To cite this chapter:

De Houwer, J., & Moors, A. (2015). Levels of analysis in social psychology. In B. Gawronski & G. V. Bodenhausen (Eds.), *Theory and explanation in social psychology* (pp. 24-40). Guilford Press.

RUNNING HEAD: LEVELS OF ANALYSIS

Levels of Analysis in Social Psychology

Jan De Houwer* and Agnes Moors

Ghent University, Belgium

*Corresponding author

Mailing address: Ghent University, Henri Dunantlaan 2, B-9000 Ghent, Belgium. Email: Jan.DeHouwer@UGent.be, Agnes.Moors@UGent.be . Tel: ++32 9 264 64 62, Fax: ++32 9 264 64 89.

Levels of Analysis in Social Psychology

Like many aspects of human behavior, scientific research is a goal-directed activity. Hence, to understand scientific behavior, we need to understand the goals that drive scientists. Scientific goals first of all refer to an object that is under investigation. One could, for instance, argue that social psychology aims at understanding social behavior. Because many questions can be asked about any single object, a scientific goal can also specify the question that is asked about that object. For instance, social cognition researchers are specifically interested in the mental processes that underlie social behavior whereas social neuroscientists might be interested primarily in the neural processes that underlie social behavior. In the present chapter, we focus on yet another way in which scientific goals differ: the level of analysis at which the scientific question is analyzed. Even a single question about a single object can be answered at different levels. In the present chapter, we illustrate this point in the context of social cognition research. However, in doing so, we also clarify the way in which social cognition research is related to other approaches in social psychology, including functional and neuroscience approaches.

The foundation of this chapter is provided by the influential ideas of David Marr (1982) about levels of analyses in cognitive psychology. In a recent paper (De Houwer & Moors, 2013), we extended the ideas of Marr by providing a more fine-grained analysis of the different levels and how these levels relate to each other. In the present chapter, we summarize these ideas and relate them to research in social psychology. We hope that reading this chapter will help researchers to make explicit the scientific goals that drive their research. Doing so not only sheds light on scientific practice but can also help facilitate research by

exposing false scientific debates and revealing links between seemingly unrelated or contradictory lines of research.

Marr's three levels of analysis

Marr's (1982) ideas about levels of analysis were developed in the context of cognitive psychology. The goal of cognitive psychology is to understand the mental processes that underlie behavior, that is, the way in which humans and other animals acquire, process, and represent information and how this processing of information guides their behavior (Gardner, 1987). Social cognition research thus shares with cognitive psychology the question that is asked about behavior (i.e., what are the underlying mental processes) but differs from other cognitive research in that it focusses on social behavior rather than behavior in general. Hence, Marr's ideas about the level of analysis of mental processes are particularly relevant for social cognition research, but as we will see later on in the chapter, they are also important for other traditions in social psychology.

In his highly influential book, Marr (1982) argued that processes can be analyzed at three levels: the computational level, the algorithmic level, and the implementational level. The computational level involves an analysis of what the process does. More specifically, it aims to specify what the output of a process will be given a certain input. In doing so, it provides insight in the function of the process. Take the example of a cash register (Marr, 1982). Its function is to add the price of different items. "Addition" is a function in that it specifies which input (e.g., 15 + 14) into the cash register leads to which output of the cash register (e.g., 29). Understanding that a cash register has the function to add up numbers provides an important insight into the nature of the cash register and thus constitutes and important level of analysis. Marr (1982) argued that, in a similar way, understanding the

function of mental processes is vital to our understanding of those processes. It specifies the task that the processes carry out.

The algorithmic level goes beyond the computational level by specifying the mechanism by which the input is transformed into the output. Again consider the example of a cash register. When entering the numbers 15 and 14 into the cash register, the output will be 29. But there are potentially different ways in which the input (15 and 14) can be transformed into the output (29). For instance, the numbers might be represented in Arabic form, first adding the units (i.e., 5 and 4), "carrying over" a number if this exceeds 10, and then adding the tens (i.e., 1 and 1). Another possible algorithm would be to represent the numbers in analog form (e.g., 15 dots and 14 dots), and then counting the total number of dots. "An algorithm thus specifies a chain of intermediate steps that starts from the input and ends at the output. When applied to mental processes, algorithms can be conceived of as mental mechanisms (Bechtel, 2005, 2008). The mechanism is mental in that the organism is assumed to represent information at each step of the chain and to transform the information from one step to the next" (De Houwer & Moors, 2013, p. 4).

The implementational level of analysis deals with the physical implementation of the algorithm. A cash register, for instance, can operate mechanically "as positions on a ten-notch metal wheel, or as binary coded decimal numbers implemented in the electrical states of digital logic circuitry" (McClamrock, 1991, p. 190). In a similar way, cognitive psychologists assume that a single information processing algorithm can be implemented in different physical systems (e.g., a human brain, an electronic computer; see Gardner, 1987).

