LUNDS UNIVERSITET

INSTITUTIONEN FÖR PSYKOLOGI

An Assessment of the Role of Individual Differences and Variable Salience in Normative Causal Reasoning

Kalle Kallio Strand
Ossian Nörthen

Kandidatuppsats VT 2018

Handledare: Jean-Christophe Rohner

Abstract

Normative causal reasoning is essential for proper decision-making to be conducted, which is highly relevant in areas such as personnel assessment and predictions of work performance. However, research has found consistent violations of the Markov Assumption, a defining aspect of normative causal models that indicates which nodes in a causal network that are conditionally independent of one another. The purpose of the study was to investigate the role of Need for Cognition (NFC) in Markov violations, and whether violations were affected by variable salience. Participants ($N = 64$) made causal inferences in causal chain networks across eight vignettes and answered a translated measurement of NFC. Results of a four-way ANOVA revealed that although participants consistently violated the Markov Assumption, violations were unaffected by NFC scores and salience. A significant medium-sized interaction effect showed that participants committed more Markov violations when the mediating variable in causal chains was present as opposed to absent, suggesting that violations are influenced by certain situational factors. Lastly, no correlation was found between Markov violations and NFC, although conclusions are limited given that analyses indicated a restriction in range for NFC. Explaining individual differences remain an important aspect of research on Markov violations, and potential directions which future researchers can take are discussed.

Keywords: Causal reasoning, Markov violations, Need for Cognition, decision-making, salience

AN ASSESSMENT OF THE ROLE OF INDIVIDUAL DIFFERENCES AND VARIABLE SALIENCE IN NORMATIVE CAUSAL REASONING

Life provides plenty of situations every day in which judgments must be made under circumstances of uncertainty. Many times, such situations require causal reasoning. An organizational manager might want to infer why the productivity has gone down following a series of changes in the workplace, and politicians might be interested in the causes of recent fluctuations in the opinion polls. Surely the reader can generate a multitude of instances where causality is central to understanding the situation. Research in the last 25 years has found several areas in which causal judgement plays a defining role; among these are decision-making, categorization, and counterfactual reasoning (Rehder, 2018). More specifically, “causal graphical models serve as accounts of how people learn categories, classify objects, and generalize properties to other categories and category members. There are accounts of how people assess counterfactual possibilities, make decisions, and generate explanations.” (Rehder, 2018, p. 43).

Causality is predominantly modelled using *Bayesian networks*, acyclic graphs in which variables are represented as nodes and are interconnected using directed arrows that indicate causal relations (Glymour, 1998; Pearl, 1988). Bayesian networks lay the foundation for a normative framework, which acts as a guideline for how to correctly make causal inferences. Three possible instances of the Bayesian network are illustrated in Figure 1: *common cause networks*, in which one variable is the causal predecessor of two (or more) variables; *common effect networks*, in which one variable has multiple causes; and *causal chain networks*, in which three (or more) variables are successively interconnected.

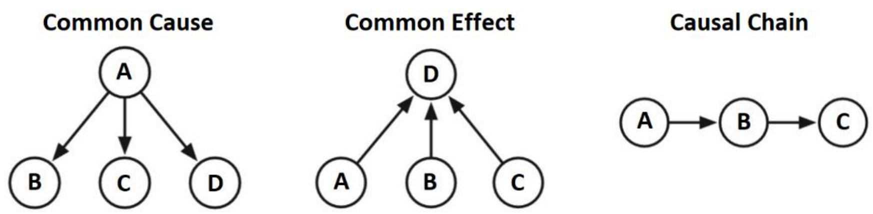

Figure 1. *Examples of causal models based on Bayesian networks.*

A defining principle of Bayesian networks is the *Markov Assumption* (also referred to as the *Markov condition*; Hausman & Woodward, 1999; Pearl, 1988, 2000), which states that a particular node in a causal network (e.g., causal chain networks) is conditionally independent of all nodes that are not direct causes or direct effects of that particular node. To illustrate, let us return to the causal chain network presented in Figure 1. The second variable (B) is directly connected to both its predecessor and its successor. However, as variables A and C are not directly connected, once the value of B is known, A and C are *conditionally independent* of one another; in other words, when B is present, variable A does not provide any useful information when inferring the state of C.

The current study will examine the tendency for people to violate the Markov Assumption and how this might be related to individual differences as well as variable salience. The remainder of this introductory chapter will present the historical perspectives and research within the subject, as well as introduce a base from which this paper's hypotheses are formed. The first section will feature the most important findings in research regarding violations of the Markov Assumption. Thereafter, an examination of individual differences in rational thinking dispositions will be provided, focusing on research on the Need for Cognition. Applications of research on this dispositional variable will demonstrate why the topic is interesting to further investigate in regard to Markov violations. Lastly, the chapter will present the purpose of the study as well as the hypotheses based on the perspectives presented.

Violations of Normative Models

The Markov Assumption can be described as a guideline for normative models, from which causal inferences are drawn. However, a considerable amount of research has revealed a tendency for people to deviate from normative models and thus violate the Markov Assumption (for an excellent review, see Rottman & Hastie, 2014). Rehder and Burnett (2005) investigated this tendency in a series of experiments, in which participants were taught causal networks in different domains (e.g., biology, chemistry, astronomy) based on the common cause network, the common effect network, and the causal chain network. They found that participants violated the Markov Assumption across all domains. To further ensure that violations were not the result of prior domain knowledge, an additional experiment was conducted in which no category information was provided. Participants were informed of an imaginary domain, *Daxes*, in which four variables (simply labelled *A*, *B*, *C*, and *D*) formed a common cause network; because no other information was presented, Rehder and Burnett argued that inferences would be exclusively based on the information provided in the causal network. However, they still found violations of the Markov Assumption.

In face of these findings, Rehder and Burnett (2005) suggested that participants relied on *domain-general underlying mechanisms* when making inferences, without necessarily knowing the nature of those mechanisms. For instance, participants could have included an additional variable (labelled *E*) to their mental representation of the causal model provided, that was directly connected to all nodes in the network. That way, it can be argued that participants still adhered to the Markov Assumption and followed normative models, because the participants' own causal representations of the models featured additional hidden knowledge structures not provided by the experimenters. However, it is unclear why people would assume such underlying mechanisms, and what kind of mechanism could flexibly account for all these violations.

