
Confounding in studies on metacognition: a preliminary causal analysis framework

Borysław Paulewicz^{1,*}, Marta Siedlecka² and Marcin Koculak²

¹*Psychology Department, SWPS University, faculty in Katowice, Poland*

²*Consciousness Lab, Institute of Psychology, Jagiellonian University, Kraków, Poland*

Correspondence*:

Borysław Paulewicz

borys.paulewicz@gmail.com

ABSTRACT

By definition, metacognitive processes may monitor or regulate various stages of first-order processing. By combining formal causal analysis with hypotheses expressed by other authors we derive the theoretical and methodological consequences of this special relation between metacognition and the underlying processes. In particular, we prove that 1) without additional causal assumptions, neither metacognitive judgements (e.g., confidence ratings) nor correlations between performance (e.g., accuracy or sensitivity) and metacognitive judgements are unbiased measures of metacognitive monitoring or regulation; 2) that without additional causal assumptions, typical methods of controlling for first-order task performance (i.e., calibration, staircase, including first-order task performance in a regression analysis, or analyzing correct and incorrect trials separately) do not deconfound measures of metacognition; 3) that the first two problems cannot be solved by using simple models of decision-making derived from Signal Detection Theory. We conclude the paper by advocating robust methods of discovering properties of latent mechanisms.

Keywords: metacognition, causal inference, confounding, Structural Causal Model, meta-theory

1 INTRODUCTION

In this paper the term metacognition denotes cognitive processes that monitor other cognitive processes, as well as the results of such monitoring, including metacognitive regulation. This broad definition seems to be in agreement with what can be found in the majority of introductory chapters of various monographs on metacognition, of which there are now many (e.g., Beran et al., 2012; Dunlosky and Metcalfe, 2008; Koriat and Shitzer-Reichert, 2002; Chambres et al., 2002; Nelson and Narens, 1994). A monitored or regulated process is sometimes called a first-order process, an object-level process, a type 1 process, or a lower-level process. This naming convention reflects the hierarchical nature of the overall cognitive process responsible for performing tasks involving metacognition. Following this convention from now on, we will use the term "hierarchical task" to denote an arbitrary cognitive task that involves metacognitive monitoring or regulation of any kind.

One of the reasons that it is difficult to study metacognition is that it is a latent mechanism which has a dual causal role, i.e., it monitors and so is influenced by the underlying cognitive process, but it may also regulate and so influence the monitored process. To further complicate the matter, not every case of a first-order process influencing metacognition represents genuine metacognitive monitoring; for example, a first-order process could become more resource consuming, thus limiting the amount of resources available for metacognition. Similarly, not every case of metacognition influencing a first-order process is automatically a case of metacognitive regulation.

The aim of this paper is to use formal causal analysis to derive the theoretical and methodological consequences of this special relation between metacognition and the underlying processes. Even though this is a theoretical paper, we made sure that it does not contain any speculative claims: instead of providing our own hypotheses about how metacognition works, we combine formal causal analysis with the hypotheses expressed by other authors. In that sense we are proposing a meta-theoretical causal framework for studying hierarchical tasks.

The usefulness of our approach is illustrated by showing how it can help identify important limitations of certain widespread practices in studies on metacognition. In particular, we prove that 1) without additional causal assumptions, neither metacognitive judgements (e.g., confidence ratings) nor correlations between performance (e.g., accuracy or sensitivity) and metacognitive judgements are unbiased measures of metacognitive monitoring or regulation; 2) that without additional causal assumptions, typical methods of controlling for first-order task performance do not deconfound measures of metacognition; 3) that the first two problems cannot be solved by using simple models of decision-making derived from Signal Detection Theory. We conclude the paper by advocating robust methods of discovering properties of latent mechanisms.

Almost one-third of our paper is devoted to introducing elements of causal analysis. It is only after we describe the relevant formalism and its interpretation that we begin to address the issues directly related to metacognition. The reason for this is that we cannot assume that a researcher interested in studying metacognition will also be acquainted with formal causal inference, and we decided not to rely on the introductory books or papers on the subject since they contain much more information than we need to derive the main results. Note also that while we are fairly specific in our criticism of the way in which metacognition is often studied, the constructive part of our paper, in which we try to provide advice on how to do some things better, is rather generic and may not be directly applicable to any specific research problem in studies on metacognition. This is a consequence of the fact that the problems that we identify are general, but the solutions to these problems depend on the particular characteristics of each study.

1.1 Structure and interpretation of causal graphs

We will rely on Pearl's Structural Causal Model (Pearl et al., 2016; Pearl, 2000). At present there is only one alternative theory of causality with a similar scope, i.e., Rubin's potential outcomes framework (Rubin, 2005); however, since the two theories are equivalent in the sense that all the axioms of one theory can be derived from the axioms of the other (Galles and Pearl, 1998), the choice is only a matter of convenience. We have chosen SCM because in our opinion it is easier to work with thanks to its reliance on graphs. It is impossible to introduce all the major results of SCM in a single paper, so we will describe only the part that seems most readily applicable in typical scenarios when doing basic research.

We will be concerned with qualitative causal structure, i.e., with the issue of the mere presence or absence of causal connections between variables. The quantitative properties of causal relations, such as how to best describe the causal effect by some deterministic function or statistical model, will be considered only to illustrate a general point. The qualitative structure of causal relations will be represented by graphs consisting of nodes (i.e., variables) and arrows. An arrow from A to B will represent the assumption that A may be a direct cause of B . Here "direct" does not mean immediate, it only means that the effect is not mediated by any other variable in the graph. The causal process represented by an arrow may be arbitrarily complex and multi-staged, but it has to go in the direction of the arrow. In fact, every arrow or node can be thought of as a collapsed causal graph.

By drawing an arrow we express the assumption that this causal relation is a theoretical possibility, i.e., that there are no good reasons to assume that A does not cause B , but the graph is still valid even if some of its arrows do not correspond to any real causal processes, as long as no real arrows connecting the modelled variables are omitted. Moreover, the presence of an arrow from A to B does not mean that A is the only thing causing changes in B , and so whenever we draw a causal graph, unless we clearly state otherwise by saying that some variables are deterministic functions of other variables, we assume that every variable is also influenced by other unspecified factors which can be safely omitted from the graph, except for special cases. In particular, we have to include unspecified factors which may be common causes of variables represented in the graph. Finally, unlike in structural equation models, the effects may be non-linear, and when two or more arrows enter the same node, the joint influence may be interactive.

Because of the dual causal role of metacognition, there will be causal loops in some of our graphs. Given that causal processes take time, the loop could be taken to mean that causality can go back in time. That is not our intention. A causal loop, which can only connect time-aggregated variables, is shorthand for mutual influence occurring over time, as illustrated by Fig. 1 below:

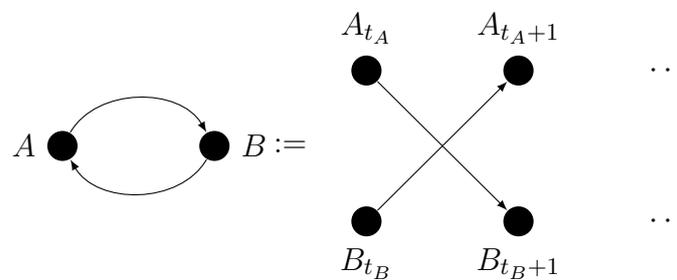


Figure 1. Interpretation of causal loops

Here t_i indexes discrete time.