In sum, Marr (1982) put forward the idea that understanding a mental process boils down to understanding what the process does (i.e., which input is transformed into which output), how the process achieves this function (i.e., the chain of intermediate information

processing steps via which an input is transformed into an output), and how the process is physically implemented. The three levels can be separated because they are not entirely interdependent. That is, conclusions at one level do not necessarily imply conclusions at another level. This follows from the fact that one function (i.e., specified at the computational level) can be achieved by different algorithms (i.e., specified at the algorithmic level) which can each be implemented in different physicals systems (i.e., specified at the implementational level).

One of the biggest advantages of Marr's (1982) distinction between levels of analysis, is that it can expose false debates in psychology. As we pointed out before, analyses at different levels are guided by different goals. A computational analysis aims to uncover the relation between the input and output of a process, an algorithmic analysis is concerned with the mechanisms by which the input is transformed in the output, and an implementational analysis has the goal to clarify how an algorithm is instantiated in a physical substrate such as the brain. Different theories may be compatible when they are directed at understanding different aspects of the same question. Hence, different answers to a scientific question (e.g., "what mental processes underlie social behavior?") are not necessarily conflicting if they are formulated at different levels of analysis.

To illustrate this point, let us consider research on causal attribution (also see De Houwer & Moors, 2013). Because we will return to this example several times during the course of this chapter, we will first provide some background about this type of research. Causal attribution research has a long history in social psychology. Perhaps most important in this respect is Kelley's (1973) covariation theory of causal attribution. A central question in this line of research is: How do people infer that the behavior of another person (e.g., Mary refuses Jim's invitation to dinner on Friday) is caused by something about the person (e.g.,

Mary is impolite), something about the entity (e.g., Jim is boring), or something about the circumstances (e.g., bad timing). The central assumption in Kelley's theory is that perceivers solve this inferential task by means of information about covariations between the presence or absence of the behavior, on the one hand, and the presence or absence of the person (i.e., consensus; e.g., do other people also refuse Jim's invitations), the presence or absence of the entity (i.e., distinctiveness; e.g., does Mary also refuse invitations from others), or variations across time (i.e., consistency; e.g., does Mary refuse invitations for Jim also at other points in time), on the other hand.

In studies on causal attribution, participants typically see a series of trials on which cues (e.g., a particular person) and outcomes (e.g., a particular social behavior) can be present or absent. Most often, they are asked to judge the extent to which a particular cue causes a particular outcome (e.g., the extent to which a behavior is caused by a particular person). Some years ago, there was a debate between so-called statistical models, such as the Delta P model, and association formation models, such as the Rescorla-Wagner model (see Cheng & Holyoak, 2011, De Houwer & Beckers, 2002, and McClure, 1998, for reviews). The Delta P model, for instance, postulates that causal judgments are a function of the difference between the probability of the outcome when a cue is present and the probability of the outcome when the cue is absent (i.e., $P(O/C) - P(O/\sim C)$). Association formation models such as the Rescorla-Wagner model, on the other hand, argue that causal judgments reflect the strength of associations that are formed between the representations of cues and the representations of outcomes in memory. The debate involved proponents of one theory raising doubts about the validity of the other theory by showing that people sometimes give causal judgments that seem to be in line only with the predictions of their own theory. In many cases, however, these attacks were countered by showing that some revised version of the criticized model

could account for those results as well (see De Houwer & Beckers, 2002, for a review). Inspired by Marr's (1982) ideas, researchers eventually realized that the two types of models are situated at different levels of analysis (Cheng, 1997; Shanks, 1995). Statistical models such as the Delta P model are computational in nature because they specify which environmental input (i.e., the presence or absence of cues and outcomes) leads to which behavioral output (i.e., a causal judgment or causal attribution). They do not make any assumptions about the mental mechanisms by which the input leads to the output. Associative models do make assumptions about the mechanism by which events lead to causal judgments (i.e., the formation of associations) and are thus situated at the algorithmic level of analysis. The fact that both models do not contradict each other was even proven mathematically: Given certain assumptions about the parameters in the Rescorla-Wagner model, this model makes exactly the same predictions as the Delta P model. In other words, both models are computationally equivalent in that they relate the same inputs to the same outputs. However, the Rescorla-Wagner model goes beyond the Delta P model in that it specifies not only which input leads to which output but also incorporates assumptions about the mental mechanism by which an input results in an output (i.e., the formation of associations in memory). This example illustrates that it is important to know whether a model is directed at describing a process in terms of input-output relations or in terms of the mechanism by which input is transformed into output. Different models at different levels of analysis are therefore not necessarily incompatible. Of course, they can be incompatible (as would be the case when an algorithmic and computational model would make different predictions about which input leads to which output) but they can also complement each other.

Beyond Marr

In a recent paper (De Houwer & Moors, 2013), we tried to go beyond the ideas of Marr (1982) in two ways. First, we argued that the computational level actually encompasses two different levels: the behavioral functional level and the informational functional level. Second, we examined in more depth how the different levels of analysis are related. In this section, we explain these ideas in relation to social cognition research and discuss their implications for social psychology as a whole.

