

**Returning advice taking to the wild:
Empirical, theoretical, and normative implications of
an ecological perspective**

Dissertation

der Mathematisch-Naturwissenschaftlichen Fakultät
der Eberhard Karls Universität Tübingen
zur Erlangung des Grades eines
Doktors der Naturwissenschaften
(Dr. rer. nat.)

vorgelegt von
Fabian Ache
aus Herdecke

Tübingen
2017

Gedruckt mit Genehmigung der Mathematisch-Naturwissenschaftlichen Fakultät der
Eberhard Karls Universität Tübingen.

Tag der mündlichen Qualifikation:	26.09.2017
Dekan:	Prof. Dr. Wolfgang Rosenstiel
1. Berichterstatter:	Jun. Prof. Dr. Mandy Hütter
2. Berichterstatter:	Prof. Dr. Wolfgang Gaissmaier
3. Berichterstatter:	Prof. Dr. Andreas Voß

Acknowledgments

The past two and a half years that led to this dissertation would not have been such an infinite jest without the people who accompanied me along the way.

Thank you, Mandy, for providing this opportunity, for appreciating my ideas, and for pushing me towards my limits and beyond.

Thank you, David, Kathi, Max, and Max, for being such great colleagues, for supporting me scientifically and emotionally, and for supplying an abundance of sweets.

Thank you, Denise, Margret, and the HiWis, for laying the ground on which all of this grew.

Thank you, drivers of bus line 826, for getting me to work safely, for outperforming the Deutsche Bahn by a big margin, and for giving me the most time to read that I had in years.

Thank you, Marion and Peter, for laying the ground on which I grew, for supporting everything I do, and for being the best parents I can imagine.

Thank you, Jan, for appreciating my figures, for discussing science and not just psychology, and for being my brother and my role model.

Thank you, Lisa, for sharing all these wonderful moments with me, for making me happier at all times, and for reminding me of what truly matters.

To all of you, I want to express my gratitude.

Fabian

Table of contents

Abstract	6
1. Returning advice taking to the wild	8
1.1 Advice taking	8
1.2 Three problems of advice taking research	10
1.3 A solution: An ecological perspective	14
1.4 Overview of thesis' structure	16
2. Exploiting the information ecology: Re-evaluating existing norms from an ecological perspective	17
2.1 Simulation	19
2.1.1 Method	20
2.1.2 Results	21
2.2 Interim discussion	23
2.2.1 Boundary conditions	26
3. Expanded information asymmetry account	28
4. Exploring the information ecology: Investigating information search in advice taking... ..	34
4.1 Overview of experiments and hypotheses	35
4.2 Data analysis	36
4.3 A-priori power analyses	39
4.4 Experiment 1	39
4.4.1 Method	40
4.4.2 Results	41
4.4.3 Discussion	43
4.5 Experiment 2	43
4.5.1 Method	44
4.5.2 Results	44
4.5.3 Discussion	49
4.6 Experiment 3	49
4.6.1 Method	50
4.6.2 Results	51
4.6.3 Discussion	52
4.7 Interim discussion	53
5. Integrating the information ecology: Converging evidence and potential boundaries	56
5.1 Overview of experiments and hypotheses	58
5.2 Data analysis	59
5.3 A-priori power analyses	61

5.4	Experiment 4.....	61
5.4.1	Method	62
5.4.2	Results	63
5.4.3	Discussion	67
5.5	Experiment 5.....	67
5.5.1	Method	68
5.5.2	Results	69
5.5.3	Discussion	76
5.6	Interim discussion.....	76
5.7	Limitations.....	77
6.	General discussion	80
6.1	Information search in advice taking	80
6.1.1	Limitations and future directions	83
6.2	Expanded information asymmetry account revisited	85
6.2.1	Limitations and future directions	86
6.3	Adaptive advice taking.....	90
6.3.1	Limitations and future directions	91
6.4	Ecological validity of advice taking: A disclaimer	92
6.5	Conclusion: Implications of an ecological perspective	95
	References	97
	Appendix	109
	Author contributions	114

Abstract

Advice taking constitutes a strategy to increase judgment accuracy that builds on existing ecological structures. Specifically, averaging the own judgment with that of others yields an increase in expected accuracy. Moreover, many decision situations are inherently interactive, thereby offering the opportunity to acquire advice before producing a final judgment. Theoretically, advice taking is thus highly adaptive.

However, research on advice taking hinges on three core problems: the uncertain adaptivity of advice taking behavior; the lack of a unifying theoretical framework; and the conceptual constraints of existing research. First, whereas averaging the own judgment and the advice yields highest expected accuracy, individuals often place a relatively low weight on the advice during judgment revision. Second, there is no theoretical framework to integrate the many determinants of advice taking already identified. Third, advice taking research almost exclusively focused on the weighting of advice once it is received, thereby neglecting information search.

The current thesis addresses these three problems from an ecological perspective. This perspective assumes that the interactive nature of advice taking affords the individual with the opportunity to seek additional information. Specifically, advice taking is conceptualized as a sequential process that can be truncated at different points in time, for example after consulting different amounts of advice. To this end, the current thesis introduces an expanded advice taking paradigm that allows individuals to sample any number of other people's judgments before revising their own. The current thesis proposes an information integration account of advice taking that expands existing theoretical ideas, incorporates information search and integrates the most important determinants of advice taking as identified by existing research. The adaptivity of advice taking is re-evaluated from this ecological perspective. A simulation assesses which integration strategy fares best in an environment that

accords to the ecological perspective advocated in this thesis. Experiments investigate whether individuals explore and adapt to the information ecology when given the opportunity.

