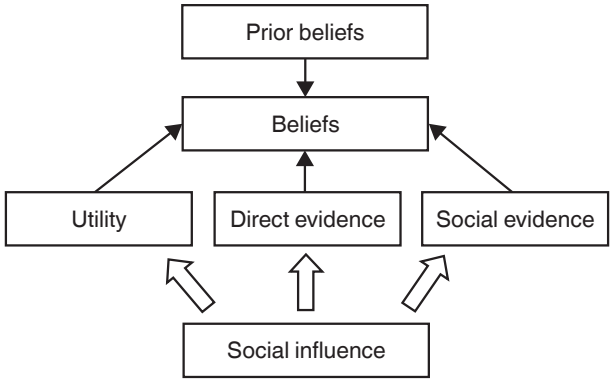

Note. Source: Developed by author.

Consensus Versus Disagreement

Let us focus on a simple scenario where the three groups identified above interact within a community: the dominant group shaping direct evidence and utility, the intellectuals influential upon social evidence, and the subaltern group encompassing people with little or no power in shaping beliefs (Gramsci, 1948–1951/2011). This scenario is obviously simplistic: in real communities, more than three groups can be identified, and, even more fundamentally, it is debatable how groups should be defined (e.g., based on class, gender, or ethnicity). However, as we shall see below, this simple scenario highlights interesting aspects emerging from the BDMB.

A community can be assessed based on whether the three groups identified above share their beliefs or not. Adopting this criterion, a community can manifest either consensus or disagreement, the former occurring when beliefs are shared, the latter when beliefs are not shared. Here, I aim to describe the characteristics of consensus and disagreement as interpreted by the BDMB.

Let us start by examining consensus. I propose to distinguish between two possible cases: *spontaneous consensus* versus *manipulated consensus*. The former occurs when the dominant group abstains from influencing the subaltern group (i.e., it abstains from affecting the latter's beliefs), and yet all groups share the same beliefs. During spontaneous consensus, the shared beliefs are grounded upon community values, rather than upon the self-interest values specific to the dominant group. Conversely, manipulated consensus occurs when beliefs are shared because the dominant group has actively modified the subaltern group's beliefs. When consensus

is manipulated, the shared beliefs are not grounded on community values, but on the self-interest values of the dominant group. Manipulated consensus occurs when the dominant group has successfully targeted the direct evidence or the utility of the subaltern group: as a consequence, the subaltern group is now convinced that, although the shared beliefs might not be palatable, they are nevertheless true. As an example of manipulated consensus, consider a dispute between employers and workers. Imagine that, despite poor supporting evidence, employers claim that a salary cut is necessary for the economy to remain aloft, and threat to fire workers if they protest against this measure. Threatened workers might end up accepting the argument that the salary cut is in fact necessary to protect the economy from unravelling. This is an example of manipulated consensus: both the dominant group (the employers) and the subaltern group (the workers) share the same belief, that is, they both believe that salary cuts are necessary; thus, we can talk about consensus. However, consensus is not spontaneous, but manipulated, because it is not grounded on evidence; rather, it arises from an attempt of the dominant group to pursue its own self-interest by targeting the utility of the subaltern group (by threatening to fire workers).

The idea of manipulated consensus can explain why subaltern groups sometimes accept stereotypes against themselves, such as when ethnic minorities embrace racist stereotypes (Jost & Banaji, 1994; Sagar & Schofield, 1980). The argument is that this occurs because, by targeting direct evidence or utility, the dominant group succeeds in persuading subaltern groups that the stereotypes are true. For example, the dominant group can arrange society in such a way that a subaltern group receives less education and thus performs worse in the future (e.g., people in this group obtain worse jobs, have poorer health, exhibit more criminal behaviour, etc.); in other words, this represents an instance in which the dominant group succeeds in affecting the direct evidence experienced by the subaltern group. Rather than being interpreted as due to environmental factors, the poorer performance of the subaltern group can in turn be interpreted as evidence consistent with a stereotype claiming that the subaltern group is congenitally inferior.