Two functional levels of analysis

In Marr's (1982) cash register example, the input (i.e., entering numbers into the register) and output of the process (i.e., the number appearing on the display of the register) are both observable events in the physical world. This is also the case for mental processes that have environmental events as their input and behavioral events as their output. We refer to these mental processes as overarching mental processes in that they span the entire black box between an environmental input and a behavioral output. For instance, causal attribution could be regarded as an overarching mental process by which events in the world result in causal judgments. Other mental processes, however, have mental representations as their input and/or as their output. We refer to these processes as mental sub-processes because they constitute only one part in a chain of mental sub-processes by which environmental input influences behavior. For instance, the overarching mental process of causal attribution most likely entails the mental sub-processes of information storage and information retrieval. Information storage has environmental events as its input and mental (memory) representations as its output. Likewise, information retrieval has mental (memory) representations as its input and behavior or other mental representations as its output.

We argued that two functional levels of analysis can be distinguished based on the type of mental processes that are analyzed. The behavioral functional level reveals the function of overarching mental processes. The informational functional level, on the other hand, uncovers the function of mental sub-processes (i.e., mental processes that have mental representations as their input or output). Crucially, both levels are functional in that they describe which input leads to which output. They are silent with regard to the mechanism by which input is transformed into output. As in Marr's (1982) framework, the analysis of the nature of the mechanism is still reserved for the algorithmic level.

The distinction between the behavioral functional level and the informational functional level is meaningful because they do not overlap completely. Most importantly, one conclusion at the behavioral functional level is compatible with several competing conclusions at the informational functional level. Hence, understanding the behavioral functional level does not provide a full understanding of the informational functional level. This point can also be illustrated in the context of research on causal attribution, more specifically on the phenomenon known as discounting (see De Houwer & Moors, 2013). Discounting relates to the ability to distinguish between genuine and spurious causes. It is related to many aspects of social behavior, including the fundamental attribution error (Ross, 1977) which describes the failure to discount the causal role of dispositional factors in the presence of situational factors that may cause the same behavior. A common effect of this failure is that perceivers draw correspondent inferences about people's dispositions on the basis of observed behavior even when the behavior might have been caused by situational factors (i.e., the correspondence bias; see Gilbert & Malone, 1995).

In studies on causal attribution, discounting refers to the observation that causal judgments are lower when relations are redundant than when they are not redundant (e.g,

McClure, 1998). Assume that Person A and Person B most often win a difficult game when they play together. Some participants also see games in which Person A plays alone and is as likely to win as when Person A and Person B play together [i.e., P(win/A) = P(win/A.B)]. Discounting occurs when judgments about the ability of Person B to win the game when playing alone are lower when Person A is seen to win games than when Person A never plays alone. In the former case, the relation between B (cue) and winning (outcome) can be described as redundant in that the presence of B does not change the probability that the game will be won (Cheng & Novick, 1990). Hence, one can conclude at the behavioral functional level that causal attribution (as an overarching mental process) depends on the redundancy of relations in the environment. There are different causal judgments (output of the causal attribution process) for redundant relations than for non-redundant relations (input of the causal attribution process).

For the present purposes, it is important to realize that this conclusion at the behavioral functional level has been attributed by some to functional properties of the information storage sub-process whereas others have attributed it to the functional properties of the information retrieval sub-process. More specifically, some have argued that in case of redundancy, the occurrence of AB-win games (input of the information storage sub-process) does not result in the storage of information about the relation between B and winning (output of the information storage sub-process; e.g., Van Overwalle & Timmermans, 2005). Others, however, argued that information about the relation between B and winning is stored in memory (and can thus serve as input of the information retrieval sub-process) but has little impact on causal judgments (output of the information retrieval sub-process) when the relation is redundant (see De Houwer, Beckers, & Vandorpe, 2005, for a review). This example illustrates that assumptions about relations between the input and output of mental

sub-processes (i.e., analyses at the informational functional level) can vary while assumptions about the input and output of overarching mental processes (i.e., analyses at the behavioral functional level) remain constant. Hence, both levels do not overlap. As noted above, both functional levels also do not overlap with the algorithmic level in that the latter level deals with how input is transformed into output, not merely with which input is transformed into which output.

How are the behavioral functional level, the informational functional level, and the algorithmic level related?

We already noted that the behavioral functional level, the informational functional level, and the algorithmic level do not overlap because a theory at one level can be compatible with multiple competing theories at other levels. Put differently, there is a one-to-many relation between certain levels. However, De Houwer and Moors (2013) pointed out that these three levels of analysis are interdependent in two important ways. First, in some cases, one theory at one level implies a single theory at another level. In other words, sometimes there is a one-to-one relation going from one level to another. Second, even when one theory at one level does not imply one theory at another level, often it does *constrain* the theories that are possible at another level.