The results suggest that advice taking is much more adaptive than previously documented. Placing a low weight on the advice yields higher accuracy than weighting the own judgment and the advice equally when advice taking is a sequential process with several steps of updating. Moreover, individuals explore and adapt to the information ecology when given the opportunity. The experimental data also support the proposed theoretical framework of advice taking. This thesis provides strong evidence that the ecological nature of advice taking can no longer be neglected by advice taking research. Further discrepancies between advice taking in the laboratory and advice taking in the wild, necessary steps to further validate the proposed theoretical account, and future directions to investigating the adaptivity of advice taking are discussed.

1. Returning advice taking to the wild

In our daily lives, accurate judgments are equally challenging and rewarding. Be it the number of refugees fleeing a country due to civil war, the exchange rate of the dollar after the presidential election, or the number of people affected by the most recent outbreak of yet another influenza virus, the truth can often only be known in retrospect.

One remarkably easy yet powerful strategy to deal with such judgments under uncertainty is to take advice from others. Indeed, averaging several judgments increases accuracy (wisdom of crowds; Davis-Stober, Budescu, Dana, & Broomell, 2014; Einhorn, Hogarth, & Klempner, 1977; Galton, 1907). Specifically, averaging reduces random error as well as opposing biases (Einhorn et al., 1977; Larrick & Soll, 2006). It performs well in a wide variety of environments (Davis-Stober et al., 2014) and in comparison to more complex models of aggregation (Clemen, 1989). Furthermore, many decisions occur in an interactive, social context (Larrick, Mannes, & Soll, 2012; Sniezek & Buckley, 1995). Specifically, many decision situations cannot be accurately described as either an individual acting in isolation or a group of equally involved individuals acting together (Sniezek & Buckley, 1995). Rather, individuals often consult others before making a judgment. In fact, cooperation has recently been suggested to be automatic (Rand, 2016; Rand, Greene, & Nowak, 2012). Likewise, the ability to effectively learn from each other has been suggested to constitute the main difference between human intelligence and that of other animals (Herrmann, Call, Hernández-Lloreda, Hare, & Tomasello, 2007). Advice taking thus offers a way to improve judgment accuracy that builds on existing ecological structures (Simon, 1956). Theoretically, it constitutes a highly adaptive strategy for judgment under uncertainty (Yaniv & Milyavsky, 2007).

1.1 Advice taking

A variety of research investigates how individuals use advice to adapt and improve their judgments (e.g., Harvey & Fischer, 1997; Yaniv & Kleinberger, 2000; for a review see

Bonaccio & Dalal, 2006). Research on judgment aggregation predicts the highest expected accuracy to result from simply averaging all available judgments (e.g., Davis-Stober et al., 2014). Advice taking research thus mainly focuses on the degree to which individuals integrate advice received into their judgment during revision (e.g., Soll & Larrick, 2009; Yaniv, 2004a). Individuals generally improve their judgment through taking advice (e.g., Mannes, 2009; Yaniv, 2004a). At the same time, however, individuals regularly place a greater weight on their own judgment than on the advice (egocentric discounting; Yaniv & Kleinberger, 2000).

Advice weighting is sensitive to a number of moderators which can be classified into characteristics of the advisee, characteristics of the task and characteristics of the advisor or advice (Gino, Brooks, & Schweitzer, 2012). For example, individuals weight advice more strongly when they are anxious (Gino et al., 2012). Likewise, increasing the difficulty of the task by decreasing the recognizability of the stimulus to be judged increases advice weighting (Gino & Moore, 2007; Yaniv & Choshen-Hillel, 2012). Advice weighting also increases when the advisor has a higher expertise (Harvey & Fischer, 1997) or when the advice is costly to obtain (Gino, 2008). One highly important determinant of individuals' advice weighting is the confidence in their own judgment (e.g., Gino et al., 2012). Indeed, individuals weight advice more strongly when they are less confident (e.g., Gino et al., 2012; Gino & Moore, 2007; Moussaïd, Kämmer, Analytis, & Neth, 2013) and confidence mediates the effect of anxiety on advice weighting (Gino et al., 2012). However, confidence can also be the outcome of advice taking (e.g., Schultze, Rakotoarisoa, & Schulz-Hardt, 2015). Specifically, individuals increase their judgment related confidence after receiving advice that is similar to their own judgment (Moussaïd et al., 2013; Schultze et al., 2015; Yaniv, Choshen-Hillel, & Milyavsky, 2009).

While research has mostly focused on advice taking as a strategy to improve judgment accuracy (Mannes, 2009; Soll & Larrick, 2009; Yaniv, 2004a, 2004b), individuals might also take advice to share responsibility (Harvey & Fischer, 1997; Kennedy, Kleinmuntz, &

Peecher, 1997) and to not reject help that they were offered (Harvey & Fischer, 1997). In this way, advice taking also constitutes a basic form of social exchange (Goldsmith & Fitch, 1997; Homans, 1958).

1.2 Three problems of advice taking research

Research on advice taking has high theoretical and practical relevance. Taking advice from others constitutes an easy to use yet powerful strategy that allows individuals to produce relatively accurate judgments in light of uncertainty. Advice taking constitutes a basic research paradigm to study interactive social decision making. Ultimately, investigating how and why individuals (fail to) gain from information obtained from other individuals might also deepen our understanding of the merits and perils of a democratic society.