Mayrhofer, Hagmayer, and Waldmann (2010) studied Markov violations in two experiments using a paradigm in which aliens communicated through telepathy. In the first experiment, four aliens formed a common cause network, in which one alien named Gonz (labelled as the cause) either sent its thoughts to the remaining three aliens (*sending condition*), or had its thoughts read by the three aliens (*reading condition*). Gonz could either think about "por" or "tus," and the aliens that either received or read Gonz's thoughts were more likely to think of whichever word Gonz had been thinking about (e.g., if Gonz thought about "por," the other aliens were consequently more likely to think about "por"). Participants were presented with specific situations in which Gonz's and two other aliens' thoughts were known, and were subsequently asked to rate the likelihood that the fourth alien would think about "por" given the specific circumstances. The findings in both conditions showed a nonindependence effect in which participants not only relied on Gonz's thoughts when making inferences, but also incorrectly based their answers on the thoughts of the two other aliens. The same tendency was found in a second experiment, in which the aliens were instead arranged in a causal chain network.

Mayrhofer et al. (2010) argued that, although most of their findings showed violations of the Markov Assumption, it was logical for participants in the sending condition of the first experiment to show a nonindependence effect. Here, Gonz can be viewed as an active agent and thus responsible for whether the three aliens were influenced by Gonz's thought ("por" or "tus"); for example, if Gonz's telepathic abilities were not functioning properly, one could reason that all aliens would be equally affected and thus conclude that the thoughts held by the three aliens were dependent of each other. Therefore, knowing the thoughts of two of the recipient aliens would provide useful information when asked to infer the state of the target alien. Stated differently, participants might have relied on an underlying mechanism (i.e., whether Gonz's telepathic abilities functioned correctly) when making inferences, similar to the findings of Rehder and Burnett (2005).

Rehder (2014) conducted several experiments to further investigate Markov violations. Participants were taught the three causal models in Figure 1, made up with sets of three binary variables from domains of economics, meteorology, and sociology, as well as information-less domains. They were then presented with two scenarios in which two of the variables took on specific values (present, absent, unknown), and were subsequently asked to determine in which scenario, if any, the remaining variable was most likely to be present. Results showed a tendency to violate the Markov Assumption that was unaffected by variations in domains, regardless of whether participants were forced to rely on fast and intuitive reasoning or if they could reason carefully and deliberately.

Knowing that previous studies (e.g., Rehder & Burnett, 2005; Mayrhofer et al., 2010) had attributed similar findings to the use of underlying knowledge structures (also referred to as *elaborated causal models*), Rehder (2014) wanted to test the extent to which these elaborated causal models could explain the Markov violations committed by participants in his experiments. He found that neither of the hypothesized models could provide a comprehensive account for the Markov violations. Essentially, what this means is that while most of the elaborated causal models could explain inferences in some situations (e.g., in common cause networks), the investigated models were considered insufficient to account for all results obtained in Rehder's experiments. In the absence of an elaborated causal model that flexibly and comprehensively accounts for causal inferences in general, Markov violations continue to hold ground as deviations from normative causal reasoning.

As indicated above, a considerable body of research suggests that violating the Markov Assumption is a general tendency that a) persists through variations of domain information and other manipulations, and b) is difficult to account for with several theories of causal representation (such as the underlying mechanism model; Rehder & Burnett, 2005). Interestingly, however, Rehder (2014) found that approximately 60% of all Markov violations in his experiments were committed by a quarter of the participants, which strongly suggests that individual differences play an important role in bringing forth this tendency. In other words, the results indicate that there might exist dispositional variables that can explain why some people violate the Markov Assumption while others do not. Investigating the role of individual differences in Markov violations is highly relevant because of the wide array of situations that revolve around causal reasoning. Thus, establishing a relation between Markov violations and individual differences would mean that measurements of these individual differences become important instruments for predicting how an individual will make causal inferences as well as the extent to which this individual will adhere to normative models. Arguably the most obvious application is how this may affect work and assessment of personnel; organizations would be able to use measurements of individual differences in recruitment processes to assess job suitability for positions that require objective and normative causal inferences to be made (e.g., program evaluation). To the extent of our knowledge, no previous study has attempted to investigate the role of individual differences in Markov violations, nor have any speculations been published that concern what individual differences might be relevant in this remark. For this reason, the following section seeks to further address this topic.

Individual Differences, Rational Thinking, and Need for Cognition

While the results of Rehder (2014) are novel in research on Markov violations, research on various cognitive biases has also found that there is a difference in the degree to which people adhere to normative models (e.g., Slovic & Tversky, 1974; Stanovich & West, 1998, 2008; Tversky & Kahneman, 1983; see also Stanovich, 2016). Stanovich (2012) brings forth an interesting discussion about individual differences in rational thinking dispositions, such as the tendency to seek out various perspectives and gather a body of information before reaching conclusions. He argues that these thinking dispositions comprise a rational thought pattern that should explain why some people deviate from normative models while others do not. According to this view, participants that failed to reason normatively in the above-mentioned studies of cognitive biases (Slovic & Tversky, 1974; Stanovich & West, 1998, 2008; Tversky & Kahneman, 1983) would be deemed irrational.

Given this apparent relation between cognitive biases and irrational thinking, it is likewise reasonable to hypothesize that a similar relation might exist between irrational thinking and Markov violations. Thus, a measurement related to rationality might be able to differentiate between people who commit Markov violations and those who do not; one such dispositional variable that is widely acknowledged in psychological research is the *Need for Cognition* (henceforth referred to as NFC; Cacioppo & Petty, 1982; see also Cacioppo, Petty, Feinstein, & Jarvis, 1996). NFC can be defined as the degree to which a person voluntarily engages in cognitive challenges, and enjoys thinking carefully and deliberately about topics regardless of their importance. Leary, Shepperd, McNeil, Jenkins, and Barnes (1986) found a correlation between NFC and objective decision-making, which they define as rational judgments based on empirical evidence. Indeed, several studies have found that people high in NFC tend to perform better at various decision-making tasks (Carnevale, Inbar, & Lerner, 2011; Smith & Levin, 1996; Stanovich & West, 1999). It has also been argued that, because NFC measures an inclination to engage in intellectual activities and deep-level reasoning, it constitutes a relevant measure of rational thinking (Epstein, Pacini, Denes-Raj, & Heier, 1996). For this reason, Epstein et al. (1996) uses a modified measure of NFC as the rationality subscale of their Rational-Experiential Inventory (REI). Taken together with the extensive testing of NFC as a dispositional variable (Cacioppo et al., 1996), we argue that there is enough evidence to warrant the use of NFC in an investigation of individual differences related to Markov violations.