One-to-one relations can exist between different levels of analysis but only if certain conditions are fulfilled and only going from some levels to some other levels. Consider the algorithmic level. For any given input, an algorithm will produce an output. By applying the algorithm to a variety of inputs, one can therefore deduce a functional model of the process, that is, a description about what the input-output relations of the mental process are. Hence, provided that all elements of the algorithm are specified (e.g., all parameter are given a value) and certain background assumptions are made (e.g., that a person can see and hear), an

algorithmic model implies a functional model. Actually, we already discussed one example of such a one-to-one relation, namely the fact that the Delta P model of causal attribution (which describes causal attribution at the behavioral functional level by specifying which environmental events lead to which judgments) can be derived from the Rescorla-Wagner model (i.e., an algorithmic model of causal attribution). We now see that the expression "be derived from" actually means that, when the parameters of the Rescorla-Wagner model are fixed in a certain way, the Rescorla-Wagner models predicts for each possible set of learning events (input), the same causal judgments (output) as the Delta P model. Likewise, assumptions at the informational functional level can sometimes imply conclusions at the behavioral functional level. For instance, based on the premises (made at the informational functional level) that (a) causal attribution involves information storage and (b) redundant events are not stored in memory, one can conclude (at the behavioral functional level) that redundant events should not influence causal judgments.

Whether there can be one-to-one relations between levels of analysis also depends on the level at which the premises are stated and the level at which the conclusions are drawn. For instance, whereas making assumptions at the algorithmic level can allow one to deduce a single conclusion at the behavioral functional level, assumptions at the behavioral functional level (e.g., that redundant events do not influence causal judgments) do not imply a single conclusion at the algorithmic level (e.g., the absence of an impact of redundant events on judgments could be due to several algorithms). In other words, whereas there is a one-to-one relation starting from the algorithmic level and going to the behavioral functional level, there is a one-to-many relation starting from the behavioral functional level and going to the algorithmic level. Hence, the interdependence between different levels of analysis is

asymmetrical. Figure 1 (taken from De Houwer & Moors, 2013) provides an overview of the types of relations between the different levels.

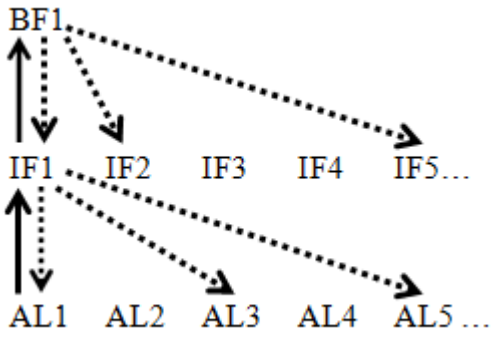

Figure 1. Schematic overview of the relations between the behavioral functional level, the informational functional level, and the algorithmic level of analysis. BF1 stands for one possible theory at the behavioral functional level. IF1 to IF5 stand for five different potential theories at the informational functional level. AL1 to AL5 stand for five different possible theories at the algorithmic level. Solid arrows symbolize one-to-one relations in that they point from a specific theory at a specific level of analysis to one theory at another level of analysis. Dotted arrows symbolize one-to-many relations in that they point from a specific theory at a specific level of analysis to several theories at another level of analysis (taken from De Houwer & Moors, 2013).

Even when a theory at one level does not imply a single theory at another level, it still constrains the possible range of theories at the other level. In more technical terms, a one-to-many relation is not the same as a one-to-all relation. For instance, assumptions at the behavioral functional level (e.g., that redundancy influences causal judgments) constrain conclusions at the algorithmic level (e.g., algorithms need to explain the effect of redundancy), even though they do not allow one to infer a single set of conclusions at those other levels (e.g., there might be several possible mechanisms that produce redundancy effects).

The previous paragraphs are related to two well-known rule-of-thumbs from methodology classes: “theories can be used to derive predictions about data” (Principle 1)

and “data can be used to reject theories but not to prove a theory” (Principle 2). In essence, our ideas go beyond these well-known principles by pointing out that (1) there are several levels of analysis at which theories can be formulated and (2) depending on the level at which it is formulated, a theory can imply a theory or merely constrain theories at other levels. Theories can address the function of an overarching mental process (i.e., a theory situated at the behavioral functional level), the function of a mental sub-process of this overarching mental process (i.e., a theory situated at the informational functional level), or the mechanism by which the function of a mental process (either overarching or sub-process) is realized (i.e., a theory at the algorithmic level). Hence, we can add to Principle 1 that some theories can be used not only to derive predictions about the data that will be obtained in a specific experiment but also to derive theories at other levels of explanation (e.g., the Delta P model can be derived from the Rescorla-Wagner model). We can add to Principle 2 that also the acceptance of certain theories at certain levels can constrain the theories that are possible at other levels without actually “proving” one specific theory at that other level (e.g., the claim that discounting is due to a lack of information storage rather than information retrieval, excludes algorithms that store information regardless of redundancy).