A causal graph can be used to predict, interpret or explain the data because causal relations have statistical implications. If A causes B then A and B are probably statistically dependent¹. A causal path is a finite sequence of adjacent arrows that may change direction along the way, but the sequence must not contain repetitions. The three most important simple cases are the chain ($A \rightarrow B \rightarrow C$), the fork ($A \leftarrow B \rightarrow C$) and the collider ($A \rightarrow B \leftarrow C$). The first two paths imply that the outermost variables are correlated and the correlations due to the given path can be broken by physically fixing the value of the middle variable or by conditioning on the middle variable (e.g., by including the middle variable as a predictor in a regression analysis). The collider represents two independent parent variables that influence a common child; it behaves in an almost opposite way to the other two paths: even if A and C are independent, they may be *dependent given B*, and so regressing A on C will show no effect; however, regressing A on C and B simultaneously will show the effect of B and the spurious effect of C . This phenomenon is known as Berkson’s paradox (Berkson, 1946) and it shows that introducing additional variables in the regression analysis requires not only statistical but also causal considerations.

We will often make use of the following important fact: variables X and Y are correlated according to graph G iff there exists at least one collider-free path between X and Y in G . Equivalently, every collider-free path between X and Y provides a causal explanation for the observed correlation between X

¹ A may cause B and yet the two variables may be statistically independent, but this happens in special cases that can be ignored most of the time.

and Y . For brevity, we will call such paths *conductive*. Different conductive paths connecting the same two variables represent alternative but non-exclusive causal explanations of the correlation between the two variables. More than one conductive path between two variables may be true, in which case each path represents a partial explanation of the statistical dependence between these variables.

1.2 Identifying confounding paths using causal graphs

The causal graph representing a study may contain many variables and many arrows, but usually the researcher will be primarily interested in only a small subset of causal paths – often just a single arrow. Following Pearl 2000 we will use the term “target causal quantity” or “target causal effect” to denote any causal relation or causal property of interest. These are impossible to derive from statistical analysis alone because, just as statistical inference requires statistical assumptions, causal inference requires causal assumptions, even in experimental studies. One consequence of this is that, in general, causal questions cannot be answered just by showing that one statistical model fits the data better than another.

The importance of causal assumptions can be illustrated by using causal analysis to elaborate on the essential difference between experimental and observational studies. If X represents some experimental manipulation and Y represents the measured effect, then, given the *causal* assumption of random assignment, the only cause of X is a random device, so the only arrow that enters X is disconnected from everything else. It follows that there can be no conductive path between X and Y that enters X , hence any conductive path between X and Y has to leave X . Such a path cannot change direction, otherwise it would contain a collider and would not be conductive. Consequently, the observed statistical dependence between any randomly assigned X and any effect measure Y can only be explained by the causal process going from X to Y : when observed, this statistical relation (i.e., observed $p(Y|X)$) is an unbiased estimate of the total causal effect of X on Y .

Often more than one conductive path corresponds to the same correlation. This is especially true if the two variables are only observed, since in general either one of the two observed variables may cause the other, or the two variables may have common causes, as illustrated in Fig. 2 below.

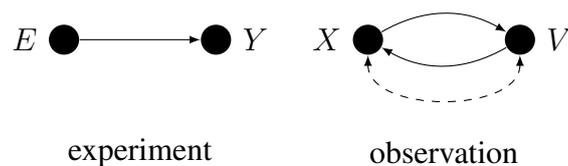


Figure 2. Conductive paths in experimental and observational studies

Here, E is randomly assigned, but X , Y and V are only observed. A dashed arc on the right is shorthand for $X \leftarrow U \rightarrow V$, where U represents all the unidentified common causes of X and V . When more than one conductive path corresponds to the given correlation, inferring causes from a correlation may require deconfounding, i.e., solving the problem of alternative causal explanations.

In SCM we usually talk about confounding paths, not variables. A path is confounding only with respect to some target causal quantity (e.g., $X \rightarrow V$) and its estimate (e.g., correlation between X and V), just as something is an alternative explanation only with respect to some other explanation and to something that is supposed to be explained. Justifying the preferred causal explanation for the obtained statistical results consists of neutralizing or arguing against the confounding paths (e.g., maybe V happens only after X has

already happened, so $X \leftarrow V$ can be safely deleted). The same goes for the justification of the preferred theoretical interpretation of the chosen measure, since validity of measurement is closely related to the relation of being caused by the subject of measurement (Borsboom et al., 2004).

If some plausible alternative causal explanations are not ruled out, i.e., if some confounding paths are not neutralized or broken, then the estimate of the target effect may be biased, i.e., the expected value of the estimate – what it actually measures – may be a mixture of the target causal effect and other confounding causal effects. In particular, even when the estimated statistical effect is different from zero, the contribution of the target causal effect to the estimate may be null, in which case the researcher will miss the target quantity entirely. Unlike noise or measurement error, bias resulting from the presence of confounding paths cannot be dealt with by increasing the sample size because it depends on what is being measured, not on how reliable the measurement is. Bias can only be dealt with – if at all – by changing the design, the method of analysis, or both.

Deconfounding is often crucial when doing basic research, especially when the study is concerned with discovering the latent mechanism, such as the mechanism of metacognition, since a mechanism is nothing more than a set of quantitative and qualitative causal relations. Of course, no study is perfect, but once the confounding paths are identified, they need to be addressed. As is commonly accepted in observational studies, the burden of proof is on the researcher, who omits certain arrows and thus dismisses alternative causal explanations.

There are several non-exclusive ways of dealing with confounding. One is by intervention, as in experimental design. However, despite their inherent strength, experimental studies rarely if ever provide definitive answers; this is partly because, especially in disciplines such as psychology, for many variables it is impossible to alter them directly, and the effects of interest may not be directly observable. For example, let the target quantity be the causal effect of short-term memory load (L) on the duration of the memory search (D). L is not directly accessible and D is not directly observable, and so the set size (S) is chosen at random on every trial as a way to indirectly determine the memory load, and recognition reaction time (RT) is used as an indirect measure of memory search duration. Thanks to random assignment, the correlation between S and RT is an unbiased estimate of the total causal effect of S on RT , but this is not the target quantity. The researcher hopes that the correlation between S and RT estimates the target quantity $S \rightarrow L \rightarrow D \rightarrow RT$. As an estimate of the target quantity, this correlation may be biased because without additional causal assumptions all that is guaranteed by random assignment of S is that the correlation of S and RT can only be explained by *some* unidirectional path from S to RT ; it does not guarantee that this is the path that the researcher has in mind. In fact if we assume that the latent effect variable (here D) cannot influence the latent cause variable (here L), there are exactly three kinds of confounding paths in an arbitrary experiment with indirect manipulation and measurement², as shown in Fig. 3 below:

² This can be proved: if S is randomly assigned, then the correlation between S and RT is an unbiased estimate of the unidirectional path from S to RT . Any unidirectional path from S to RT has to leave S and it either enters L or not. If it does enter L , it either enters D or not. From D it can only go to RT . If it does not enter D , it has to go to RT directly. If it does not enter L , it either enters D or not. If it does not enter D , it has to go to Y directly. Here we are assuming that L and D are not time-aggregated variables that can hide mutual influence, otherwise there are additional possible confounding paths.

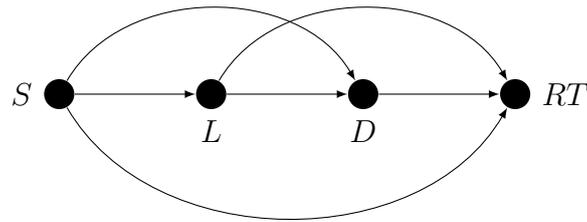


Figure 3. Confounding paths in an experiment with indirect manipulation and measurement

For example, the $S \rightarrow D$ path could represent the causal effect of perceived set size on memory search duration mediated by motivation or effort, but not mediated by memory load.