Applications of Need for Cognition in Business and Everyday Life

The applications of individual differences in relation to normative causal reasoning are many. Burke and Litwin (1992) provides a causal model of how transformational and transactional factors affect organizational motivation and performance. The model is a complex network of different variables that are causally related. If we are to understand how this may be applicable, it may also be important to understand how people resonate in their causal judgement. Previous research has found a tendency for people to violate the Markov Assumption, as well as finding some people more prone to do so than others (Rehder, 2014). This paper will aim to find out if NFC could be one individual difference that may affect this tendency. The aspect of individual differences in Markov violations is relevant to explore because of its application for measurement of rational judgement and decision-making. According to Highhouse (2001), research in judgement and decision-making has increased the past decades. The research has been embraced in a wide range of areas including, but not limited to, consumer psychology, medicine, and law. However, the research has not rooted itself as deeply into industrial and organizational psychology. To understand which people might draw the wrong conclusions, research has to be done on what trait might have an effect on rational judgement, which in this case refers to a person’s adherence to, or deviations from, normative models. Further connection of individual dispositions to rational judgement and decision-making might open up the possibility for more precise measurement of performance, for example in recruitment for businesses where causal judgement and decision-making is highly prioritized.

Additionally, research has explored the connection between NFC and various dispositional variables that can be beneficial in aspects of work as well as everyday life. A study carried out by Haugtvedt and Petty (1992) explored how NFC can impact judgement. Two groups, one consisting of those with high NFC and one with those of low NFC, both received a message that argued that a food additive was unsafe. Secondly, the groups received a countermessage arguing that the food additive was not unsafe. The results showed that there was no significant difference between the groups after the first message. The second message, on the other hand, generated a significant difference. The high NFC group showed much resistance to the countermessage, whereas it had a significant impact on the low NFC group. The researchers argued that this difference can be explained with the fact that people with high NFC tend to recall more of the initial message and therefore elaborate on them; this argument was further supported with an additional test for recall and thought measure, where in fact people with high NFC recalled more information on average.

Another study conducted by Kearney, Gebert, and Voelpel (2009) aimed to research when and how diversity benefits teams for better team performance. They sampled 83 different teams and measured a wide variety of variables, such as age and educational diversity. They found that diversity in both age and education had a positive impact on both collective team identification and team performance, but only when the average NFC in the team was high. By contrast, high NFC tended to have less effect in homogeneous groups, which, according to this study, could be explained by the fact that those with high NFC tend to discuss in length what they have in common rather than what they do not. This, as opposed to heterogeneous groups with a high NFC average, leads to a counterproductive process. In other words, diversity was more likely to be beneficial for teams in certain aspects when the team on average had high NFC. It was not, however, beneficial for homogeneous groups with a high NFC average.

Furthermore, Smith, Kerr, Marcus, and Stasson (2001) researched whether NFC correlated with social loafing. They sampled 160 students and applied a 2 × 2 × 2 design with groupings of NFC (high vs. low), task description (engaging vs. not engaging), and different work conditions (coactive vs. collective). The coactive work condition meant the subjects learned that the individual work performance would be measured, while the collective condition meant they learned the group performance would be measured. The task description showed no significant effect. However, the results revealed that those with high NFC performed as well under the coactive condition as the collective (if ever so slightly better in the collective condition). Those with low NFC performed significantly worse in the collective condition than the coactive, which suggests that people with low NFC are more prone to social loafing.

As presented above, research has shown that those with high NFC tend to perform better in several areas, such as judgement, cooperation in diverse groups, and team performance. Research also suggests a connection between NFC and rationality. Finally, rational thinking can, as previously shown, explain why some people deviate from normative models, which also might apply to Markov violations. In this regard, we find it interesting to investigate how NFC might be one of the individual differences that affect rational thinking measured in Markov violations.

Purpose of the Study

In light of previous studies on normative reasoning (e.g., Stanovich & West, 1998, 2008), violations of the Markov Assumption can be described as a cognitive bias in that people erroneously deviate from normative models when making causal inferences. Although extensive research has been carried out testing the relation between rational thinking and several cognitive biases (see Stanovich, 2016; Stanovich, West, & Toplak, 2016), to the extent of our knowledge, no study has attempted to map which individual differences that might be relevant in Markov violations. The main purpose of the current study is to take the first step towards solving this issue, by investigating whether NFC can explain people's tendencies to violate the Markov Assumption in causal chain networks. Our first hypothesis posits that, in line with previous studies (e.g., Mayrhofer et al., 2010; Rehder, 2014; Rehder & Burnett, 2005), participants in our study will show a general tendency to violate the Markov Assumption. To illustrate using the causal chain network in Figure 1, this prediction means that, when the value of variable B is known, participants in general are expected to incorrectly infer that variable C is more likely to be present if variable A is also present (compared to absent), even though the two variables are conditionally independent. However, we also hypothesize that participants high in NFC will make fewer violations compared to participants low in NFC. Because of the apparent relation between NFC and rationality, it is reasonable to believe that participants high in NFC will show a stronger adherence to normative models in our study.

Additionally, although previous studies have found violations of the Markov Assumption in various domains and with abstract categories (e.g., Rehder, 2014; see also Rottman & Hastie, 2014), we found it interesting to explore the role of *salience* on participants' tendencies to violate the Markov Assumption. Highly salient factors in a causal model should per definition be more prominent in the network, and this increase in noticeability might cause participants to incorrectly assume dependence in a situation where the Markov Assumption screens off two variables from each other and makes them conditionally independent. To the extent of our knowledge, no study has attempted to investigate variable properties on this particular basis. Within the term salience we chose to further distinguish between *overt* stimuli, defined as observable and shown in an obvious way, and *covert* stimuli, defined as hidden and non-observable. The secondary purpose of the study is therefore to investigate whether overt and covert variables in causal chain networks differ in the degree to which they invoke Markov violations. We argue that overt stimuli, given their highly salient nature, should cause more violations of the Markov Assumption than variables that are covert.