As another means to realise manipulated consensus, the dominant group can co-opt the intellectual group (Gramsci, 1948–1951/2011): for instance, the dominant group can reward intellectuals (e.g., in terms of economic wealth) if they sponsor the beliefs promoted by the dominant group. Given the trust typically attributed to intellectuals, this strategy is often highly effective in persuading the subaltern group. Historical examples of this are many, such as when, at the start of the 19th century, European church hierarchies claimed that the Ancient Regime, and not liberal regimes, conformed to God's will, encouraging the bulk of the peasantry to resist political change (Hitchcock, 2012; Vidmar, 2014)—the interpretation being that church hierarchies took this position in order to protect the interests of the ruling class in exchange of protection.

Now that we have examined the notion of consensus, let us focus on disagreement, which arises when groups do not share beliefs. This typically occurs when the dominant group, driven by self-interest values and neglecting community values, attempts to persuade the subaltern group but fails to do so. Here, there is a struggle to

persuade (from the perspective of the dominant group) and to resist persuasion (from the perspective of the subaltern group). As above, the dominant group's effort can target the direct evidence or the utility experienced by the subaltern group. Examples where the dominant group intervenes upon the direct evidence experienced by the subaltern groups are when employers rearrange the workplace with the aim to discourage interactions among workers, and when governments drift towards war as a way to distract the population from domestic concerns. Sometimes these attempts fail and leave the subaltern groups unconvinced about the claims of the dominant group, producing a condition of disagreement. The role of intellectuals is critical here too. During disagreement, intellectuals are sometimes co-opted by the dominant group—thus supporting the latter's beliefs—thanks to the rewards they receive in exchange. Other times, especially when they perceive that the rewards offered to them by the dominant group are inadequate, intellectuals support the subaltern group's beliefs.

In summary, the BDMD offers a framework to explain how influence dynamics can result in a state of consensus or in a state of disagreement among groups. Whether a community is characterised by consensus or by disagreement is important because, while disagreement spurs conflict, consensus allows cooperation to be established, and cooperation in turn is necessary to allow a community to mobilize all its available resources. Yet, the reasoning outlined above suggests that cooperation is not always beneficial for everyone: all can benefit from cooperating only when consensus is spontaneous, that is, when consensus is grounded on fostering community values. On the contrary, when consensus is the product of manipulation, that is, when it is instrumental to protecting the self-interest of the dominant group, cooperation benefits the dominant group alone while damaging the subaltern group.

Discussion

Previous literature has introduced the BDMB to examine the psychological processes at play when beliefs are formed (Rigoli, 2021a, 2021b, 2022; Rigoli et al., 2021). Yet, this literature has so far neglected the question of how social processes contribute to the formation of beliefs. Can the BDMD offer any insight on this question? The paper addresses this by proposing that, following the BDMD, different groups within a community are characterised by specific direct evidence and utility, and thus end up entertaining divergent beliefs. At the same time, groups are predicted to influence other groups' direct evidence, social evidence, and utility, thus impacting on the beliefs embraced by other groups. Dominant groups are defined as those particularly capable to affect other groups' direct evidence and utility, while intellectual groups are regarded as those who are powerful in shaping other groups' social evidence. Based on this, the BDMB offers a perspective to interpret the notion of consensus and disagreement among the different groups of a community.

Although the BDMB shares characteristics with many previous works, I highlight Gramsci's theory of ideology as particularly relevant (Gramsci, 1948–1951/2011). This is based on a Marxist outlook concerning the structure of society, grounded on the notion that classes, defined by their economic position, compete with one another in

the social arena. Gramsci's analysis focuses on the social superstructure, namely on the set of values and beliefs embraced by the different classes. The argument is that the dominant class (e.g., the bourgeoisie in the age of capitalism) builds an ideology which justifies the current social structure: although the dominant class exploits the dominated class (e.g., the workers), the argument is that the dominant class conceals this exploitation, or at least depicts it as just and unavoidable. The concept of hegemony is adopted by Gramsci to describe circumstances when the dominant class succeeds in persuading the dominated class: when the dominant class is hegemonic, it has no need to exert coercion, because the dominated class does not resist exploitation. Conversely, when the dominated class is not persuaded, class conflict ensues: in this scenario, the dominant class can preserve the status quo only through coercion. If coercion fails, the dominated class can revolt and overthrow the dominant class.