We would also like to add a third point that goes beyond traditional ideas about the relation between theory and data (also see De Houwer, 2011; De Houwer, Barnes-Holmes, & Moors, in press; De Houwer & Moors, 2013). We agree with the axiom that data merely constrain rather than prove theory but believe that data constrain theories at different levels of analysis to different degrees. Data constrain theories at the behavioral functional level to the largest extent because these theories specify only which aspects of the environment lead to which aspects of behavior. Whether something in the environment influences a behavior cannot be observed directly (which is why the behavioral functional level is a level of

analysis rather than mere description) but it can be inferred from data generated from experiments. Hence, experimental research puts strong constraints on theories at the behavioral functional level and thus provides a good source for conclusions at that level (also see De Houwer, 2011; De Houwer & Moors, 2013).

Data put much less constraints on informational functional theories and even less on algorithmic theories. Whereas data are essentially behaviors that are observed under certain environmental conditions, informational functional theories contain assumptions about input-output relations that involve mental representations that intervene between the environment and behavior (De Houwer, 2011; Gardner, 1987). Hence, informational theories go beyond data in two ways: (1) they imply conclusions about which input leads to which output (as is the case for theories at the behavioral functional level); (2) they make assumptions about the content of mental representations in the organism that performs the behavior (whereas data only specify environmental conditions and behaviors). [1]

Algorithmic models go even further beyond the data by adding assumptions about the mental mechanism via which an input is transformed into an output and the format of the representations on which these mechanisms operate. In addition, algorithmic models often include assumptions about how these mental representations are formed, transformed, and influence behavior. Hence, at the algorithmic level, even more needs to be inferred from the data. It seems reasonable to assume that the more that needs to be inferred from the data, the more difficult it is to arrive at firm conclusions. De Houwer and Moors (2013) introduced the term “inductive penetrability” to refer to the extent to which data constrain theories at the different levels of analysis. They concluded that the behavioral functional level is more penetrable than the informational functional level which is in turn more penetrable than the algorithmic level.

Our analysis raises the question of how one can arrive at conclusions at the informational functional and algorithmic levels. One option is to look for so-called proxies of mental representations and mental mechanisms. Proxies are specific behavioral effects that can be treated as if they are a perfect reflection of some mental construct. For instance, ratings on an attitude scale can be taken as a proxy for mental attitudes. If the rating is positive, then the attitude is positive. If the rating is negative, then the attitude is negative. However, the use of proxies requires the unlikely assumption that the proxies are determined only by one mental construct (e.g., De Houwer, 2011). For instance, attitude ratings as a proxy of attitudes makes sense only if the ratings are determined only by the attitudes in memory. However, it has been shown that attitude ratings depend on a wide variety of factors that are not related to attitudes, such as the order in which ratings have to be completed (e.g., Schwarz & Bohner, 2001). Many more proxies of attitudes have been proposed, but none provides a pure index of attitudes, probably because behavior is never determined only by one type of representation or process (see De Houwer, 2011; De Houwer, Garwronski, and Barnes-Holmes, 2013, for a detailed discussion of these issues in the context of attitude research).

There is, however, another way to draw conclusions at the informational functional and algorithmic levels. We know that data constrain the behavioral functional level. We also know that the behavioral functional level constrains the informational functional level which in turn constrains the algorithmic level. Hence, if we proceed carefully from the data to behavioral functional theories, conclusions at the behavioral functional level can be used to learn more about the informational functional level which can then inform us about the algorithmic level. It will be difficult to arrive at definite conclusions at the informational functional and (even more so) algorithmic levels, but we can make progress at those levels by

carefully considering how conclusions at these levels can be constrained by conclusions at other levels. As we shall see in the next section, these ideas have profound implications not only for social cognition research but for social psychology as a whole.

Implications for social psychological research

Benefits of focusing on the functional levels in addition to the algorithmic level.

Many social cognition researcher try to formulate theories at the algorithmic level. This is very appealing if one's goal is to understand the mental processes that underlie social behavior (as is the case for social cognition researchers). Not only does the algorithmic level offer an understanding of the mental mechanisms by which the input of a process is transformed into an output, it also offers insight into the function of the mental (sub-)process by allowing one to predict for any particular input into the algorithm what the output would be.

As is often the case, however, the best things in life are the hardest to get. Our analysis suggests that algorithmic theories have a low inductive penetrability. In other words, it is difficult to arrive at firm conclusions at this level on the basis of empirical data. Illustrations of this problem in social psychology are manifold. Just think of the seemingly unresolvable debates between dual-process and single-process models of social information processing (e.g., Deutsch & Strack, 2006; Kruglanski et al., 2007), competing accounts of dissociations between implicit and explicit attitude measures (e.g., Greenwald & Nosek, 2008), and representational versus constructive theories of attitudes (e.g., Fazio, 2007; Schwarz, 2007). Some have even wondered whether any substantial progress has been made at the algorithmic level of understanding mental processes (e.g., Garcia-Marques & Ferreira, 2011; Meiser, 2011).