Advice taking research produced a number of interesting and highly relevant findings. However, further progress hinges on three core problems that have to be addressed. First, whereas advice taking offers the potential of efficiently and effectively improving judgment accuracy, individuals' realization of this potential appears suboptimal in most advice taking studies. Second, despite the many identified determinants of individuals' advice weighting, there is yet no unifying theoretical framework that allows to integrate the most central ones. Third, existing research (almost) exclusively focused on advice weighting while neglecting advice seeking, constraining the ecological validity and thereby the development of a comprehensive account as well as any valid assessment of the adaptivity of advice taking.

Despite the general improvement in accuracy through taking advice, individuals' advice weighting is widely assumed to be suboptimal (e.g., Mannes, 2009; Yaniv, 2004a). Specifically, highest accuracy is expected for weighting the own judgment and the advice received equally (e.g., Mannes, 2009). However, individuals regularly place a relatively low weight on the advice (e.g., 0.2; Harvey & Fischer, 1997), placing the higher complement (i.e.,

1-0.2) on their own judgment (e.g., Bonaccio & Dalal, 2006; Yaniv & Kleinberger, 2000).¹

This egocentric discounting thus appears to preclude optimal gains from taking advice, casting doubt on the idea that individuals are able to realize the potential adaptivity of advice taking.

One account of this underweighting of advice assumes an information asymmetry in individuals' minds (Yaniv, 2004a, 2004b). Whereas individuals have at least some access to the reasoning underlying their own judgment, this is true to a lesser degree for the reasoning underlying some other individual's judgment. Consequently, the own judgment is weighted more strongly than that of others (Tversky & Koehler, 1994). Support for this assumption stems from studies reporting stronger advice weighting when the recognizability of the target stimulus is decreased (Gino & Moore, 2007; Yaniv & Choshen-Hillel, 2012), a manipulation that likely also affects the number of reasons individuals have in support of their own judgment. Some empirical evidence also appears to oppose the information asymmetry account. Reminding individuals of the cues their judgment was based on during judgment revision does not decrease advice weighting (Soll & Mannes, 2011). From this null-effect, one could conclude that the differential availability of reasons does not explain underweighting. Alternatively, one could conclude that the *mere repetition* of information underlying the initial judgment was not effective in increasing underweighting tendencies. This empirical criticism does not render the information asymmetry account obsolete. However, there are two important determinants of advice weighting that the information asymmetry account in its current form cannot explain.

It fails to account for the increase in advice weighting when advice is dissimilar from rather than similar to the own judgment (Moussaïd et al., 2013; Schultze et al., 2015).

Specifically, while initial research assumed a monotone negative relationship between the

¹ This pattern of underweighting has been suggested to actually result from aggregation over several trials in which individuals either choose to stick with the own judgment or adopt the advice, with the former being more frequent than the latter (Soll & Larrick, 2009).

distance of advice from the own judgment and advice weighting (Minson, Liberman, & Ross, 2011; Yaniv, 2004a), recent research provides ample evidence that this relationship is actually curvilinear. That is, advice is weighted most strongly when it is intermediately distant from the own judgment but is weighted less strongly when it is close to the own judgment (Moussaïd et al., 2013; Schultze et al., 2015). Rather, close advice leads to an increase in confidence (Moussaïd et al., 2013; Schultze et al., 2015).² The reasons in individuals' minds should correspond less to advice the more it differs from the own judgment. Therefore, if advice weighting was merely a function of the underlying reasons, advice weighting should decrease linearly with increasing advice distance.

Furthermore, the information asymmetry account so far has not been related to confidence. Although the accessibility of reasons in support of the own judgment can be assumed to be related to the associated level of confidence (Koriat, 2012a; Tversky & Koehler, 1994) the account does not explicitly predict this relation. The information asymmetry account therefore does not explain how confidence mediates for example the effect of anxiety on advice taking (Gino et al., 2012). Likewise, the information asymmetry account does not explain how and why advice taking might lead to an increase in confidence rather than judgment revision (e.g., Moussaïd et al., 2013; Schultze et al., 2015).

While the information asymmetry account is one of the very few existing theoretical approaches to advice taking, it has to be expanded to account for the above described important determinants of advice taking as identified by existing research.

However, there is an even more pressing problem of both current theoretical accounts of advice taking as well as all existing advice taking research, namely its conceptual constraints. Advice taking research almost exclusively focused on the weighting of advice

² This research also suggests a decrease of advice weighting when advice is very distant from the own judgment, presumably because very distant advice is implausible (Moussaïd et al., 2013; Schultze et al., 2015). However, given the recent criticism of testing u-shaped relations using quadratic functions (Simonsohn, in preparation) this evidence should be treated with caution.

once it is received (e.g., Harvey & Fischer, 1997; Mannes, 2009). While two studies investigated whether individuals solicited advice at all (Gino et al., 2012; Gino & Moore, 2007), no research has yet investigated the process of advice seeking (Bonaccio & Dalal, 2006; Gino et al., 2012). Specifically, in the classical research paradigm, the individual produces an initial estimate for a quantitative estimation task. The individual then receives a single piece of advice and in light of this information has the opportunity to revise his or her initial estimate. This approach constitutes a serious limitation of the ecological validity of advice taking research which is highly problematic with regard to existing theoretical accounts as well as any assessment of the adaptivity of advice taking.