The similarities between Gramsci's perspective and the BDMB are numerous. In both accounts, different groups tend to develop their own specific beliefs. The psychological mechanisms underlying this remain unspecified in Gramsci's writings, while the BDMB interprets these mechanisms in terms of experiencing different direct evidence and utility. Moreover, both accounts postulate that groups influence and are influenced by other groups, with some groups being more influential than others; again, while the nature of this influence is not fully specified in Gramsci's writings, the BDMB offers an explicit analysis by identifying influences upon direct evidence, social evidence, and utility. Another similarity is between the idea of hegemony and the idea of consensus or, more precisely, between the idea of hegemony and the idea of manipulated (but not spontaneous) consensus. Finally, both theories highlight the key role played by intellectuals in shaping the beliefs embraced by other groups.

Notwithstanding many similarities, Gramsci's view and the BDMB have also important differences. While the former is firmly grounded on a Marxist notion of social structure, the latter is more flexible and independent of whether a Marxist social structure is postulated or not. In addition, consistent with a Marxist outlook, Gramsci presupposes that self-interest is the driver of belief formation, while the BDMB argues that both self-interest and community values can be important. Finally, the psychological processes driving belief formation remain poorly specified in Gramsci's writings: it remains obscure to what extent beliefs reflect an attempt to be accurate or to fulfil other motives (in the terminology of this paper, it is not clear whether Gramsci's proposal can be classified as an explanation or as a motivation theory). The BDMB offers an explicit description of this aspect, clarifying the specific contribution of epistemic motives captured by evidence and prior beliefs and non-epistemic drives captured by utility.

Limitations

The argument proposed here relies on simplifications that should be carefully reconsidered by future research. For example, when discussing the notion of consensus and disagreement, I assumed the existence of three social groups only (the dominant, intellectual, and dominated group). This is obviously simplistic. For example, the dominant group is rarely monolithic, encompassing subgroups such as political, military, or economic elites, each with its own self-interest values and its own life experience.

The same applies to the dominated and intellectual groups, often including subgroups based on region, age, gender, and class. Moreover, in realistic cases, it is debatable if a clear-cut distinction between dominant and dominated group, with only the former affecting beliefs of the latter, is meaningful; adopting a more fine-grained approach where, to some degree, all groups influence other groups is likely to be preferrable in most cases. Likewise, the idea that social evidence is affected only by intellectuals is dubious: one can even debate whether intellectuals should be considered as a separate group, especially in recent times where social media allow everyone to play the role of intellectuals to some degree. The notions of consensus and disagreement appear simplistic as well: with numerous groups at play, consensus can rarely be total, and spontaneous consensus might emerge for some groups but not for others. Despite these shortcomings, simplifications are helpful to highlight the key points raised by the BDMB. A promising research avenue is nonetheless to progressively add complexity to the problem and explore the theory in more elaborated settings.

Future Directions

The framework articulated in the paper opens up a set of potential avenues for future research. First, the BDMB can inspire empirical research on social influence processes. So far, research on this has focused primarily on which features of a message are most effective in changing the beliefs embraced by a target person (Turner, 1991). In the context of the BDMB, this prior research can be interpreted as being focused on social evidence alone, which is only one component of the picture offered by the theory. Indeed, the BDMB not only identifies social evidence, but it also stresses the role of direct evidence and utility as potential targets of social influence. These two targets have been largely neglected by research so far. The BDMB encourages empirical investigation to address this by exploring how social influence can affect direct evidence and utility.

The BDMB can inspire further research also by providing a conceptual framework to interpret dynamics characterising society in the present or in the past, and to examine how these dynamics shape people's beliefs. Above, I have offered a sketch of how this analysis may look like in the context of European Medieval Society. A possibility is to extend this analysis further by exploring in detail the historical circumstances in which the actions performed by one group affected the direct evidence, social evidence, or utility of other groups, and how this had an impact on the beliefs embraced by the other groups. Of course, this type of analysis can be applied to Medieval societies as well as to any other form of social organization in the past or in the present.