Another approach to deconfounding is by conditioning, i.e., by selecting observations or subjects with some property (e.g., only correct trials), or by introducing additional variables in the possibly non-linear regression analysis (see Pearl et al., 2016, for a more comprehensive treatment). For example, if X and Y are observed and Y cannot possibly influence X , then conditioning by regressing Y on X and all the common causes, if any, without introducing any spurious correlations, would neutralize all the confounding paths. In this case, even though the design is observational, given the causal assumptions it would be possible to estimate $X \rightarrow Y$ without bias.

By now it should already be clear why we take the arrow to mean that the causal effect is merely possible. All it takes for some causal path to provide a valid candidate explanation of the observed correlation is for the path to be theoretically possible and conducive. That is why the fewer arrows there are in the graph, the stronger the causal assumptions: there are fewer alternative explanations and more can be inferred from data about the generating process.

It follows that the more theoretically possible forms of monitoring or regulation there are, the harder it is to deconfound measures of metacognition in general. As we will now show, the relevant literature clearly indicates that it is more difficult to list cognitive operations that cannot possibly be monitored or regulated than it is to list ones that can be.

2 A CAUSAL ANALYSIS OF A GENERIC HIERARCHICAL TASK

In the case of metacognition, foundational questions are often answered using tasks in which the stimuli or their properties can be experimentally controlled and both the first-order (e.g., classification or free recall) and the second-order (some form of metacognitive judgement) responses are measured, sometimes simultaneously. Because of the random assignment of stimulus properties, as a first approximation we can represent such tasks as unidirectional causal paths going from the stimulus S to some first-order response R : this is the first-order process. Metacognition can be represented by a node connected causally with the nodes along the $S \rightarrow R$ path.

For our purposes it will be useful to divide the $S \rightarrow R$ path into six stages: stimulus S , stimulus encoding SE , decision- or evidence-accumulation process DP , the making of the decision D , motor response preparation MP , and first-order task response R . However, this subdivision into stages is in no way final and only serves as an example of how a preliminary sketch of a causal graph for a given study could look.

We will restrict our attention to monitoring or regulatory processes that operate during a task trial. If the data are not aggregated over trials, then the effects of any trial-level events on subsequent trials can often be safely ignored, which simplifies the graph considerably. For example, by performing statistical analysis on data that are not aggregated over trials, we can ignore the possibility of the alteration of decision criterion used on the following trial caused by the perceived distribution of the stimuli on previous trials, or the possible effect of confidence in a given trial on confidence in the following trials (Rahnev et al., 2015). This is because a path of the form $X_i \rightarrow Y_{i+1}$, where i is the trial number, cannot connect trial-level variables (i.e., variables with the same i value). As far as statistical relations within trial are concerned, such effects rightly belong to the omitted unspecified factors category, unless some trial level variable may influence two or more trial level variables on the following trial, thus forming a conductive path between them.

When deriving the causal graph for a generic hierarchical task, we will not assume anything about the stimulus or the response other than that the stimulus is randomly assigned. In this way, our causal model can be applied both to finite alternative forced choice tasks as well as to tasks where the space of valid responses is not clearly defined, such as learning tasks with a free recall stage.

In the following section we will provide a non-comprehensive list of theoretical and empirical arguments for introducing specific arrows in the causal graph that represents the overall causal process responsible for performing a generic hierarchical task. Note that the fact that we mention a study or a hypothesis does not necessarily mean that we agree with the interpretation of the results given by the authors; it only means that the authors expressed a hypothesis that has causal implications and should be taken into account when designing or interpreting the results of an experiment on metacognition.

Some variables that influence metacognition change asynchronously with task stimuli. For example, Samaha et al. (2017) observed that fluctuations in prestimulus alpha-band power are strongly negatively related to confidence ratings, although the relation to accuracy was not detected. Another example is the level of arousal, which fluctuates while a person is doing the task, and may influence cognitive performance on many different levels. Such factors can cause spurious correlations between confidence judgments in different trials, which can be interpreted, for example, as evidence of influence of confidence in a given trial on confidence in the following trial (see for example: Rahnev et al., 2015). However, because we restrict our attention to processes occurring within trial, here we will ignore such effects.

2.1 Some candidate metacognitive monitoring processes

Most theories of confidence assume that metacognitive assessments are informed by stimulus-related information, such as the quality of a perceptual item, its intensity or its size (e.g., Rhodes and Castel, 2008; Galvin et al., 2003; Higham et al., 2009; Vickers and Lee, 1998), although the degree to which stimuli-related evidence translates to metacognitive assessments varies between theories (see e.g. Kiani et al., 2014; Moran et al., 2015). In some studies it was found that confidence correlates with experimentally manipulated characteristics of the stimuli, such as presentation time (e.g. Lou et al., 2011), SOA (e.g. Del Cul et al., 2009) and motion coherence/strength (Kiani et al., 2014).

The first-order decision process that follows the stimulus-encoding stage may also be monitored by metacognitive processes. For example, an important class of hypotheses in metamemory studies concerns the relation between fluency or ease of processing and metacognitive judgments (Dunlosky and Metcalfe, 2008; Kelley and Lindsay, 1993; Koriat, 1997; Koriat and Ma'ayan, 2005). Moreover, all the models of choice confidence based on dynamic generalizations of Signal Detection Theory that we are aware of assume that confidence is a function of the history of evidence accumulation (sometimes referred to as

”random walk”), such as the drift rate (Ratcliff and Starns, 2009), the distance traversed by the decision-accumulation process scaled by the discriminability parameter (Link and Heath, 1975), post-decisional evidence accumulation (Pleskac and Busemeyer, 2010), or the ratio of smoothed sampled discriminial differences obtained when the response is made (Juslin and Olsson, 1997).

The very act of making a decision may also affect metacognition, for example, by reducing uncertainty (Busemeyer et al., 2006; Kvam et al., 2015; Wang and Busemeyer, 2016; Yearsley and Busemeyer, 2016). Stages of the process of translating the decision to the motor response may also be monitored. Motor response allows action monitoring and control and it seems implausible that the results of performance monitoring (e.g., a failure to execute an intended motor response) would not affect confidence judgments. It has been shown, for example, that response-locked error-related neural activity covaries with confidence level (Boldt and Yeung, 2015; Scheffers and Coles, 2000). A number of studies report correlations between confidence level and action characteristics, such as reaction time (Dougherty et al., 2005; Faivre et al., 2018; Fleming et al., 2015; Gajdos et al., 2019; Kelley and Lindsay, 1993; Kiani et al., 2014; Koriat and Ma’ayan, 2005; Siedlecka et al., 2019a; Wokke et al., 2019) or the presence of preparatory motor activity (Gajdos et al., 2019). Also, the model of self-evaluation proposed by Fleming and Daw (2017) assigns a crucial role to action by assuming that it provides information about one’s own decisional process that might not be accessible otherwise.

2.2 Some candidate metacognitive regulatory processes

It seems that the majority of studies on metacognition are concerned with monitoring, while metacognitive regulation is studied less frequently, especially in the foundational context. Many authors even omit the regulatory role when defining the term ”metacognition”, stating, for example, that it refers to the ability to monitor one’s cognitive processes or to knowledge about ongoing task performance (Fleming and Daw, 2017; Fleming and Dolan, 2012; Metcalfe and Shimamura, 1994; Siedlecka et al., 2016).

Metacognitive regulation during stimulus-encoding stages is probably ubiquitous, given the assumption that perception is an active process (for review see: Findlay and Gilchrist, 2001; Henderson, 2003, 2007; Stark and Ellis, 1981). The central idea in active perception theories (Bajcsy, 1988; Gibson, 1966) is that behaviors are selected based on the expected information content of the sensory data obtained by those behaviors, and expected information content can be thought of as a metacognitive property because it is relative to current knowledge and to the goals of an agent. A more trivial example of metacognitive regulation of stimulus encoding is the use of mnemonic techniques to improve future memory performance.