Lastly, it is possible that an interaction effect might exist between NFC and salience; for example, overt stimuli might evoke more Markov violations exclusively for participants low in NFC. While there is indeed a possibility, we proceed with caution when making assumptions given that the inclusion of salience in the research design primarily is exploratory and thus not grounded in previous research.

Method

Participants

Sixty-four participants (54.7% female, 42.2% male, 3.1% not specified) were recruited on Lund University campus grounds using a convenience sampling method. Age varied between 19 and 28 years (M = 22.78, SD = 1.88). The vast majority of the sample were students (98.4%). Participants were not given any compensation and were thus considered volunteers.

Materials

The survey consisted of a written questionnaire in which eight vignettes were presented followed by a measurement of NFC.

Vignettes. Eight vignettes were created to measure Markov violations. Each vignette featured a causal model in which three binary variables formed a causal chain network (as seen in Figure 1). Thus, participants learned that A caused B, and that B caused C. To investigate whether the likelihood of Markov violations would vary based on variable salience, in half of the vignettes the variable A was overt (e.g., temperature), and in the other half covert (e.g., intelligence).

Each causal model was accompanied with one specific scenario where the variables A and B took on specific values (present, absent). In total, four combinations of the values of the variables were possible, resulting in four different scenarios. In each vignette, the participants were asked to rate the likelihood that C would be present given the specific values of A and B in the scenario. Table 1 illustrates how the information provided in a vignette was presented differently across the four scenarios (for a complete list of all vignettes used, see Appendix).

To ascertain that participants could not base their responses on their previous knowledge, all vignettes were designed to concern causal chain networks about extraterrestrial beings on a previously unknown planet; additionally, the instructions emphasized that our laws of physics were not applicable to the alien planet and its residents, further reminding the participants to solely base their responses on the information provided in each vignette.

Table 1. *Example of information provided across the four different conditions of a vignette.*

Causal model
The following is true for aliens in general: • High intelligence causes high physical strength. • High physical strength causes high-pitched voice.

	Condition[a]	Scenario information about A Value and B Value
1.	A = Present, B = Present	*A certain alien has high intelligence.* *This alien has high physical strength.*
2.	A = Absent, B = Present	*A certain alien does* ***not*** *have high intelligence.* *This alien has high physical strength.*
3.	A = Present, B = Absent	*A certain alien has high intelligence.* *This alien does* ***not*** *have high physical strength.*
4.	A = Absent, B = Absent	*A certain alien does* ***not*** *have high intelligence.* *This alien does* ***not*** *have high physical strength.*

[a] Condition description is only featured for clarity; in the actual vignettes, only the causal models and the scenarios were presented.

Need for Cognition. In order to measure NFC, the 18-item Need for Cognition Scale (NCS; Cacioppo, Petty, & Kao, 1984) was used. Since there is currently no validated Swedish version of the scale, the original scale was translated from English to Swedish; this was accomplished through a *back-translation*, in which the 18-item NCS was translated to Swedish and then back into English to ensure that the scale was accurate enough to be comparable with the original scale.

Design

The study uses a 2 × 2 × 2 × 2 mixed between-within subject design. The between-subject independent variable is NFC (2: low, high), which is measured by the translated 18-item NCS and transformed into a binary variable using a mean split.[1] The within-subject variables are the following: salience (2: overt, covert), A Value (2: present, absent), and B Value (2: present, absent). The dependent variable is the estimated likelihood rating that C is present, measured using a 7-point Likert scale ranging from “Not at all likely” (1) to “Very likely” (7).

Additionally, two between-subject counterbalancing factors were used. First, half of the participants were presented with the four overt vignettes followed by the four covert vignettes, and vice versa for the other half of the participants. Second, the scenario conditions in Table 1 were counterbalanced by designing two 4 × 4 Latin Squares – one for the overt vignettes, and one for the covert vignettes. Thus, each vignette appeared equally often in all four scenario conditions. In total, this counterbalancing method resulted in eight versions of the written survey, each consisting of eight vignettes: two of each scenario condition, distributed evenly among the two salience values. Participants were assigned in equal numbers to the between-subject counterbalancing conditions.

Procedure

Participants first read the instructions page in which the causal chain network was introduced. The instructions emphasized that all causal models presented in the vignettes would follow the presented causal chain network. Additionally, participants were informed that the causal models were general principles, whereas the scenarios were special cases with specific

[1] Note that the division into high and low Need for Cognition is based on relations between scores; thus, a participant in the low NFC group does not necessarily have an objectively low score on the NCS, and vice versa.

circumstances; this description was meant to stress the fact that scenarios did not always adhere to the causal models (for example, B could be present although A was absent).

After reading the instructions page, in each of the eight vignettes the participants were taught a specific causal model followed by a scenario in which the values of A and B were specified. For each vignette they were asked to rate the likelihood that C would be present on a 7-point Likert scale. Lastly, the participants filled in the translated 18-item NCS. The study was estimated to take 5-10 minutes to complete.

Ethics

Participants were informed of the purpose of the study, which was to investigate how people reason about causal models in different situations. Additionally, they were informed of their right to withdraw from the study at any time, that participating in the study was voluntary, and that all participants remain anonymous throughout the study. In the instructions of the written survey, a contact email-address was added in case anyone had questions regarding the study. The only personal information collected was age, sex, and occupation. Participating in the study did not involve more than minimal risk.

Results

An index for the 18-item NCS was obtained by calculating the mean rating for the 18 items. Two participants failed to complete the 18-item NCS, and were excluded from the sample; thus, analyses were conducted with 62 of the original 64 participants.

Results revealed a high internal reliability of the translated 18-item NCS (α = .84). Initial analyses showed that likelihood ratings were unaffected by the two between-subject counterbalancing factors, as well as by gender; likelihood ratings are therefore presented in Table 2 collapsed over these factors.

Table 2. *Summary of Means and Standard Deviations for probability ratings in the between-group and within-group conditions.*

Within-Subject Factors			Low NFC		High NFC	
Salience	B Value	A Value	*n*	*M (SD)*	*n*	*M (SD)*
Overt	Present	Present	27	6.37 (1.31)	35	6.43 (1.04)
		Absent	27	5.00 (1.90)	35	5.11 (2.10)
	Absent	Present	27	2.78 (1.31)	35	2.77 (1.42)
		Absent	27	1.93 (1.39)	35	2.34 (1.63)
Covert	Present	Present	27	6.81 (0.40)	35	5.94 (1.70)
		Absent	27	5.30 (1.73)	35	5.54 (1.80)
	Absent	Present	27	2.52 (1.19)	35	2.54 (1.54)
		Absent	27	1.67 (1.11)	35	2.06 (1.43)

Note: Total $N = 62$.