Although we believe that social cognition researchers can continue to strive for progress at the algorithmic level, we also believe that it is important to simultaneously pay more attention to the behavioral functional level and informational functional levels. First, these levels are inductively more penetrable than the algorithmic level of analysis, which is why it should be easier to reach conclusions at these levels and thus to make genuine scientific progress. Second, focusing on the two functional levels will also allow one to consolidate progress that is made at the algorithmic level. In those cases where empirical data clearly contradict a specific algorithmic model, it is important to consider what these data imply at the informational and behavioral functional levels. Those data can be used not only to constrain the algorithmic level but also to make progress at the informational and/or behavioral functional levels. In this way, ideas at the algorithmic level can continue to fluctuate and evolve while at the same time a stable base of knowledge is formed at the informational and behavioral functional level. This not only increases our understanding of the function of mental (sub-)processes but also provides the best way to constrain future developments at the algorithmic level and thus make progress also at that level of analysis.

Perhaps the informational functional level provides the most fertile ground for social cognition research. As De Houwer and Moors (2013, p. 24) pointed out, “just like there is merit in understanding environment-behavior relations without understanding the (mental) mechanisms by which elements in the environment lead to certain behaviors, there is merit in understanding which elements in the environment lead to which mental representations or which mental representations lead to which behaviors without understanding the mechanisms by which an input leads to an output”. Importantly, this level combines some of the advantages of the other two levels. On the one hand, it has more deductive value than the behavioral functional level. That is, theories at the behavioral functional level can in principle

be derived from theories at the informational functional level whereas the reverse is not true. On the other hand, it is inductively more penetrable than the algorithmic level. Hence, the informational functional level can be a middle ground at which important and genuine progress can be made in understanding mental processes.

Different levels of analysis are mutually reinforcing provided that they are conceptually separated. Because different levels of analysis are interdependent, progress at one level can facilitate progress at another level. In other words, the different levels of analysis are mutually supportive. However, in order to be optimally supportive, the different levels need to be separated conceptually. It is particularly important that concepts at levels with low inductive penetrability (e.g., the algorithmic level) are not used in analyses at levels with high inductive penetrability (e.g., the functional levels). The lower the inductive penetrability, the more difficult it is to reach conclusions and thus the more likely it is that conclusions at that level will fluctuate while conclusions at other level remain constant. If theories at a level with high penetrability are formulated in terms of theories at a lower level of penetrability, then changes in the latter theories might jeopardize the former theories.

Consider the phenomenon of discounting in causal attribution (also see De Houwer & Moors, 2013). Discounting is sometimes also referred to as “blocking”, a term that stems from the idea that discounting it is due to a the lack of association formation for redundant cue-outcome relations (e.g., Fanselow, 1998). The use of such a term might discourage researchers from considering algorithmic models that explain discounting not in terms of a lack of association formation but in terms of a lack of information retrieval. Moreover, if discounting is defined as blocking (i.e., a lack of association formation), then support for other explanations of blocking would imply that “real” discounting (i.e., discounting that is due to a lack of association formation) might not exist.

On the relation between different approaches in social psychology. As noted in the introduction, the unifying aim of social psychology is to understand social behavior. Until now in this chapter, we have focused on the social cognition approach in social psychology which tries to understand the mental processes underlying this social behavior. There are, however, also other approaches to studying social behavior. These other approaches differ from the social cognition approach with regard to the kind of question that they ask about social behavior. In this section, we discuss two other approaches and examine how they are related to the social cognition approach. More specifically, we focus on the functional and the neuroscience approach in social psychology.

The functional approach in social psychology aims to understand social behavior by identifying the environmental causes of that behavior. One example of this approach is provided by research on "the power of the situation" in the 1960s and 1970s, which primarily focused on how subtle features of social situations can have powerful effects on social behavior. The most prominent examples include Milgram's (1963) studies on obedience, Zimbardo's Stanford prison experiment (Haney, Banks & Zimbardo, 1973), and Darley and Latané's (1968) research on bystander intervention. Most of the initial research on these forms of social influence was not directed at understanding the metal processes by which the environment influences behavior (although in some cases subsequent research did attempt to do so)." .

It is quite obvious that this functional approach in social psychology resonates well with the behavioral functional level of analysis in social cognition research. Both approaches try to uncover functional relations between environment and social behavior. In fact, the difference between the two merely concerns the goals to which the discovery of functional relations is a means. Whereas traditional functional researchers want to uncover functional

relations between environment and social behavior in order to understand the environmental determinants of social behavior, social cognition researchers who engage in behavioral functional analyses want to uncover those same functional relations in order to better understand the overarching mental processes by which the environment influences social behavior.