The generalizations of Signal Detection Theory provide theoretical arguments for the existence of a regulatory arrow from metacognition to the decision process as well as for the existence of an arrow that enters the stage of making of the decision. According to the common interpretation of the diffusion model (Ratcliff and McKoon, 2008), which is a dynamic generalization of the standard SDT model, the decision process is a kind of noisy evidence accumulation that starts from a possibly biased state and stops when accumulated evidence crosses a decision threshold. There are theoretical and empirical reasons to believe that both the initial bias and the decision thresholds can be metacognitively regulated (Ratcliff and McKoon, 2008). For example, studies on performance monitoring have shown that after encountering a difficulty (e.g., a conflicting stimuli) or after committing an error the subsequent response tends to be slower, which may be an indicator of engaging in a more cautious strategy (Dutilh et al., 2012; Gratton et al., 1992; Ullsperger and Von Cramon, 2001; Veen and Carter, 2006). A similar effect has been shown with real and false accuracy feedback: participants took longer to respond in a trial following negative feedback (Derryberry, 1991; Siedlecka et al., 2019b). Finally, Desender et al. (2019) have found that decision bounds

that regulate the speed-accuracy tradeoff in the diffusion model are related to the confidence judgement on the preceding trial.

Metacognitive regulation has also been studied in the context of learning. These studies indicate that the allocation of learning time or the selection of learning strategies may be guided by metacognitive monitoring and metacognitive knowledge. For example, feeling-of-knowing judgements positively correlate with the time spent on a question before giving up (e.g., Costermans et al., 1992; Gruneberg et al., 1977; Nelson et al., 1990; Reder, 1987, 1988). Judgements of Learning can be inversely related to self-paced study time, a result which may indicate that the time spent studying an item may depend on monitored or expected changes in the encoding strength (Mazzoni and Cornoldi, 1993; Mazzoni et al., 1990; Dunlosky and Connor, 1997; Thiede and Dunlosky, 1999).

Finally, correlations between motor response properties and confidence judgments found in many studies may also be interpreted as manifestations of the regulatory role of metacognition. For example, the positive correlation between confidence and reaction time may be at least partially explained by the hypothesis that when confidence in a decision is high there is little need to be cautious and the motor execution of the response can be relatively fast (see Gajdos et al., 2019, for a related result).

We are now in a position to draw, in Fig. 4 below, a partial causal graph representing the process of performing a generic hierarchical task. We want to stress that this graph is only meant as a simplified illustration of the problem of confounding in studies on metacognition. Note that the fact that there is only one node representing metacognition does not imply that there is only one metacognitive process or module, since, as we have already explained, every node may represent a collapsed causal graph.

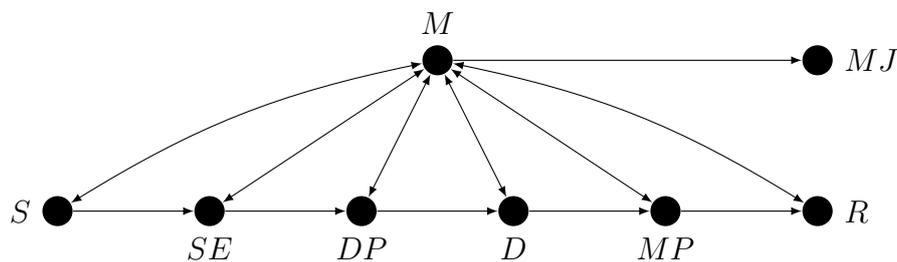


Figure 4. A partial causal graph representing a generic hierarchical task.

To improve the readability on this graph, causal loops are represented by bi-directional edges. Note that the $S \rightarrow M$ arrow represents the possibility of influence of S on M mediated by stimulus encoding processes unrelated to the first-order task. An immediate consequence of the presence of all the conductive paths in Fig. 4 is that without additional strong causal assumptions any single arrow or any path corresponding to a proper subset of conductive paths, such as all the paths involving only monitoring, is *not estimable*. To better illustrate the problems involved we will use the graph from Fig. 4 to identify potential confounding paths in studies on metacognitive monitoring.

3 CONFOUNDING IN STUDIES ON METACOGNITIVE MONITORING

The majority of studies on metacognition target metacognitive monitoring. The results of metacognitive monitoring, such as choice confidence, are only observed – they are not experimentally manipulated – and the sources of monitored information are also often not subject to experimental manipulation, at least not

directly. Because of this, any measure of metacognitive monitoring may be confounded in one way or another.

Imagine that a researcher was interested in metacognitive monitoring, or metacognitive "resolution", or "accuracy", but interpreted as a property of metacognitive monitoring. This researcher measured both accuracy and confidence and interpreted their correlation as a measure of metacognitive monitoring. This situation is so common in studies on metacognition that it deserves a graph, shown in Fig. 5 below.

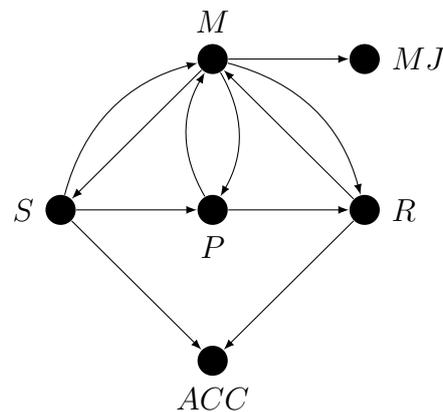


Figure 5. Confounding paths in a generic metacognitive monitoring study

Here S is the stimulus, P represents all the first-order processing stages between S and R , R is the response, M represents all the metacognitive processes, and MJ is some metacognitive judgement. The $S \rightarrow M$ arrow represents the possibility of monitoring the properties of the stimulus (e.g., if it looks familiar, if it is clearly visible, etc.), $P \rightarrow M$ represents the possibility of monitoring arbitrary stages of the first order decision process, and $R \rightarrow M$ represents the possibility of monitoring the first-order response (e.g., if it was quick or slow, or if there was an error in motor execution).

The arrows from S and R to ACC represent the fact that observed accuracy is a deterministic function of the stimulus and the response ($ACC = 1$ iff $S = R$). Note that nothing changes in the graph if an estimate of accuracy is replaced by an estimate of sensitivity (i.e., d') since it is also a deterministic function of S and R . The arrow from M to the first-order processing stages S , P and R represents the assumption that these stages can be metacognitively regulated. For example, $M \rightarrow S$ could represent metacognitively guided active perception, and $M \rightarrow R$ could represent metacognitively regulated response bias.

If the researcher was interested only in the correlation between MJ and ACC , then the graph would be irrelevant, but this is unlikely since this correlation by itself has no psychological meaning. If, however, the researcher interpreted this correlation as a measure of monitoring, e.g., of the amount of first-order information available for metacognition, then there are conductive paths that need to be addressed. The only way to justify the claim that this correlation represents metacognitive monitoring is to argue, based on the design of the study and the chosen method of statistical analysis, against the relevance of all the conductive paths between MJ and ACC that do not correspond to metacognitive monitoring. For example, both the $ACC \leftarrow S \rightarrow P_1 \rightarrow M_1 \rightarrow P_2 \rightarrow M_2 \cdots M_n \rightarrow MJ$ path, where n is the number of feedback loop iterations, and the $MJ \leftarrow M \rightarrow P \rightarrow R \rightarrow ACC$ path are conductive and connect MJ to ACC , but the former represents a metacognitive feedback loop involving not only monitoring but also regulation;

the latter has nothing to do with metacognitive monitoring, instead this path represents one way in which metacognitive regulation may contribute to the correlation between MJ and ACC .