Advice taking has to be preceded by advice seeking in most cases (Bonaccio & Dalal, 2006; Gino et al., 2012), especially because advice that was not solicited can be seen as “butting in” (Goldsmith & Fitch, 1997). Furthermore, this constraint of the current research limits individuals’ options to react to advice in an unduly manner. Specifically, in the existing research individuals could only take the advice or reject it as reliable information. The assessment of reliability of the advice in this case is constrained to the information in individuals’ minds, which can be assumed to relate to the own judgment more than to the advice (Yaniv, 2004a, 2004b). Consequently, the discounting of advice does not appear surprising and could even be seen as rational (Ravazzolo & Røisland, 2011). However, given that advice taking is grounded in the assumption that most decision situations are interactive (Sniezek & Buckley, 1995), the individual’s search for information should not be constrained to his or her own mind. Specifically, individuals outside the laboratory have the opportunity to investigate the reliability of the advice by exploring their information ecology (e.g., by consulting additional advisors). The classical research paradigm thus deprives individuals of an option to react to advice received that they have outside the laboratory, namely to seek additional information.

This constraint is highly problematic on a theoretical level. While the assumed skewness of initial information towards supporting the own judgment appears a plausible prior for many advice taking situations, the above reasoning renders it highly artificial to assume that this skewness cannot be altered. Specifically, information search might increase the amount of information relating to the advice received and thereby increase reliance on this advice. Existing theoretical accounts have to be expanded by this opportunity.

This constraint also has two important implications for the adaptivity of advice taking. First, if individuals expanded the information in their minds by the information sampled from the ecology as suggested above, advice taking might be much more flexible than documented by existing research. Second, existing norms for advice taking have only ever considered a single instance of judgment updating. In contrast, the opportunity to seek additional information renders advice taking a sequential process with consecutive instances of updating. Potentially, placing a greater weight on the own judgment than on the advice yields high accuracy when several pieces of advice are encountered sequentially.

1.3 A solution: An ecological perspective

The goal of the current thesis is to advance research on advice taking by addressing the above identified three problems: the uncertain adaptivity of advice taking behavior; the lack of a unifying theoretical framework; and the conceptual constraint of the existing research. As laid out above, both the evaluation of adaptivity as well as the development of a comprehensive theoretical account of advice taking so far were constrained by the lack of any conceptualization of information search in advice taking (Bonaccio & Dalal, 2006; Gino et al., 2012). The current thesis thus addresses advice taking from an ecological perspective. The core assumption of this ecological perspective is that the interactive nature of advice taking affords the individual with the opportunity to seek additional information. It is assumed that advice taking often is a sequential process that can be truncated at different points in time. Thus, depending on characteristics of the decision task, personal factors (e.g., knowledge), or

the advice received, people could seek smaller or larger samples of information. For example, individuals could consult additional advisors when advice is dissimilar from their own judgment. Such interactions between cognition and the information ecology have long been recognized (Brunswik, 1955; Simon, 1956) and have recently been studied in related paradigms as well (e.g., Busemeyer & Townsend, 1993; Denrell, 2005; Hertwig, 2015; Hertwig & Erev, 2009). By introducing this ecological perspective into advice taking research, the current thesis strives to contribute to existing research in three ways.

First, an expanded advice taking paradigm is proposed that allows to incorporate information search into advice taking research. This provides the opportunity to investigate advice taking in an interactive situation, which is one of the assumed core properties of advice taking outside the laboratory (Sniezek & Buckley, 1995). This also provides a conceptual link between research on advice taking and related literatures on judgment and decision making (e.g., decisions from experience; Hertwig, 2015; Hertwig & Erev, 2009).

Second, an information integration account is proposed as a theoretical framework for advice taking. This is the first theoretical account that allows to derive predictions for precursors and consequences of information search in advice taking (Bonaccio & Dalal, 2006). This account constitutes a necessary starting point to integrate the various findings on determinants and outcomes of advice taking and provides a theoretical link to related literatures (e.g., Koriat, 2012a; Newell, 2005).

Third, the adaptivity of advice taking is re-evaluated from this perspective that is less constrained than previous conceptualizations. Here, adaptivity will be investigated in two ways. First, it will be assessed how far advice utilization is sensitive to the information provided by the information ecology. Second, the existing normative considerations for advice taking will be re-evaluated in light of the current ecological perspective. Consequently, this constitutes the first assessment of the adaptivity of advice taking as it might occur outside the laboratory.

1.4 Overview of thesis' structure

To this end, the current thesis presents one theoretical and three empirical chapters. Chapter two focuses on the normative considerations for advice taking from an ecological perspective. Data from a simulation is reported that investigates the accuracy resulting from different advice weighting strategies in an environment that accords to the ecological perspective advocated in the current thesis.

Chapter three gives a detailed account of the proposed information integration framework of advice taking and its core assumptions. It is also illustrated how existing findings can be integrated into this framework.

Chapter four introduces the expanded research paradigm that allows to incorporate information search in advice taking. Data from three experiments serve to investigate whether participants are willing to seek additional information in the advice taking paradigm and whether the opportunity to do so affects advice weighting behavior as assumed by the proposed theoretical framework.

Chapter five investigates participants' confidence as a proxy for their information samples to provide converging evidence for the proposed theoretical framework. To this end, data from two additional experiments is reported, in which participants' confidence was measured before and after receiving advice.