Now that we have made explicit the link between traditional functional research and the behavioral functional level of analysis within social cognition research, it also becomes clear how the functional and the social cognition approach in social psychology are related. Because the functional approach and the behavioral functional analysis in the social cognition approach use the same means for different goals (i.e., they both look for functional relations between environment and social behavior, be it for different reasons), they can directly contribute to each other. The only requirement is that one reaches agreement about which overarching mental processes underlie which kinds of environment-behavior relations. For instance, one could agree that the overarching mental process that drives obedience in Milgram's (1963) experiments is persuasion (in the sense of the overarching mental process by which persuasive messages influence behavior). In that way, by demonstrating how elements of the environment influence obedience behavior, traditional functional research also contributes to understanding of persuasion at the behavioral functional level. Vice versa, social cognitive research on persuasion as an overarching mental process can help increase our understanding of obedience behavior by uncovering new aspects of the environment that influence the output of the persuasion process.

Because of this overlap between the traditional functional approach and the behavioral functional analysis within the social cognition approach, both approaches can also interact in other ways. As we have discussed above, behavioral functional analyses constrain

analyses at the informational functional and the algorithmic level. The traditional functional approach can therefore contribute also to the informational functional and algorithmic level of the social cognition approach by expanding knowledge at the behavioral functional level. In turn, the informational functional and the algorithmic level can lead to new predictions and insights at the behavioral functional level. Testing these predictions can therefore also contribute to the traditional functional approach. In fact, by combining the traditional functional and the cognitive approach, a new functional-cognitive approach emerges that facilitates progress both at the functional and cognitive level (see De Houwer, 2011, De Houwer, Barnes-Holmes, & Moors, in press, and De Houwer, Gawronski, & Barnes-Holmes, submitted, for more details).

Another approach in social psychology that is becoming increasingly popular is the neuroscience approach. Social neuroscience aims to understand the neural processes that underlie social behavior. There are several ways in which it interacts with both the functional and social cognition approach in social psychology. From the perspective of a functional approach, the brain can be seen both as a part of the environment and as the locus at which neural behavior takes place. As such, within a functional approach, one can examine how properties of the brain (as part of the environment) influence behavior (e.g., by examining the effects of brain lesions on behavior) or how other parts of the environment influence the behavior of the brain (e.g., by examining the influence of stimuli on electrical activity or blood flow in the brain). This functional neuroscience approach actually exists under the name of behavioral (social) neuroscience (e.g., Carlson, 2010).

From the perspective of social cognition research, neuroscience research can help understand how mental processes are implemented in the brain. Hence, neuroscience can contribute to the implementational level of analysis as introduced by Marr (1982). As Marr

already pointed out, the algorithmic level cannot be reduced to the implementational level because the same algorithm can in principle be implemented in different physical systems. In our terms, there is a one-to-many relation going from the algorithmic level to the implementational level. De Houwer and Moors (2013) argued that there is also a one-to-many relation when going from theories at the implementational level to algorithmic theories. Without going into details, their argument is based on the fact that there can never be a one-to-one relation between aspects of the physical reality (such as neural processes) and aspects of the non-physical reality (such as the processing of information). What is most important for the present purposes, is that these ideas imply that the (social) neuroscience approach will never fully replace the (social) cognition approach. Both approaches can certainly inform each other but one cannot be reduced to the other (see Craver & Bechtel, 2007, and Chiesa, 1992, 1994, for excellent discussions of how different questions about a single object relate to each other).

Conclusion

Our scientific goals determine our research and the way in which we interact with our colleagues. We believe that making these goals explicit will promote scientific progress. In this chapter, we have argued not only that a scientific goal can specify an object and a specific question about that object, but also that a specific questions about a specific object can be studied at different levels. Hence, to promote scientific progress, it is good to also be aware of the level of analysis at which a scientific question is examined and to know how the different levels relate. We hope that our chapter not only raises awareness of these issues but also helps researchers to engage in these analyses of their research and the research of their peers.

Acknowledgements

Jan De Houwer and Agnes Moors, Ghent University, Belgium. Preparation of this paper was made possible by Methusalem Grant BOF09/01M00209 from Ghent University that was awarded to Jan De Houwer. We thank Bertram Gawronski and Galen Bodenhausen for their extensive comments on an earlier draft of this paper, especially with regard to how we could illustrate the relevance of our ideas for social psychology.

References

Bechtel, W. (2005). The challenge of characterizing operations in the mechanisms underlying behavior. *Journal of the Experimental Analysis of Behavior, 84*, 313-325.

Bechtel, W. (2008). *Mental mechanisms: Philosophical perspectives on cognitive neuroscience*. London: Routledge.

Carlson, N.R. (2010). *Foundations of behavioral neuroscience*. New York: Pearson.

Cheng, P.W. (1997). From covariation to causation: A causal power theory. *Psychological Review, 104*, 367–405.

Cheng, P.W., & Holyoak K.H. (2011). Causal attribution and inference as a rational process: The new synthesis. *Annual Review of Psychology, 62*, 135-163.

Cheng, P.W., & Novick, L.R. (1990). A probabilistic contrast model of causal induction. *Journal of Personality and Social Psychology, 58*, 545–567.

Chiesa, M. (1992). Radical behaviorism and scientific frameworks: From mechanistic to relational accounts. *American Psychologist, 47*, 1287-1299.