Imagine also that the $MJ-ACC$ correlation was significantly different in the two experimental conditions, and this difference was interpreted as a measure of the influence of the experimental manipulation on metacognitive monitoring. We could represent this on the graph by introducing the experimental manipulation variable E that emits an arrow to M . The assumption that the amount of first-order information available for metacognition depends on E corresponds to the assumption that the joint influence of E and P (or S or R , since these are also stages of first-order processing) on M is interactive (e.g., $E(M) = a_0 + a_1P + a_2E + a_3PE$, $a_3 \neq 0$, assuming linearity). This would certainly explain the between-group difference in the $MJ-ACC$ correlation, but the researcher does not get to choose what is affected by E – nature does. If the researcher is interested only in a specific monitoring path such as $P \rightarrow M$ or an arrow from a specific stage of P to M , then this correlation is also confounded with other forms of metacognitive monitoring. Moreover, random assignment of E does not change the fact that, given the graph, as a measure of metacognitive monitoring the $MJ-ACC$ correlation is confounded with metacognitive regulation. Thus, the difference in the $MJ-ACC$ correlation could also be explained by the effect of experimental manipulation on metacognitive regulation, i.e., interactive effects $E \rightarrow P \leftarrow M$ or $E \rightarrow R \leftarrow M$ (but not $E \rightarrow S \leftarrow M$, since S was randomly assigned).

With some modifications, the graph from Fig. 5 can also be used to identify potential confounding paths in studies on the influence of heuristic cues such as fluency, response time, memorizing effort, or familiarity on metacognitive judgement formation. In many such studies the cues are not directly manipulated, although there are clear exceptions, such as font size, which when directly manipulated correlates with metacognitive judgement (Rhodes and Castel, 2008). The results of such studies are sometimes interpreted as evidence that changes in cues cause changes in metacognitive judgements by informing the monitoring process. However, when the cues are not directly manipulated, the correlation between the cues and metacognitive judgements does not necessarily measure the influence of the cues since there may be common causes of both. Finally, regardless of whether the cues are directly manipulated or not, the correlation between the cues and metacognitive judgments may reflect a complicated process involving iterations of a metacognitive feedback loop.

Common use of simple deconfounding strategies such as controlling for first-order task performance clearly indicates that researchers who study metacognition are well aware of the critical importance of deconfounding. However, as we will now demonstrate, the most popular simple deconfounding strategies fail to address this issue in its full generality.

4 WHY CONTROLLING FOR FIRST-ORDER TASK PERFORMANCE MAY NOT DECONFOUND MEASURES OF METACOGNITION

A popular approach to deconfounding measures of metacognition, or measures of effects of various manipulations on metacognition, is by attempting to make some chosen performance measure equal between the conditions, either by intervention, as in calibration or staircase, or by statistically controlling for the effect of first-order task performance.

The basic idea, which dates back at least to Nelson (1984), seems simple enough: if the experimental conditions differ in first-order task performance, then any differences in measures of metacognition can be attributed at least in part to the differences in first-order processing, which makes the latter a confound. If we force the performance measure to be equal in different conditions by calibration, or by using some

form of staircase procedure, or if we control for performance in statistical analysis, then any remaining differences in metacognition will be deconfounded from the effects of first-order task processing. We will now prove that this is not true without additional strong causal assumptions, and that it is in fact unlikely to be true in general.

We will only consider two popular ways of controlling for first-order task performance, namely calibration and including the performance estimate in a regression analysis, but with minor modifications our reasoning can be easily generalized to other cases.

Controlling for first-order task performance by calibration in metacognition studies usually consists of altering the stimuli in the preliminary stage of the experiment in such a way as to make the chosen performance measure more or less equal between the conditions. As long as the performance does not change during the experiment, calibration forces any observed differences in measures of metacognition to be unrelated to the calibrated performance measure.

Calibration certainly limits the set of possible causal paths between the stimulus and the response to those that correspond to the fixed performance score. However, this is a *quantitative* restriction: it changes how difficult the task is, but additional causal assumptions are necessary to infer that calibration makes any of the arrows in the graph from Fig. 4 disappear. In particular, this procedure is of no help to the researcher who claims that the observed effect is due to any specific arrow, such as metacognitive monitoring of a specific stage along the $S \rightarrow R$ path, because without additional causal assumptions all the other monitoring arrows may still be relevant.

More importantly, just because the first-order task *performance measure* was equalized between the conditions does not mean that first-order *processing* was equalized, nor does it mean that *only* first-order processing was affected by calibration. Without additional causal assumptions it is impossible to say if calibration affects only some first-order processing stages, or if it affects both some first-order stages and some metacognitive processes but does so in a way that makes the performance measure more or less equal between the conditions. In particular, calibration is a stimulus-level intervention and since stimulus properties can be metacognitively monitored, calibration may influence metacognitive monitoring. Because the main function of metacognitive monitoring is to guide metacognitive regulation, calibration may also influence metacognitive regulation. It follows, perhaps surprisingly, that without additional strong causal assumptions calibration deconfounds nothing. With or without calibration, the difficulty of deconfounding the effects observed for the chosen measure of metacognition remains essentially the same: the researcher either provides theoretical or empirical reasons to believe that all the confounding paths have been dealt with or the observed effects cannot be interpreted in terms of any specific causal path that connects metacognition to the first-order process.

Another important limitation of calibration is that there may be more than one performance measure for the given task, and different performance measures may not be interchangeable. For example, assuming that it is impossible for all participants to have the same decision bias in terms of SDT criterion placement, equalizing d' will not equalize average accuracy because participants with different biases will require different d' values to achieve the same accuracy. For the same reason, equalizing accuracy will not equalize d' .

Task performance can also be controlled for statistically. When there is uncertainty as to which performance measure is most relevant, the researcher can perform separate analyses, each time controlling for a different performance measure to see if the results hold. One way to statistically control for first-order task performance is by introducing the performance measure as a predictor in the regression analysis that

is aimed at estimating the metacognitive effects of interest. This method succeeds only if 1) it breaks all the confounding paths that are not dealt with by other means and 2) the first-order performance measure is not influenced by any stage along the target path. The second condition is necessary partly because conditioning on the descendant of a stage along the target path takes away some (or all, if the variable conditioned on *is* a stage along the target path) of the variance due to this path.

To see when conditioning on first-order performance results in deconfounding, consider a study in which some stimulus-level manipulation (S) is assumed to influence metacognition (M) by affecting some latent cause (C). This situation can be represented by a modified graph for an experiment with indirect manipulation and measurement shown in Fig. 6 below. For the sake of illustration, we assumed that task performance mediates the $S \rightarrow M$ confounding path, but the same line of reasoning would apply if it mediated any of the other two confounding paths.

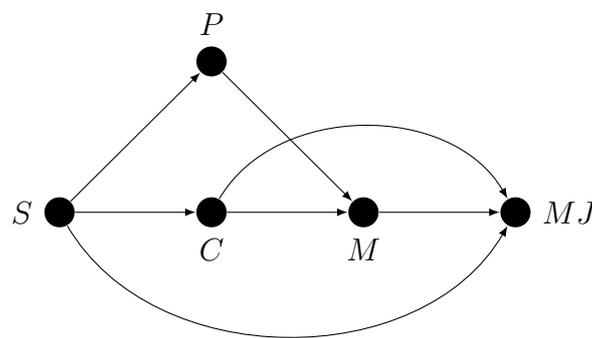


Figure 6. Deconfounding by conditioning on first-order task performance

Note that here we generously assume that P represents directly observable purely first-order performance. As we have explained before, indirect manipulation and measurement allows three types of possible confounding paths. Importantly, the confounding paths may be non-overlapping and it is impossible for the same variable (here P) to break (i.e., mediate) more than one distinct path. It immediately follows that without additional causal assumptions which imply that the other two confounding paths can be deleted, deconfounding by conditioning on performance is only partial, even when the chosen performance measure somehow estimates purely first-order processing without error.