Conclusion

To summarize, the paper develops a theory describing how social processes underpin the formation of beliefs. The aim of this theory is to bridge perspectives focusing on the individual domain with perspectives examining the social domain. Such an integrated approach, I argue, is ultimately necessary, because individual and social processes are deeply intertwined during belief formation.

References

Abrams, D. E., & Hogg, M. A. (1990). *Social identity theory: Constructive and critical advances*. Springer-Verlag Publishing.

Adolphs, R. (2009). The social brain: Neural basis of social knowledge. *Annual Review of Psychology*, *60*, 693–713. https://doi.org/10.1146/annurev.psych.60.110707.163514

Balliet, D., & van Lange, P. A. M. (2013). Trust, conflict, and cooperation: A meta-analysis. *Psychological Bulletin*, *139*(5), 1090–1112. https://doi.org/10.1037/a0030939

Bartels, L. M. (2016). *Unequal democracy: The political economy of the new gilded age* (2nd ed.). Princeton University Press.

Boudon, R. (1989). Subjective rationality and the explanation of social behavior. *Rationality and Society*, *1*(2), 173–196. https://doi.org/10.1177/1043463189001002002

Castelfranchi, C., & Falcone, R. (2010). *Trust theory: A socio-cognitive and computational model*. John Wiley & Sons. https://doi.org/10.1002/9780470519851

Converse, P. E. (1964). The nature of belief systems in mass publics. In D. E. Apter (Ed.), *Ideology and discontent* (pp. 206–261). Free Press of Glencoe.

Duby, G. (1982). *The three orders: Feudal society imagined* (A. Goldhammer, Trans.). University of Chicago Press. (Originally published in French 1978)

Dunbar, R. I. M. (2003). The social brain: Mind, language, and society in evolutionary perspective. *Annual Review of Anthropology*, *32*, 163–181. https://doi.org/10.1146/annurev.anthro.32.061002.093158

Edwards, P. (Ed.). (1967). *The encyclopedia of philosophy* (Vols. 1–8). Macmillan.

Festinger, L. (1962). *A theory of cognitive dissonance*. Stanford University Press. (Originally published 1957)

Friedkin, N. E., & Johnsen, E. C. (1990). Social influence and opinions. *The Journal of Mathematical Sociology*, *15*(3–4), 193–206. https://doi.org/10.1080/0022250X.1990.9990069

Friston, K. (2005). A theory of cortical responses. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *360*(1456), 815–836. https://doi.org/10.1098/rstb.2005.1622

Gigerenzer, G., & Selten, R. (Eds.). (2001). *Bounded rationality: The adaptive toolbox*. The MIT Press.

Goldthorpe, J. H. (1998). Rational action theory for sociology. *British Journal of Sociology*, *49*(2), 167–192. https://doi.org/10.2307/591308

Gramsci, A. (2011). *Prison Notebooks* (J. A. Buttigieg & A. Callari, Eds., & Trans., Vols. 1–3). Columbia University Press. (Originally published in Italian 1948–1951)

Harmon-Jones, E., & Mills, J. (2019). An introduction to cognitive dissonance theory and an overview of current perspectives on the theory. In E. Harmon-Jones (Ed.), Cognitive dissonance: Reexamining a pivotal theory in psychology (2nd ed., pp. 3–24). American Psychological Association. https://doi.org/10.1037/0000135-001

Harris, S., Sheth, S. A., & Cohen, M. S. (2008). Functional neuroimaging of belief, disbelief, and uncertainty. *Annals of Neurology*, *63*(2), 141–147. https://doi.org/10.1002/ana.21301

Hitchcock, J. (2012). *History of the Catholic Church: From the Apostolic Age to the third millennium*. Ignatius Press.