A 2 × 2 × 2 × 2 mixed between-within subject ANOVA was conducted to test our hypotheses. Results are displayed in Table 3. There was a large significant main effect for A Value, which lends support to the hypothesis that participants would tend to violate the Markov Assumption. Another main effect was found concerning B Value, with a substantial effect size; this suggests that participants correctly inferred the crucial role of variable B when estimating the likelihood that C is present.

Additionally, a significant interaction effect between A Value and B Value was found, with a medium effect size. The means are displayed in Figure 2. Participants based their judgments on the level of variable A to a higher degree in situations in which B was present as opposed to absent.

No significant interaction effect was found between A Value and salience, nor was there a significant interaction effect between A Value and NFC. Based on these results, we cannot reject the null-hypotheses that Markov violations were unaffected by salience and participants' NFC values. Additionally, there was no significant three-way interaction effect between A Value, salience, and NFC. Remaining main and interaction effects, including those not presented in Table 3, were nonsignificant, ps > .05.

Table 3. *Relevant main effects and interaction effects from the four-way mixed between-within subject Analysis of Variance (Salience × A Value × B Value × Need for Cognition).*

Factors	$F(1, 60)$	p	η_p^2
A Value	28.07	< .001	.32
B Value	338.50	< .001	.85
A Value × B Value	4.20	.045	.07
A Value × Salience	0.67	.42	.01
A Value × NFC	2.08	.16	.03
A Value × B Value × Salience	1.04	.31	.02
A Value × Salience × NFC	1.35	.25	.02
A Value × B Value × NFC	0.16	.69	.003
A Value × B Value × Salience × NFC	1.92	.17	.03

Note: The same analysis was conducted a second time after removing six outliers whose scores were more than three standard deviations from the mean. All significant main and interaction effects presented in the table above remained significant in the second analysis.

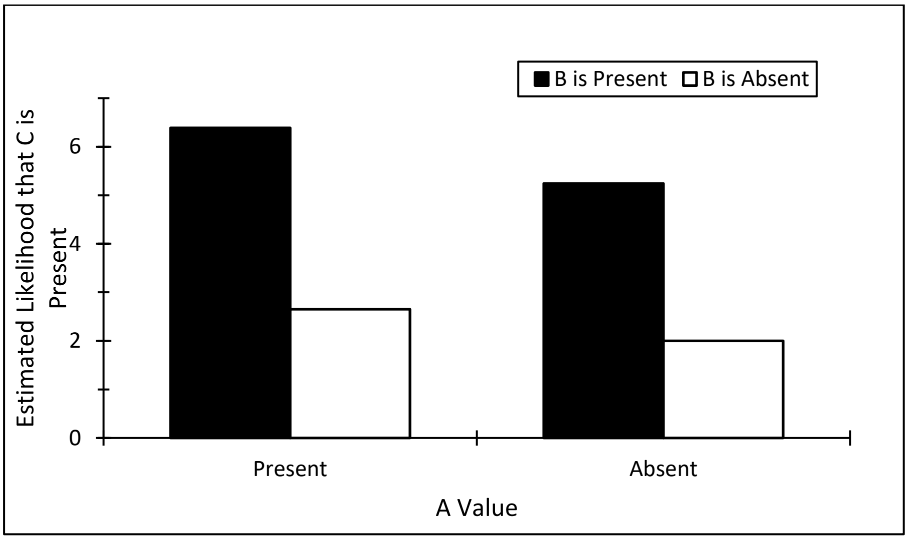

Figure 2. *Bar graph showing the interaction effect between A Value and B Value.*

Since a nonsignificant interaction effect between NFC and A Value was found, we reasoned that it might be appropriate to investigate whether there was a correlation between Markov violations in general and NFC, using the full range of NCS scores rather than the dichotomized variable. Markov violation scores were calculated by subtracting the sum of all probability scores when variable A is absent from the sum of all probability scores when A is present. A nonsignificant Pearson correlation was found between NFC (M = 1.21, SD = 1.07) and mean

Markov violations (M = 3.48, SD = 5.37), r(60) = -.13, p = .30, indicating no relation between the two variables. However, it was also found that the NCS scores were non-normally distributed, with skewness of -0.99 (SE = 0.30) and kurtosis of 2.42 (SE = 0.60). The distribution of NCS scores is displayed in Figure 3. Taken together, the descriptive statistics indicate a restriction in range for NCS scores. Because the NCS scores were non-normally distributed, a Spearman correlation was conducted to further test the relation between the variables; in line with previously reported results, no significant correlation was found between NCS scores and Markov violations, $r_s(60)$ = -.09, p = .50.

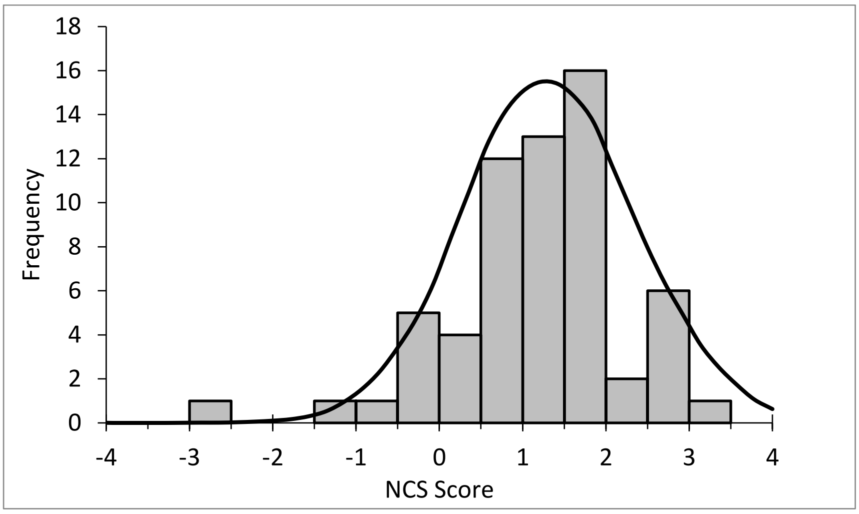

Figure 3. *Distribution of NCS scores displayed using a histogram. The black line represents how scores would be distributed had they followed a normal distribution curve.*