Our reasoning can be generalized to statistical control of first-order task performance when it is built in simplified models of decision-making, such as models of metacognitive judgement based on Signal Detection Theory. In fact, one such model, called meta- d' (Maniscalco and Lau, 2012; Fleming and Lau, 2014), seems to have been explicitly designed to deconfound metacognition: according to its authors, one of the meta- d' parameters provides "a bias-free measure of metacognitive sensitivity" that is deconfounded from the effects of first-order performance and bias. We mention these models here to better illustrate the point, but a detailed analysis of their limitations deserves a separate paper. We should stress, however, that our critique applies only to the extent that these models are used to provide estimates of specific components of metacognition.

In every model of metacognitive judgement based on Signal Detection Theory that we are aware of, the process of arriving at a decision is represented either by an internal evidence sample, as in the meta- d' model, or, in the case of dynamic SDT-like models (e.g., Pleskac and Busemeyer, 2010; Ratcliff and

Starns, 2009; Juslin and Olsson, 1997; Link and Heath, 1975), by some form of an evidence-accumulation process. Importantly, in these models the cognitive mechanism is seen as an *abstract* evidence-sampling or evidence-accumulation process and the mechanism by which the evidence was obtained is not explicitly represented. This abstract internal evidence may just as well be the result of purely first-order processing or of an arbitrarily complex interplay of first-order and metacognitive processes. In effect, the problem of controlling for purely first-order processing that we have already encountered reappears: unless there are good reasons to believe that metacognition plays no regulatory role in the process of computing the internal evidence, the models of metacognitive judgement derived from Signal Detection Theory which do not account for the possible role of metacognition in the process of calculating the evidence are unable to deconfound metacognitive monitoring of first-order task information from metacognitive regulation, or to deconfound one form of monitoring from another.

5 WHY DECONFOUNDING METACOGNITION IS HARD AND WHAT CAN BE DONE ABOUT IT

As we hope we have already demonstrated, it is not easy to see when successful deconfounding of metacognition is achieved without formal causal analysis, even in the case of widely practiced, intuitively sound and seemingly straightforward control of first-order task performance. The limitations of performance equalization or of fitting simplified models of decision-making as methods of studying metacognition are a consequence of several properties that make metacognition a challenging subject of study: little is known about the mechanism of metacognition, therefore the researcher is forced to consider many possible arrows, which in turn may force the researcher to address many confounding paths. Moreover, these confounding paths can be particularly problematic because by definition metacognitive processes may be connected uni- or bi-directionally with many different stages of the first-order process.

5.1 Robust approaches to deconfounding

Ultimately, the limitations of all the approaches to deconfounding metacognition that we have analyzed so far are consequences of strong causal assumptions which are implicit in simplified models or in simple statistical corrections. Fortunately, there are several general-purpose approaches to deconfounding which can be used in studies on metacognition and which are robust in the sense that they do not require strong unsubstantiated causal assumptions. We describe these methods here because compared to a fully-fledged causal analysis targeted at a particular research problem and study design, they are relatively easy to apply. In a way all these methods revolve around the idea of deleting arrows or paths, which makes them well suited to the task of deconfounding measures of latent cognitive mechanisms. However, their robustness comes at a price: because these methods are not based on strong causal assumptions about the target latent mechanism, they may not allow for particularly strong causal conclusions.

5.1.1 Breaking confounding paths by design

Sometimes confounding paths can be guaranteed to be broken because of the design of the study. One example is studies on the effect of response order that some of us were involved in in the past (Siedlecka et al., 2019a, 2016). These studies were not aimed specifically at deconfounding, but this is irrelevant to the point that we are now making. The main manipulation was the order in which the metacognitive judgement and the first-order response were provided. Although this intervention does not break many confounding paths, when metacognitive ratings are provided first it certainly does break one path, i.e., from the first-order motor response to the metacognitive judgement. In principle, thanks to the simple

manipulation of order, such studies are well suited to eventually showing when metacognitive judgement can be influenced by motor response execution.

5.1.2 Deleting arrows by arguing for the null hypothesis

To the extent that it is possible that the arrows belonging to a confounding path are not real, it makes sense to try to demonstrate this empirically. Interestingly, demonstrating that some conductive path does not exist does not require an unbiased estimate of the path. To see why, imagine that a researcher was interested in the causal effect of X on Y , but the two variables were only observed and given what is known about them it was also possible that Y affects X . It follows that the correlation between X and Y is not an unbiased estimate of either arrow. However, if the researcher managed to demonstrate that X and Y are statistically independent, then the most likely explanation of this fact would be that neither arrow really exists. The downside is that obtaining evidence of statistical independence is not nearly as straightforward as obtaining evidence of an effect being different from zero.

We are aware of two ways of solving the problem of obtaining evidence for the null, but due to the lack of space we will only mention them briefly since this is not a paper on statistical analysis. One popular solution is to use Bayesian inference. The null hypothesis significance testing framework is ill suited to the task of arguing for the null hypothesis because a lack of statistical significance in no way indicates that the effect does not exist, it only means that it was not reliably detected. Moreover, in frequentist inference it is impossible to obtain a probabilistic statement about the null hypothesis, because in frequentist inference point hypotheses such as a null hypothesis are not points of some probability space, and so frequentist point hypotheses can only be true or false. In Bayesian inference a set of mutually exclusive and exhaustive hypotheses may form a probability space associated with a prior probability distribution and, once the data are obtained, a posterior probability distribution. A common approach to arguing for the null in Bayesian inference is by using the Bayes Factor in the form of the Savage-Dickey ratio (Wagenmakers et al., 2010). Another solution is to collect enough data points so that the resulting frequentist confidence intervals will be so narrow that if they contain zero it will make sense to say that the effect is either nonexistent or so small as to be negligible. The downside to the Bayes Factor is that it is sensitive to the choice of the prior distribution (Sinharay and Stern, 2002), while the downside to the frequentist approach is that it forces the researcher to justify the choice of the threshold below which the effect size can be considered to be negligible.

5.1.3 Identifying functionally distinct parts of the latent mechanism by selective modification

Arguing for the null is also an essential part of Sternberg's method of demonstrating separate modifiability by selective influence (Sternberg, 2001), which is a method of process and structure decomposition that has proved useful in the past (see Sternberg, 2001, for examples) and can be reconciled with Pearl's theory of causality. We are aware of five studies on metacognition or consciousness that were interpreted by the authors as demonstrating (partial) selective influence either on performance but not on metacognitive judgement or on metacognitive judgement but not on performance (Fleming et al., 2015; Rounis et al., 2010; Busey and Arici, 2009; Wilimzig et al., 2008; Lau and Passingham, 2006). We should point out that only one of these studies (i.e., Busey and Arici, 2009) contained a discussion of the inherent problems associated with arguing for the null hypothesis; it was also the only study in which the sample size was substantial. In every other case, the authors of the studies claimed – typically both in the title and in the abstract – that one of the effects (either on performance or on confidence) was zero based solely on the fact that it was not significant! Moreover, in some of these studies the effect on confidence was found to be

non-significant when conditioning on performance, which is problematic since, as we have already pointed out, first-order task performance may causally depend on metacognitive monitoring and regulation.

Given all of the above, it seems worthwhile to briefly introduce the method of separate modifiability. In its most basic form this method consists of finding two distinct randomly assigned factors, F and G , such that 1) given the hypothetical nature of the latent mechanism, F and G could potentially influence distinct aspects of the mechanism (e.g., stages, processes, or modules), and 2) the effects of each factor are demonstrably independent. The premise is that if there exist functionally distinct parts of a latent process or structure, then it may be possible to selectively influence them, which could be established if there were also distinct measures, each sensitive to one of the distinct parts. It is perhaps worth noting that none of the studies that we have just mentioned demonstrated that two factors selectively influenced two different measures in the same task.