Hutto, D., & Ravenscroft, I. (2021). Folk psychology as a theory. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (Fall 2021 ed.). Stanford University. https://plato.stanford.edu/archives/fall2021/entries/folkpsych-theory/

Jackson, F., & Pettit, P. (1990). In defence of folk psychology. *Philosophical Studies*, *59*(1), 31–54. https://doi.org/10.1007/bf00368390

James, W. (1889). The psychology of belief. *Mind*, *14*(55), 321–352. https://doi.org/10.1093/mind/xiv.55.321

Jost, J. T., & Banaji, M. R. (1994). The role of stereotyping in system-justification and the production of false consciousness. *British Journal of Social Psychology*, *33*(1), 1–27. https://doi.org/10.1111/j.2044-8309.1994.tb01008.x

Jost, J. T., Federico, C. M., & Napier, J. L. (2009). Political ideology: Its structure, functions, and elective affinities. *Annual Review of Psychology*, *60*, 307–337. https://doi.org/10.1146/annurev.psych.60.110707.163600

Knill, D. C., & Pouget, A. (2004). The Bayesian brain: The role of uncertainty in neural coding and computation. *Trends in Neurosciences*, *27*(12), 712–719. https://doi.org/10.1016/j.tins.2004.10.007

Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin*, *108*(3), 480–498. https://doi.org/10.1037/0033-2909.108.3.480

Marx, K., & Engels, F. (1970). *The German ideology* (C. J. Arthur, Ed.). International Publishers Co. (Originally published in German 1932)

Oaksford, M., & Chater, N. (2007). *Bayesian rationality: The probabilistic approach to human reasoning*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780198524496.001.0001

Opp, K.-D. (2012). Collective identity, rationality and collective political action. *Rationality and Society*, *24*(1), 73–105. https://doi.org/10.1177/1043463111434697

Rigoli, F. (2021a). A computational perspective on faith: Religious reasoning and Bayesian decision. *Religion, Brain & Behavior*, *11*(2), 147–164. https://doi.org/10.1080/2153599X.2020.1812704

Rigoli, F. (2021b). Masters of suspicion: A Bayesian decision model of motivated political reasoning. *Journal for the Theory of Social Behaviour*, *51*(3), 350–370. https://doi.org/10.1111/jtsb.12274

Rigoli, F. (2022). Deconstructing the conspiratorial mind: The computational logic behind conspiracy theories. *Review of Philosophy and Psychology*. Advance online publication. https://doi.org/10.1007/s13164-022-00657-7

Rigoli, F., Martinelli, C., & Pezzulo, G. (2021). I want to believe: Delusion, motivated reasoning, and Bayesian decision theory. *Cognitive Neuropsychiatry*, *26*(6), 408–420. https://doi.org/10.1080/13546805.2021.1982686

Sagar, H. A., & Schofield, J. W. (1980). Racial and behavioral cues in Black and White children's perceptions of ambiguously aggressive acts. *Journal of Personality and Social Psychology*, *39*(4), 590–598. https://doi.org/10.1037/0022-3514.39.4.590

Schwartz, S. H. (1992). Universals in the content and structure of values: Theoretical advances and empirical tests in 20 countries. *Advances in Experimental Social Psychology*, *25*, 1–65. https://doi.org/10.1016/S0065-2601(08)60281-6

Shermer, M. (2012). *The believing brain: From spiritual faiths to political convictions—How we construct beliefs and reinforce them as truths*. Robinson. (Originally published 2011)

Spillman, L. (2020). *What is cultural sociology?* John Wiley & Sons.

Swidler, A., & Arditi, J. (1994). The new sociology of knowledge. *Annual Review of Sociology*, *20*, 305–329. https://doi.org/10.1146/annurev.so.20.080194.001513

Tajfel, H. (1982). Social psychology of intergroup relations. *Annual Review of Psychology*, *33*, 1–39. https://doi.org/10.1146/annurev.ps.33.020182.000245

Turner, J. C. (1991). *Social influence*. Thomson Brooks/Cole Publishing Co.

Vidmar, J. (2014). *The Catholic Church through the ages: A history* (2nd ed). Paulist Press.