Chapter six discusses the documented evidence in light of the identified three problems of advice taking research. It is considered carefully how far the current thesis advances research on advice taking, which new questions emerge, and how they can be answered by future research.

2. Exploiting the information ecology: Re-evaluating existing norms from an ecological perspective

A plethora of research on the wisdom of crowds attests to the fact that simply averaging several judgments yields high accuracy (Davis-Stober et al., 2014; Einhorn et al., 1977; Galton, 1907). Consequently, a variety of studies investigated how individuals make use of this strategy, incorporating judgments obtained from other individuals (Yaniv, 2004a), groups (Mannes, 2009) or different sources of knowledge (Herzog & Hertwig, 2009). However, this research has greatly overlooked the dynamic nature of updating. Specifically, with the exception of a few rather peculiar topics, the crowd is not a fixed entity readily available to individuals in need of an accurate judgment. Outside of artificial laboratory settings, judgment updating most likely resembles a sequential process in which the individual consecutively consults different sources of information (cf., decision-field theory, Busemeyer & Townsend, 1993; decisions from experience, Hertwig & Erev, 2009). This process has two characteristics that have been ignored by existing research. First, the individual accumulates information on each instance of updating. Second, the number of judgments to be integrated cannot be known in advance.

Under such circumstances, averaging is more complex than in the idealized setting of a fixed number of judgments (N) that need to be accounted for. Specifically, true averaging for an a-priori unknown N requires a Bayesian updating algorithm that places a dynamic weight of $1/N$ on the newly encountered information (and the complement of $1 - 1/N$ on the previous judgment). This is not only much more challenging computationally, it also necessitates keeping a counter of the previous instances of updating to determine N . Consider for example an individual that gives an estimate of 10 (to a fictitious judgment task) and then encounters an advisory judgment of 30. Averaging would imply weighting both judgments by 0.5, or shifting the own judgment half the way towards the advice, resulting in an updated judgment of 20. If the individual next encountered a second advisory judgment of 50,

Bayesian updating dictated a weight of 0.66 on the individual's updated judgment and a weight of 0.33 on this second advice, or shifting one third towards the second advice.

Bayesian updating would thus result in an updated judgment of 30. This updated judgment is the average of all three judgments. By contrast, if the individual failed to account for the fact that it already updated its judgment after receiving the first advice and thus weighted also the second advice by 0.5, the individual's final judgment would be 35, which misses the average by 5 points.

Research has only investigated a single instance of updating. In this case the own and the additional judgment should be weighted equally. The classical finding is that individuals assign only low weights to additional judgments, demonstrating "underweighting". This phenomenon has been deemed a serious shortcoming (e.g., Mannes, 2009; Soll & Mannes, 2011; Yaniv, 2004a). However, as the example illustrates, already at the second instance of updating the own judgment should receive twice the weight of the additional judgment. At a potential third instance of updating, the individual's updated judgment ought already be weighted three times as strong as the additional judgment. In fact, the optimal weight placed on an additional judgment decreases rapidly (and approximates 0; see Figure 1).

In most instances, the number of judgments (N) ultimately acquired cannot be known in advance and keeping an accurate counter of the number of judgments already integrated constitutes a demanding task. Under these circumstances, dynamic Bayesian updating is not manageable. Instead, it is more feasible to rely on a heuristic strategy that assigns the same static weight to each encountered judgment. Such a heuristic might prove a powerful integration strategy that emulates the accuracy of Bayesian updating.

In the current chapter, it is suggested that for a wide range of instances of updating N , assigning a low weight (0.1; 0.2) to additional information (and a complementary high weight to the own judgment) constitutes such a satisficing solution (Simon, 1956) that approximates the accuracy of Bayesian updating. By contrast, if individuals were to weight every additional

judgment equally to their own (updated) judgment, as considered optimal in the current literature that focuses on single updating instances, accuracy should be much lower compared to both Bayesian updating and updating by low weights. To support this claim with empirical data and quantify the magnitude of the resulting accuracy differences, a simulation was run.