Discussion

The purpose of the study was to investigate the role of NFC and variable salience in Markov violations. We hypothesized that, in line with previous research (e.g., Mayrhofer et al., 2010; Rehder, 2014; Rehder & Burnett, 2005), participants would show a tendency to violate the Markov Assumption regardless of the properties of the variables. This prediction was confirmed by our substantial main effect of A Value on participants' likelihood ratings that C would be present. Interestingly, however, the current study also found a significant medium-sized interaction effect between A Value and B Value, which suggests that Markov violations are more likely to occur for certain values of the variable directly connected to the to-be-inferred variable; in this case, the value of B appeared to affect the extent to which participants used variable A to estimate the likelihood of C. This novel finding suggests that although violating the Markov Assumption is a general tendency that appears regardless of context, there are some circumstances under which violations are more likely to occur. Before proceeding, we want to remind the reader that the current study chose to focus on Markov violations in causal chain networks, given the limiting circumstances regarding time and resources. Therefore, although previous research has found Markov violations in both common cause, common effect, and causal chain networks (e.g., Rehder, 2014), we are cautious about generalizing the interaction effect between A Value and B Value beyond the causal model used in our experiment.

Nevertheless, this finding is relevant because of the importance to not only know in which situations people are more likely to violate the Markov Assumption, but also in which situations violations are committed to a higher degree. A number of implications are possible. For instance, one situation that comes to mind concerns program evaluations on the workplace, and how expectations of causal relations can prevent people from making objective and correct inferences. Imagine a workplace that implements a program (A) known to cause high motivation (B) among employees, because research (fictive in this example) has found that high motivation leads to high productivity (C). The interaction effect found in our study posits that, in a situation where the program is implemented (A is present) and high motivation is observed (B is present), people would be inclined to expect a higher productivity compared to if the program were not implemented but the motivation remained high. This higher-than-normal expectation would likely cause disappointment when evaluators find out that the productivity, albeit high, did not turn out as promising as was incorrectly assumed. We imagine that the implemented program might in this case be deemed less efficient, and thus a well-functioning program would risk abandonment due to the tendency to violate the Markov Assumption and falsely overestimate the value of C when both A and B are present. Awareness of the tendency to commit Markov violations might help a practitioner to identify and address such incorrect inferences, so that objective evaluations can be made.

This study was a first attempt to test the role of individual differences in explaining the tendency for some people to violate the Markov Assumption while others do not (Rehder, 2014), in order to better predict performance regarding judgement and decision-making. We argued that, because NFC has previously been connected to areas highly relevant for work (Haugtvedt & Petty, 1992; Kearney et al., 2009), decision-making (Carnevale et al., 2011; Smith & Levin, 1996; Stanovich & West, 1999), and rational thinking (Epstein et al., 1996;

Leary et al., 1986), it might also be able to explain the individual differences in Markov violations. Additionally, the size of the 18-item NCS further strengthened the suitability of the test for our study. That said, no significant effects of NFC were found on participants' Markov violations, neither when using a binary variable nor the full spectrum of the NCS index. The conclusions that can be drawn from the current study are however limited by the non-normal distribution of NCS scores, which showed clear signs of restriction in range. One potential explanation for such a finding could be that our sample mainly consisted of university students; arguably, choosing to deepen one's knowledge by means of further education would imply a preference for deeper-level thinking. The possibility remains that a different result would be obtained had the sample been drawn from a more heterogeneous population; given the time and resource constraints of the current study, however, recruiting students was considered the best option. Since NFC already has a number of real world applications, especially work-related (e.g., Kearney et al., 2009), a future study with a more heterogeneous sample may find the connection we hypothesized about and thereby make NFC a valuable predictor for normative causal reasoning, applicable to assessment of personnel as well as work performance.

A different approach would be to examine the relationship between Markov violations and rationality directly; Stanovich et al. (2016) describe their journey towards a prototype of a rationality test that they named the Comprehensive Assessment of Rational Thinking (CART). There is reason to believe that a full-fledged test of rationality might be related to Markov violations, given the different cognitive biases and heuristics that the CART is based on (Stanovich, 2016; Stanovich et al., 2016). Until the CART is fully developed, however, the Adult Decision-Making Competence index (A-DMC; Bruine de Bruin, Parker, & Fischhoff, 2007) might be a useful instrument for scientist seeking to further examine individual differences in relation to Markov violations. Finding a dispositional variable related to Markov violations could potentially generate highly applicable tools for predicting job performance, especially for business centred around causal judgement and decision-making.

No effect of salience was found on Markov violations, as opposed to what was hypothesized. One potential explanation could be that salience was only manipulated on variable A and therefore the salience of variables B and C were not taken into consideration when constructing the vignettes. This was motivated by the fact that in our experiment the Markov Assumption was violated when the information of variable A was incorrectly used to infer the value of C. Therefore, variable salience should only be relevant for variable A, because it is the variable that the Markov Assumption is based on in this specific example. At the same time, we do reserve ourselves to the possibility that perhaps a different result could be obtained if future researchers distribute salience evenly across all three variables, so that A, B and C all take on the same value of either overt or covert.

Furthermore, we chose to divide salience into overt and covert, where overt variables were described as explicit and openly observable whereas covert variables were hidden and non-observable. While this distinction is novel to the extent of our knowledge, one could apply these definitions to previous studies on Markov violations in an attempt to explain the nonsignificant effects of the current study. For instance, Walsh and Sloman (2007) found Markov violations in common cause networks in which the cause was covert (worrying), while the domains tested in Rehder (2014) arguably are observable and thus overt. The fact that previous studies have found Markov violations independent of domains further supports the view that violating the Markov Assumption is a general tendency. At the same time, although previous studies have found Markov violations regardless of domain properties, one could hypothesize that there are situations in which normative models are more frequently violated. Take the interaction effect between A Value and B Value that was found in the current study, for example. Although a substantial main effect of A Value was observed, which implies a general tendency to violate the Markov Assumption, violations were committed more frequently when B was present as opposed to absent. The same principle can be applied to variable properties in general, meaning that there might exist variable domains that evoke more violations of normative models than others. Perhaps a different categorization than the one used in the current study, preferably with more theoretical grounds, might be relevant to investigate.