Chiesa, M. (1994). *Radical behaviorism: The philosophy and the science*. Boston, MA: Authors' Cooperative.

Craver, C. F., & Bechtel, W. (2007). Top-down causation without top-down causes. *Biology and Philosophy, 22*, 547-563.

Darley, J. M., & Latané, B. (1968). Bystander intervention in emergencies: Diffusion of responsibility. *Journal of Personality and Social Psychology, 8*, 377-383.

De Houwer, J. (2011). Why the cognitive approach in psychology would profit from a functional approach and vice versa. *Perspectives on Psychological Science, 6*, 202-209.

De Houwer, J., Barnes-Holmes, D., & Moors, A. (in press). What is learning? On the nature and merits of a functional definition of learning. *Psychonomic Bulletin & Review*.

De Houwer, J., & Beckers, T. (2002). A review of recent developments in research and theory on human contingency learning. *Quarterly Journal of Experimental Psychology, 55B,* 289-310.

De Houwer, J., Beckers, T., & Vandorpe, S. (2005). Evidence for the role of higher-order reasoning processes in cue competition and other learning phenomena. *Learning & Behavior, 33,* 239-249.

De Houwer, J., Gawronski, B., & Barnes-Holmes, D. (2013). A functional-cognitive framework for attitude research. Manuscript submitted for publication.

De Houwer, J. & Moors, A. (2013). Levels of analysis in psychology: Marr and beyond. Manuscript submitted for publication.

Deutsch, R., & Strack, F. (2006). Duality models in social psychology: From dual processes to interacting systems. *Psychological Inquiry, 17,* 166-172.

Fanselow, M. S. (1998). Pavlovian conditioning, negative feedback, and blocking: Mechanisms that regulate association formation, *Neuron, 20,* 625 - 627

Fazio, R. H. (2007). Attitudes as object-evaluation associations of varying strength. *Social Cognition, 25,* 603-637.

Garcia-Marques, L., & Ferreira, M. B. (2011). Friends and foes of theory construction in psychological science: Vague dichotomies, unified theories of cognition, and the nex experimentalism. *Perspectives on Psychological Science, 6,* 192-201.

Gardner, H. (1987). *The mind's new science: A history of the cognitive revolution.* New York: Basic Books.

Gilbert, D. T. & Malone, P. S. (1995). The correspondence bias. *Psychological Bulletin, 117,* 21-38.

Greenwald, A. G., & Nosek, B. A. (2008). Attitudinal dissociation: What does it

mean? In R. E. Petty, R. H. Fazio, & P. Brinol (Eds.), *Attitudes: Insights from the New Implicit Measures* (pp. 65-82) . Hillsdale, NJ: Erlbaum.

Haney, C., Banks, W. C., & Zimbardo, P. G. (1973). Interpersonal dynamics of a simulated prison. *International Journal of Criminology and Penology, 1*, 69-97.

Kelley, H. H. (1973). The processes of causal attribution. *American Psychologist, 28*, 107-128.

Kruglanski, A. W., Erb, H.-P., Pierro, A., Mannetti, L., & Chun, W. Y. (2006). On parametric continuities in the world of binary either ors. *Psychological Inquiry, 17*, 153-165.

Marr, D. (1982) Vision. San Francisco: W.H. Freeman.

McClamrock, R. (1990). Marr's Three Levels: A Re-Evaluation. *Minds and Machines, 1*, 185-196.

McClure, J. (1998). Discounting of causes of behavior: Are two reasons better than one*? Journal of Personality and Social Psychology*, *74*, 7–20.

Meiser, T. (2011). Much pain, little gain? Paradigm-specific models and methods in experimental psychology. *Perspectives on Psychological Science, 6*, 183-191.

Milgram, S. (1963). Behavioral study of obedience. *Journal of Abnormal Psychology, 67*, 371-378.

Ross, L. D. (1977). The intuitive psychologist and his shortcomings: Distortions in the attribution process. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 10, pp. 173-220). New York: Academic Press.

Shanks, D.R. (1995). *The psychology of associative learning*. Cambridge: Cambridge University Press.

Schwarz, N. (2007). Attitude construction: Evaluation in context. *Social Cognition, 25*, 638-656.

Schwarz, N., & Bohner, G. (2001). The construction of attitudes. In A. Tesser & N. Schwarz (Eds.), *Intrapersonal Processes* (pp. 436-457). Oxford, UK: Blackwell.

Van Overwalle, F., & Timmermans, B. (2005). Discounting and the role of the relation between causes. *European Journal of Social Psychology, 35*, 199-224.

Footnotes

1. One could argue that analyses at the behavioral functional level also go beyond the observed environment and observed behavior. This is true because the observer always perceives the world in a way that is shaped by his or her concepts and theories. However, drawing conclusions about mental representations requires an extra level of subjectivity and uncertainty because mental representations, as non-physical entities, can be inferred only by observing the physical environment and behavior.