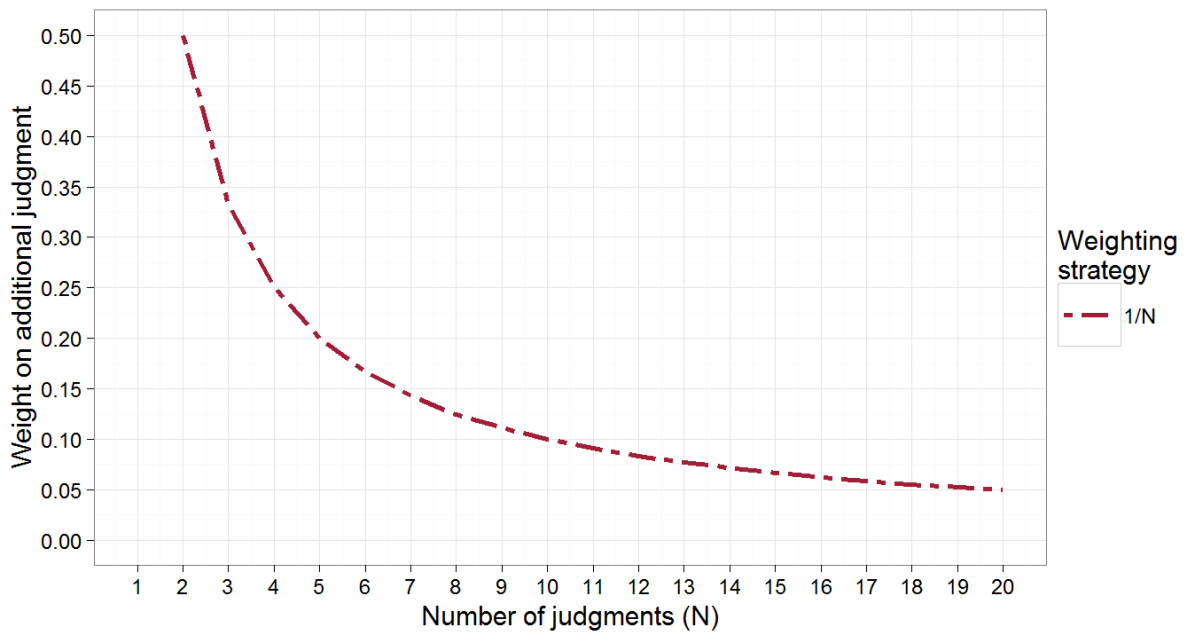


Figure 1. Optimal weight for a given additional judgment depending on the number of judgments encountered previously ($1/N$).

2.1 Simulation

1000 wise crowds were simulated (i.e., groups of 20 judgments whose average resembled the truth). From every crowd, one judgment was randomly chosen and assigned the role of the *individual* updating his or her own judgment. This *individual* was sequentially paired with each other member of the respective crowd. Along these pairings the *individual* cumulatively updated his or her judgment in accord with each of several weights (e.g., 0.5, 0.2, $1/N$). As the main dependent measure, the accuracy resulting from each weighting strategy was tracked along the pairings.

2.1.1 Method

Wise crowds. 1000 groups of 20 judgments each were created. These judgments were drawn at random from a normal distribution with $M = 0$ (to resemble the truth) and $SD = 25$, which served as a model for a wise crowd. The first judgment was always assigned the role of the *individual* updating his or her judgment. All other judgments served as the judgments the *individual* would be paired with.

Judgment updating. The *individual* was sequentially paired with each judgment from the respective crowd. For each pairing, the *individual's* judgment was updated in accord with the following formula:

$$judgment_t = (1 - x) * judgment_{t-1} + x * advice,$$

with *advice* being the newly encountered judgment and x being set to each of the following weighting strategies: $1/N$, 0.1, 0.2, 0.3, 0.4, 0.5. For the Bayesian strategy, x was dynamically set to $1/N$. That is, both the prior judgment and the newly encountered information were weighted with regard to the amount of information already integrated. This serves as a comparison standard. All other weights assigned were static. That is, the same weight was applied independent of whether the *individual* had previously integrated only 1 or already 10 other judgments. According to the assumption, low static weights (0.1; 0.2) should perform equally well as the Bayesian strategy for a wide range of N . To test whether this applies only to low static weights, the accuracy resulting from $x = 0.3$, $x = 0.4$ or $x = 0.5$, respectively, was also tested.

Updating was cumulative. For example, given the weighting strategy of 0.5, if the *individual's* initial judgment of 10 was first paired with a judgment of 20, the *individual's* updated judgment would be 15. This updated judgment (15) was then paired with the next judgment (e.g., 25) and resulted in the updated judgment of 20. This process continued until the *individual* had been paired with each other crowd member.

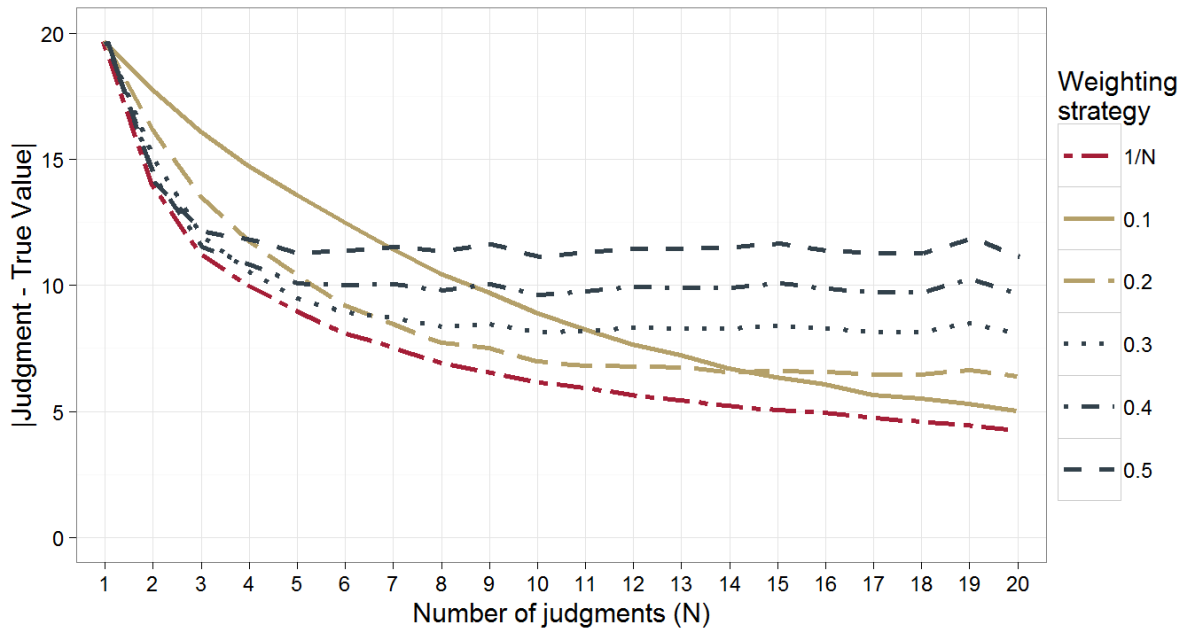


Figure 2. Absolute deviation of resulting judgment from true value (0) depending on the number of judgments encountered previously (N).