Conclusion

The current study was a first attempt to investigate the role of individual differences in Markov violations. This focus was regarded as highly relevant, given that a connection between Markov violations and specific dispositional variables could provide a useful tool for measuring performance in areas related to judgement and decision-making. NFC was chosen because it showed promising connections to rational thinking (Epstein et al., 1996; Leary et al., 1986); however, it did not influence participants' violations of the Markov Assumption. In the framework of our limitations we provided insights on how future researchers might wish to further investigate the relation between individual differences and Markov violations. Furthermore, no effect of salience was found, as participants violated the Markov Assumption regardless. The study makes a novel contribution to research about causal reasoning in that there appear to be situations in which participants violate the Markov Assumption to a higher degree; in our case, more violations were committed when variable B was present as opposed to absent. We believe that future research in this particular aspect of Markov violations might unravel why people tend to deviate from normative causal models.

Acknowledgements

We wish to thank our supervisor, Jean-Christophe Rohner, for his continuous support throughout the research process. Our sincere gratitude also goes to Fredrik Björklund, for providing helpful comments during the initial phase of the project.

References

Bruine de Bruin, W., Parker, A. M., & Fischhoff, B. (2007). Individual differences in adult decision-making competence. *Journal of Personality and Social Psychology*, *92*(5), 938-956. doi:10.1037/0022-3514.92.5.938

Burke, W. W., & Litwin, G. H. (1992). A causal model of organizational performance and change. *Journal of Management*, *18*(3), 523-545. doi:10.1177/014920639201800306

Cacioppo, J. T., & Petty, R. E. (1982). The need for cognition. *Journal of Personality and Social Psychology*, *42*(1), 116-131. doi:10.1037/0022-3514.42.1.116

Cacioppo, J. T., Petty, R. E., Feinstein, J. A., & Jarvis, W. B. G. (1996). Dispositional differences in cognitive motivation: The life and times of individuals varying in need for cognition. *Psychological Bulletin*, *119*(2), 197-253. doi:10.1037/0033-2909.119.2.197

Cacioppo, J. T., Petty, R. E., & Kao, C. F. (1984). The efficient assessment of need for cognition. *Journal of Personality Assessment*, *48*(3), 306-307. doi:10.1207/s15327752jpa4803_13

Carnevale, J. J., Inbar, Y., & Lerner, J. S. (2011). Individual differences in need for cognition and decision-making competence among leaders. *Personality and Individual Differences*, *51*, 274-278. doi:10.1016/j.paid.2010.07.002

Epstein, S., Pacini, R., Denes-Raj, V., & Heier, H. (1996). Individual differences in intuitive-experiential and analytical-rational thinking styles. *Journal of Personality and Social Psychology*, *71*(2), 390-405. doi:10.1037/0022-3514.71.2.390

Glymour, C. (1998). Learning causes: Psychological explanations of causal explanation. *Minds and Machines*, *8*, 39-60. doi:10.1023/A:1008234330618

Haugtvedt, C. P., & Petty, R. E. (1992). Personality and persuasion: Need for cognition moderates the persistence and resistance of attitude changes. *Journal of Personality and Social Psychology*, *63*(2), 308-319. doi:10.1037/0022-3514.63.2.308

Hausman, D. M., & Woodward, J. (1999). Independence, invariance and the causal Markov condition. *British Journal for the Philosophy of Science*, *50*, 521-583. Retrieved from https://www-jstor-org.ludwig.lub.lu.se/stable/40072257

Highhouse, S. (2001). Judgment and decision-making research: Relevance to industrial and organizational psychology. In N. Anderson, D. S. Ones, H. K. Sinangil, & C. Viswesvaran (Eds.), *Handbook of Industrial, Work & Organizational Psychology: Vol. 2. Organizational Psychology* (pp. 314-331). London: SAGE Publications.

Kearney, E., Gebert, D., & Voelpel, S. C. (2009). When and how diversity benefits teams: The importance of team members' need for cognition. *Academy of Management Journal*, *52*(3), 581–598. doi:10.5465/AMJ.2009.41331431

Leary, M. R., Shepperd, J. A., McNeil, M. S., Jenkins, T. B., & Barnes, B. D. (1986). Objectivism in information utilization: Theory and measurement. *Journal of Personality Assessment*, *50*(1), 32-43. doi:10.1207/s15327752jpa5001_5

Mayrhofer, R., Hagmayer, Y., & Waldmann, M. (2010). Agents and causes: A Bayesian error attribution model of causal reasoning. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *32*, 925-930. Retrieved from https://escholarship.org/uc/item/294533f0

Pearl, J. (1988). *Probabilistic reasoning in intelligent systems*. San Francisco, CA: Morgan Kaufmann Publishers, Inc.

Pearl, J. (2000). *Causality: Models, reasoning, and inference*. Cambridge, MA: Cambridge University Press.

Rehder, B. (2014). Independence and dependence in human causal reasoning. *Cognitive Psychology*, *72*, 54-107. doi:10.1016/j.cogpsych.2014.02.002

Rehder, B. (2018). Beyond Markov: Accounting for independence violations in causal reasoning. *Cognitive Psychology*, *103*, 42-84. doi:10.1016/J.COGPSYCH.2018.01.003

Rehder, B., & Burnett, R. (2005). Feature inference and the causal structure of categories. *Cognitive Psychology*, *50*, 264-314. doi:10.1016/j.cogpsych.2004.09.002

Rottman, B. M., & Hastie, R. (2014). Reasoning about causal relationships: Inferences on causal networks. *Psychological Bulletin*, *140*(1), 109-139. doi:10.1037/a0031903

Slovic, P., & Tversky, A. (1974). Who accepts Savage's axiom? *Behavioral Science*, *19*(6), 368–373. doi:10.1002/bs.3830190603

Smith, B. N., Kerr, N. A., Markus, M. J., & Stasson, M. F. (2001). Individual differences in social loafing: Need for cognition as a motivator in collective performance. *Group*

Dynamics: Theory, Research, and Practice, *5*(2), 150-158. doi:10.1037/1089-2699.5.2.150

Smith, S. M., & Levin, I. P. (1996). Need for cognition and choice framing effects. *Journal of Behavioral Decision Making*, *9*, 283-290. doi:10.1002/(SICI)1099-0771(199612)9:4<283::AID-BDM241>3.0.CO;2-7

Stanovich, K. E. (2012). On the distinction between rationality and intelligence: Implications for understanding individual differences in reasoning. In K. Holyoak & R. Morrison (Eds.), *The Oxford handbook of thinking and reasoning* (pp. 433-455). New York, NY: Oxford University Press.