The purpose of separate modifiability is to decompose a latent mechanism by providing information about the separate parts from which it could be composed. By itself, this method does not deconfound anything, but it may help in understanding the problem of confounding by providing information about the latent causal structure.

6 CONCLUDING REMARKS

In this paper we have demonstrated the limitations of common approaches to studying metacognition, including methods specifically aimed at deconfounding. Our analysis shows that detailed questions about metacognition are unlikely to be answered using simple statistical corrections such as conditioning on performance, or by fitting overly simplified mathematical models, such as various generalizations of Signal Detection Theory.

Because by definition metacognitive processes may be connected uni- or bi-directionally with arbitrary stages of first-order processing, confounding is a major problem and formal causal analysis may be required to correctly identify all the theoretically possible alternative causal explanations of the obtained statistical results, or to design a study that can potentially provide unbiased estimates of target causal quantities. The reader who believes that addressing all plausible alternative causal explanations is too tall an order should be reminded that it is exactly what is commonly required when observational studies are interpreted causally.

As the understanding of metacognition advances, some confounding paths may become irrelevant while new confounding paths may appear, thus making studies that once seemed valid look unconvincing, or vice versa. In fact, the theoretical analysis that we have presented in this paper led us to question what we thought our own past studies on metacognition indicated.

We believe that it is not unreasonable to expect that every study provides results which are valid given the explicitly stated assumptions. To this end we have advocated modesty when interpreting the data, arguing for the null, using selective influence and special designs that break confounding paths in order to better identify distinct parts of metacognition, and, if possible, replacing intuitive understanding of causality with formal analysis.

CONFLICT OF INTEREST STATEMENT

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

AUTHOR CONTRIBUTIONS

The main results were derived by BP with some help from MS, most of the text was written by BP with great help from MS, with the exception of sections describing the hypotheses about metacognitive monitoring or regulation expressed by other authors, which was written mostly by MS. MK provided valuable feedback and was partially responsible for finding the relevant literature.

FUNDING

Funded by Natinal Science Center, Poland, with Sonata grants given to BP (2013/09/D/HS6/02792) and MS (2017/26/D/HS6/ 00059).

REFERENCES

- Bajcsy, R. (1988). Active perception. *Proceedings of the IEEE* 76, 966–1005
- Beran, M. J., Perner, J., and Proust, J. (2012). *Foundations of metacognition* (Oxford University Press)
- Berkson, J. (1946). Limitations of the application of fourfold table analysis to hospital data. *Biometrics Bulletin* 2, 47–53
- Boldt, A. and Yeung, N. (2015). Shared neural markers of decision confidence and error detection. *Journal of Neuroscience* 35, 3478–3484
- Borsboom, D., Mellenbergh, G. J., and van Heerden, J. (2004). The concept of validity. *Psychological Review* 111, 1061
- Busemeyer, J. R., Wang, Z., and Townsend, J. T. (2006). Quantum dynamics of human decision-making. *Journal of Mathematical Psychology* 50, 220–241
- Busey, T. A. and Arici, A. (2009). On the role of individual items in recognition memory and metacognition: Challenges for Signal Detection Theory. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 35, 1123
- Chambres, P., Izaute, M., and Marescaux, P.-J. (2002). *Metacognition: Process, function, and use* (Springer Science & Business Media)
- Costermans, J., Lories, G., and Ansay, C. (1992). Confidence level and feeling of knowing in question answering: The weight of inferential processes. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 18, 142
- Del Cul, A., Dehaene, S., Reyes, P., Bravo, E., and Slachevsky, A. (2009). Causal role of prefrontal cortex in the threshold for access to consciousness. *Brain* 132, 2531–2540
- Derryberry, D. (1991). The immediate effects of positive and negative feedback signals. *Journal of Personality and Social Psychology* 61, 267
- Desender, K., Boldt, A., Verguts, T., and Donner, T. H. (2019). Confidence predicts speed-accuracy tradeoff for subsequent decisions in humans. *eLife* 8
- Dougherty, M. R., Scheck, P., Nelson, T. O., and Narens, L. (2005). Using the past to predict the future. *Memory & cognition* 33, 1096–1115
- Dunlosky, J. and Connor, L. T. (1997). Age differences in the allocation of study time account for age differences in memory performance. *Memory & Cognition* 25, 691–700
- Dunlosky, J. and Metcalfe, J. (2008). *Metacognition* (Sage Publications)
- Dutilh, G., Vandekerckhove, J., Forstmann, B. U., Keuleers, E., Brysbaert, M., and Wagenmakers, E.-J. (2012). Testing theories of post-error slowing. *Attention, Perception, & Psychophysics* 74, 454–465

- Faivre, N., Filevich, E., Solovey, G., Kühn, S., and Blanke, O. (2018). Behavioral, modeling, and electrophysiological evidence for supramodality in human metacognition. *Journal of Neuroscience* 38, 263–277
- Findlay, J. M. and Gilchrist, I. D. (2001). Visual attention: The active vision perspective. In *Vision and attention*, eds. M. Jenkins and L. Harris (Springer). 83–103
- Fleming, S. M. and Daw, N. D. (2017). Self-evaluation of decision-making: A general bayesian framework for metacognitive computation. *Psychological review* 124, 91
- Fleming, S. M. and Dolan, R. J. (2012). The neural basis of metacognitive ability. *Philosophical Transactions of the Royal Society B: Biological Sciences* 367, 1338–1349
- Fleming, S. M. and Lau, H. C. (2014). How to measure metacognition. *Frontiers in Human Neuroscience* 8, 443
- Fleming, S. M., Maniscalco, B., Ko, Y., Amendi, N., Ro, T., and Lau, H. (2015). Action-specific disruption of perceptual confidence. *Psychological science* 26, 89–98
- Gajdos, T., Fleming, S. M., Saez Garcia, M., Weindel, G., and Davranche, K. (2019). Revealing subthreshold motor contributions to perceptual confidence. *Neuroscience of consciousness* 2019, niz001
- Galles, D. and Pearl, J. (1998). An axiomatic characterization of causal counterfactuals. *Foundations of Science* 3, 151–182
- Galvin, S. J., Podd, J. V., Drga, V., and Whitmore, J. (2003). Type 2 tasks in the theory of signal detectability: Discrimination between correct and incorrect decisions. *Psychonomic Bulletin & Review* 10, 843–876
- Gibson, J. (1966). *The senses considered as perceptual systems* (Boston: Houghton Mifflin)
- Gratton, G., Coles, M. G., and Donchin, E. (1992). Optimizing the use of information: strategic control of activation of responses. *Journal of Experimental Psychology: General* 121, 480
- Gruneberg, M. M., Monks, J., and Sykes, R. N. (1977). Some methodological problems with feeling of knowing studies. *Acta Psychologica* 41, 365–371
- Henderson, J. M. (2003). Human gaze control during real-world scene perception. *Trends in cognitive sciences* 7, 498–504
- Henderson, J. M. (2007). Regarding scenes. *Current directions in psychological science* 16, 219–222
- Higham, P. A., Perfect, T. J., and Bruno, D. (2009). Investigating strength and frequency effects in recognition memory using type-2 signal detection theory. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 35, 57
- Juslin, P. and Olsson, H. (1997). Thurstonian and brunswikian origins of uncertainty in judgment: A sampling model of confidence in sensory discrimination. *Psychological Review* 104, 344
- Kelley, C. M. and Lindsay, D. S. (1993). Remembering mistaken for knowing: Ease of retrieval as a basis for confidence in answers to general knowledge questions. *Journal of Memory and Language* 32, 1–24
- Kiani, R., Corthell, L., and Shadlen, M. N. (2014). Choice certainty is informed by both evidence and decision time. *Neuron* 84, 1329–1342
- Koriat, A. (1997). Monitoring one's own knowledge during study: A cue-utilization approach to judgments of learning. *Journal of experimental psychology: General* 126, 349
- Koriat, A. and Ma'ayan, H. (2005). The effects of encoding fluency and retrieval fluency on judgments of learning. *Journal of memory and Language* 52, 478–492
- Koriat, A. and Shitzer-Reichert, R. (2002). Metacognitive judgments and their accuracy. In *Metacognition*, eds. J. Dunlosky and J. Metcalfe (Springer). 1–17