2.1.2 Results

The main dependent measure was the absolute deviation of the *individual's* updated judgment from the truth (here: 0) after having been paired with a given number of additional judgments (i.e., for increasing N ; Figure 2). As apparent from Figure 2, the dynamic weighting ($1/N$) strategy performs better than all static weighting strategies. While higher static weights (0.3, 0.4, 0.5) show a relatively steep increase in accuracy for updating in the range of $N = [2, 5]$, lower weights (0.1, 0.2) become more accurate for higher N , in the range of $N = [6, 20]$. Specifically, whereas higher weights reach a relatively stable level of accuracy in the range of $N = [5, 8]$, lower weights show an increase in accuracy even for higher numbers of N . Interestingly, assigning a static weight of 0.2 appears to closely mirror the accuracy of Bayesian updating across the whole range of N investigated.

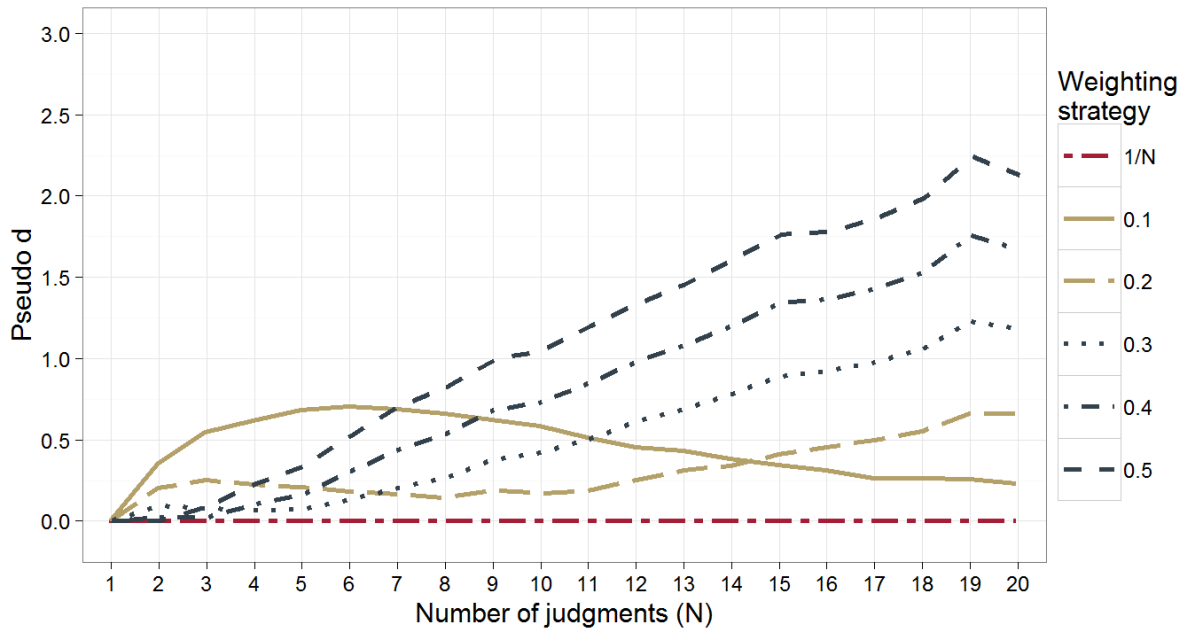


Figure 3. Standardized deviation of resulting judgment from optimal accuracy ($1/N$) depending on the number of judgments encountered previously (N). Standardization was performed for every judgment for a given number of additional judgments (N) by subtracting $M_{1/N}$ and dividing by $SD_{1/N}$.

To quantify the resulting differences in accuracy, this deviation was standardized in relation to the highest accuracy achievable in analogy to Cohen's d . To this end, the mean and standard deviation for the Bayesian strategy ($1/N$) were calculated for every N across all 1000 groups. Every judgment (for each remaining strategy and each N) was then standardized by subtracting this mean and dividing it by the respective standard deviation. Consequently, the resulting score indicates the number of standard deviations (of the optimal strategy) that a given strategy's accuracy deviates from the accuracy of the optimal strategy (pseudo d ; Figure 3). This loss in accuracy in comparison to the optimal strategy supports the above interpretation. Higher weights (0.3, 0.4, 0.5) perform well for small numbers of N (range of $N = [2, 5]$) but rapidly deteriorate beyond that range. In contrast to that, lower weights (0.1, 0.2) perform slightly less well for small numbers of N but do not deteriorate (and partly even improve performance) for the investigated range of $N = [2, 20]$. Importantly, whereas the higher weights all deviate by more than $d = 1$ in the range of N investigated, the lower

weights remain relatively accurate. Interestingly, at least for the range of $N = [2, 15]$, weighting by 0.2 never deviates by more than $d = 0.5$ from Bayesian updating.