Stanovich, K. E. (2016). The comprehensive assessment of rational thinking. *Educational Psychologist*, *51*(1), 23-34. doi:10.1080/00461520.2015.1125787

Stanovich, K. E., & West, K. F. (1998). Individual differences in framing and conjunction effects. *Thinking and Reasoning*, *4*(4), 289-317. doi:10.1080/135467898394094

Stanovich, K. E., & West, K. F. (1999). Discrepancies between normative and descriptive models of decision making and the understanding/acceptance principle. *Cognitive Psychology*, *38*, 349-385. doi:10.1006/cogp.1998.0700

Stanovich, K. E., & West, R. F. (2008). On the relative independence of thinking biases and cognitive ability. *Journal of Personality and Social Psychology*, *94*(4), 672-695. doi:10.1037/0022-3514.94.4.672

Stanovich, K. E., West, R. F., & Toplak, M. E. (2016). *The rationality quotient: Toward a test of rational thinking*. Cambridge, MA: MIT Press.

Tversky, A., & Kahneman, D. (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review*, *90*(4), 293-315. doi:10.1037/0033-295X.90.4.293

Walsh, C. R., & Sloman, S. A. (2007). Updating beliefs with causal models: Violations of screening off. In M. A. Gluck, J. R. Anderson, & S. M. Kosslyn (Eds.), *Memory and mind: A festschrift for Gordon H. Bower* (pp. 345–358). New York, NY: Erlbaum. Retrieved from https://www.researchgate.net/publication/242182215

Appendix

Vignettes Used in the Study

In this version of the survey, four overt vignettes were presented first, followed by four covert vignettes. The A Values and B Values were counterbalanced across the eight vignettes and the eight versions of the survey.

Vinjett 1

Orsaksmodell

För aliens gäller i allmänhet följande:

- Långa fingrar orsakar blå naglar.
- Blå naglar orsakar hög attraktivitet.

Fall

En viss alien har långa fingrar. Denna alien har **inte** blå naglar. Hur stor tror du sannolikheten är att denna alien har hög attraktivitet?

1	2	3	4	5	6	7
Inte alls sannolikt			Varken sannolikt eller osannolikt			Väldigt sannolikt

---

Vinjett 2

Orsaksmodell

För aliens gäller i allmänhet följande:

- Hög konsumtion av ämnet zrxium orsakar hög hjärtfrekvens.
- Hög hjärtfrekvens orsakar hög signalstyrka vid telepati.

Fall

En viss alien har **inte** hög konsumtion av zrxium. Denna alien har hög hjärtfrekvens. Hur stor tror du sannolikheten är att denna alien har hög signalstyrka vid telepati?

1	2	3	4	5	6	7
Inte alls sannolikt			Varken sannolikt eller osannolikt			Väldigt sannolikt

Vinjett 3

Orsaksmodell

När vissa aliens blir vuxna börjar deras kroppar att permanent utstråla ett grönt ljus.

För aliens gäller i allmänhet följande:

- Utstrålning av grönt ljus orsakar hög fantasiförmåga.
- Hög fantasiförmåga orsakar lång livslängd.

Fall

En viss alien utstrålar ett grönt ljus. Denna alien har hög fantasiförmåga. Hur stor tror du sannolikheten är att denna alien har lång livslängd?

1	2	3	4	5	6	7
Inte alls sannolikt			Varken sannolikt eller osannolikt			Väldigt sannolikt

---

Vinjett 4

Orsaksmodell

På en utomjordisk planet gäller i allmänhet följande:

- Hög temperatur i luften orsakar tjockt kristallskikt i atmosfären.
- Tjockt kristallskikt i atmosfären orsakar högt lufttryck.

Fall

I ett land på den utomjordiska planeten är det **inte** hög temperatur i luften. På denna plats är kristallskiktet i atmosfären **inte** tjockt. Hur stor tror du sannolikheten är att det är högt lufttryck?

1	2	3	4	5	6	7
Inte alls sannolikt			Varken sannolikt eller osannolikt			Väldigt sannolikt

---

Vinjett 5

Orsaksmodell

För aliens gäller i allmänhet följande:

- Hög intelligens orsakar hög fysisk styrka.
- Hög fysisk styrka orsakar ljus röstfrekvens.

Fall

En viss alien har hög intelligens. Denna alien har hög fysisk styrka. Hur stor tror du sannolikheten är att denna alien har ljus röstfrekvens?

1	2	3	4	5	6	7
Inte alls sannolikt			Varken sannolikt eller osannolikt			Väldigt sannolikt

Vinjett 6

Orsaksmodell

För aliens gäller i allmänhet följande:

- Hög ängslighet orsakar hög nivå av musikalitet.
- Hög nivå av musikalitet orsakar hög makt i samhället.

Fall

En viss alien har hög ängslighet. Denna alien har **inte** en hög nivå av musikalitet. Hur stor tror du sannolikheten är att denna alien har hög makt i samhället?

1	2	3	4	5	6	7
Inte alls sannolikt			Varken sannolikt eller osannolikt			Väldigt sannolikt

---

Vinjett 7

Orsaksmodell

Utomjordiska havssniglar kan antingen tänka på mat eller inte.

För havssniglar gäller i allmänhet följande:

- Att havssniglar tänker på mat orsakar att de rör sig med hög hastighet.
- Att de rör sig med hög hastighet orsakar att de får hårda skal.

Fall

En viss havssnigel tänker **inte** på mat. Denna havssnigel rör sig med hög hastighet. Hur stor tror du sannolikheten är att denna havssnigel får ett hårt skal?

1	2	3	4	5	6	7
Inte alls sannolikt			Varken sannolikt eller osannolikt			Väldigt sannolikt

---

Vinjett 8

Orsaksmodell

För aliens gäller i allmänhet följande:

- Hög nivå av empati orsakar hög popularitet.
- Hög popularitet orsakar stora huvuden.

Fall

En viss alien har **inte** en hög nivå av empati. Denna alien har **inte** hög popularitet. Hur stor tror du sannolikheten är att denna aliens huvud är stort?

1	2	3	4	5	6	7
Inte alls sannolikt			Varken sannolikt eller osannolikt			Väldigt sannolikt