- Kvam, P. D., Pleskac, T. J., Yu, S., and Busemeyer, J. R. (2015). Interference effects of choice on confidence: Quantum characteristics of evidence accumulation. *Proceedings of the National Academy of Sciences* 112, 10645–10650
- Lau, H. C. and Passingham, R. E. (2006). Relative blindsight in normal observers and the neural correlate of visual consciousness. *Proceedings of the National Academy of Sciences* 103, 18763–18768
- Link, S. and Heath, R. (1975). A sequential theory of psychological discrimination. *Psychometrika* 40, 77–105
- Lou, H. C., Skewes, J. C., Thomsen, K. R., Overgaard, M., Lau, H. C., Mouridsen, K., et al. (2011). Dopaminergic stimulation enhances confidence and accuracy in seeing rapidly presented words. *Journal of Vision* 11, 15–15
- Maniscalco, B. and Lau, H. (2012). A signal detection theoretic approach for estimating metacognitive sensitivity from confidence ratings. *Consciousness & Cognition* 21, 422–430
- Mazzoni, G. and Cornoldi, C. (1993). Strategies in study time allocation: Why is study time sometimes not effective? *Journal of Experimental Psychology: General* 122, 47
- Mazzoni, G., Cornoldi, C., and Marchitelli, G. (1990). Do memorability ratings affect study-time allocation? *Memory & Cognition* 18, 196–204
- Metcalfe, J. and Shimamura, A. P. (1994). *Metacognition: Knowing about knowing* (MIT press)
- Moran, R., Teodorescu, A. R., and Usher, M. (2015). Post choice information integration as a causal determinant of confidence: Novel data and a computational account. *Cognitive Psychology* 78, 99–147
- Nelson, T. O. (1984). A comparison of current measures of the accuracy of feeling-of-knowing predictions. *Psychological Bulletin* 95, 109
- Nelson, T. O., Dunlosky, J., White, D. M., Steinberg, J., Townes, B. D., and Anderson, D. (1990). Cognition and metacognition at extreme altitudes on Mount Everest. *Journal of Experimental Psychology: General* 119, 367
- Nelson, T. O. and Narens, L. (1994). Why investigate metacognition. In *Metacognition: Knowing about knowing* (MIT press). 1–25
- Pearl, J. (2000). *Causality: Models, reasoning and inference* (Cambridge University Press)
- Pearl, J., Glymour, M., and Jewell, N. P. (2016). *Causal inference in statistics: A primer* (John Wiley & Sons)
- Pleskac, T. J. and Busemeyer, J. R. (2010). Two-stage dynamic signal detection: A theory of choice, decision time, and confidence. *Psychological Review* 117, 864
- Rahnev, D., Koizumi, A., McCurdy, L. Y., D'Esposito, M., and Lau, H. (2015). Confidence leak in perceptual decision making. *Psychological Science* 26, 1664–1680
- Ratcliff, R. and McKoon, G. (2008). The diffusion decision model: Theory and data for two-choice decision tasks. *Neural computation* 20, 873–922
- Ratcliff, R. and Starns, J. J. (2009). Modeling confidence and response time in recognition memory. *Psychological Review* 116, 59
- Reder, L. M. (1987). Strategy selection in question answering. *Cognitive Psychology* 19, 90–138
- Reder, L. M. (1988). Strategic control of retrieval strategies. In *Psychology of Learning and Motivation* (Elsevier), vol. 22. 227–259
- Rhodes, M. G. and Castel, A. D. (2008). Memory predictions are influenced by perceptual information: Evidence for metacognitive illusions. *Journal of Experimental Psychology: General* 137, 615
- Rounis, E., Maniscalco, B., Rothwell, J. C., Passingham, R. E., and Lau, H. (2010). Theta-burst transcranial magnetic stimulation to the prefrontal cortex impairs metacognitive visual awareness. *Cognitive Neuroscience* 1, 165–175

- Rubin, D. B. (2005). Causal inference using potential outcomes: Design, modeling, decisions. *Journal of the American Statistical Association* 100, 322–331
- Samaha, J., Iemi, L., and Postle, B. R. (2017). Prestimulus alpha-band power biases visual discrimination confidence, but not accuracy. *Consciousness & Cognition* 54, 47–55
- Scheffers, M. K. and Coles, M. G. (2000). Performance monitoring in a confusing world: Error-related brain activity, judgments of response accuracy, and types of errors. *Journal of Experimental Psychology: Human Perception and Performance* 26, 141
- Siedlecka, M., Hobot, J., Skora, Z., Paulewicz, B., Timmermans, B., and Wierzchoń, M. (2019a). Motor response influences perceptual awareness judgements. *Consciousness & Cognition* 75, 102804
- Siedlecka, M., Paulewicz, B., and Wierzchoń, M. (2016). But i was so sure! metacognitive judgments are less accurate given prospectively than retrospectively. *Frontiers in Psychology* 7, 218
- Siedlecka, M., Wereszczyński, M., Paulewicz, B., and Wierzchoń, M. (2019b). Visual awareness judgments are sensitive to the outcome of performance monitoring. *bioRxiv* , 572503
- Sinharay, S. and Stern, H. S. (2002). On the sensitivity of bayes factors to the prior distributions. *The American Statistician* 56, 196–201
- Stark, L. and Ellis, S. R. (1981). Scanpaths revisited: Cognitive models direct active looking. In *Eye movements: Cognition and visual perception* (Routledge). 193–226
- Sternberg, S. (2001). Separate modifiability, mental modules, and the use of pure and composite measures to reveal them. *Acta Psychologica* 106, 147–246
- Thiede, K. W. and Dunlosky, J. (1999). Toward a general model of self-regulated study: An analysis of selection of items for study and self-paced study time. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 25, 1024
- Ullsperger, M. and Von Cramon, D. Y. (2001). Subprocesses of performance monitoring: A dissociation of error processing and response competition revealed by event-related fmri and erps. *Neuroimage* 14, 1387–1401
- Veen, V. v. and Carter, C. S. (2006). Conflict and cognitive control in the brain. *Current Directions in Psychological Science* 15, 237–240
- Vickers, D. and Lee, M. D. (1998). Dynamic models of simple judgments: I. properties of a self-regulating accumulator module. *Nonlinear Dynamics, Psychology, and Life Sciences* 2, 169–194
- Wagenmakers, E.-J., Lodewyckx, T., Kuriyal, H., and Grasman, R. (2010). Bayesian hypothesis testing for psychologists: A tutorial on the Savage–Dickey method. *Cognitive Psychology* 60, 158–189
- Wang, Z. and Busemeyer, J. (2016). Comparing quantum versus Markov random walk models of judgements measured by rating scales. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 374, 20150098
- Wilimzig, C., Tsuchiya, N., Fahle, M., Einhäuser, W., and Koch, C. (2008). Spatial attention increases performance but not subjective confidence in a discrimination task. *Journal of Vision* 8, 7–7
- Wokke, M. E., Achoui, D., and Cleeremans, A. (2019). Action information contributes to metacognitive decision-making. *bioRxiv* , 657957
- Yearsley, J. M. and Busemeyer, J. R. (2016). Quantum cognition and decision theories: A tutorial. *Journal of Mathematical Psychology* 74, 99–116