One potential determinant of the loss in accuracy of a given strategy in comparison to the optimal strategy could be the summed deviation of the applied weights from $1/N$. For example, the strategy of .5 does not deviate from $1/N$ for $N = 2$, but deviates by .17 for $N = 3$ and by .25 for $N = 4$. Thus, the summed deviation for the strategy of .5 and $N = 4$ is .42. To test whether the summed deviation predicts the standardized loss in accuracy, the summed deviation for each strategy for a range of $N = [2, 20]$ was entered into a regression model as a predictor for the standardized loss in accuracy. Summed deviation from $1/N$ strongly predicted loss in accuracy, $b = .28$, $se = .01$, $t = 20.60$, $R^2 = .67$.

2.2 Interim discussion

A variety of research acknowledges the immense potential of averaging and investigates how individuals use this strategy to improve their judgments in one-shot (laboratory) situations. However, this research failed to take into account the dynamic nature of updating outside the laboratory. When information is encountered sequentially, over several instances of updating, true averaging amounts to a rather complex, dynamic Bayesian updating algorithm. It is argued and demonstrated by means of simulation, that a satisficing solution to dynamic, sequential updating need not be complicated. Simply assigning a low weight to newly encountered information yields accuracy comparable to that of optimal Bayesian updating across a wide range of N .

Interestingly, this advantage of low weights opposes the normative reference frame frequently applied in the existing literature. Specifically, it is often considered a serious shortcoming that individuals underweight other individuals' (Yaniv, 2004a) or groups' judgments (Mannes, 2009) when updating their own. For example, in the classical advice taking paradigm, individuals often weight the single piece of advice they receive by roughly 0.2 (e.g., Harvey & Fischer, 1997). Considering this a one-shot situation this resembles

underweighting and causes a loss in accuracy in comparison to weighting the advice by 0.5. However, the above simulation shows that shifting the frame of reference to a dynamic, sequential process that involves several instances of updating, as more likely outside the laboratory, habitually assigning such low weights actually yields higher accuracy and lower variability thereof (compare upper and lower panel; Figure 4) than the high weights that have been used as benchmarks in the existing literature. Consequently, what appeared a shortcoming given the laboratory situation might actually be highly adaptive given the real world analog.




Figure 4. Variability of resulting judgment around true value (0; red line) depending on the number of judgments encountered previously (N). Every black line indicates accuracy for one of the 1000 individuals simulated.

The assumption that static, one-shot updating fails to capture updating outside the laboratory depends on the following two premises. First, individuals accumulate information with each instance of updating and, second, updating frequency is greater than $N = 1$. It was carefully considered how often these premises hold in the real world. The first premise would not be met if individuals' judgments showed compliance rather than conversion. If individuals were compliant, they might adapt their judgment in a given situation but would afterwards

turn back to their initial judgment (Moscovici, 1980). However, for several reasons it appears more likely that judgment updating constitutes information accumulation. Specifically, individuals often seek information and advice as a form of informational influence (Deutsch & Gerard, 1955), which is more likely to lead to conversion than compliance (Mannes, 2009; Moscovici, 1980). Moreover, laboratory experiments investigating judgment updating generally leave little room for compliance. For example, in the classical advice taking paradigm, individuals work on the judgment task anonymously and in the absence of others and should thus not feel a great need to comply with others (Yaniv & Choshen-Hillel, 2012).

Second, as apparent from the simulation, both the accuracy differences as well as the optimal static weight depend on N . That is, if individuals would not update their judgment on more than 4 instances, almost all static weights yield accuracy comparable to that of Bayesian updating (except 0.1). However, it appears unlikely that judgments are only updated on such rare occasions. For example, literature on decisions from experience documents average sample sizes of 11-19 instances (e.g., Hertwig, 2015; Phillips, Hertwig, Kareev, & Avrahami, 2014). Furthermore, in a digital world that abounds with sources of information easily accessible, opportunities for judgment updating are numerous. Consequently, it appears likely that individuals seek out a number of additional pieces of information that greatly surpasses $N = 1$.

The present notion of adaptive small weights in sequential information uptake applies to the acquisition of objective knowledge as usually used in advice taking paradigms implemented in the lab (e.g., distances between cities, the year of historical events, or the caloric content of foods; e.g., Yaniv, 2004a). The present simulation however has implications for any type of numerical or quantifiable judgment be it in social or non-social domains. For instance, the present notion may also be applied to the formation of accurate impressions of other people's attitudes and beliefs, where it is not the "true" value that needs to be

determined, but rather a representative value that reflects a person's state of mind or a group's position on a given issue.

2.2.1 Boundary conditions

Under what conditions do the present results hold? Outside the laboratory, individuals' judgments likely have already been updated several times, thereby representing more than a single piece of information. However, while the simulation investigated updating along pairings with individual judgments, the same logic (and thus a similar pattern of results) applies when every judgment resembles a (random) amount of information (i.e., a group judgment; Mannes, 2009). Only if the amount of information a given judgment contains and the order by which these judgments are encountered were strongly positively correlated, the here documented advantage of conservative updating would not emerge. That is, only if every judgment encountered was based on as many single judgments as the own updated judgment, both should be weighted equally.

Likewise, while research on the wisdom of crowds reliably documents averaged crowd judgments to be more accurate than individual judgments (e.g., Davis-Stober et al., 2014), the case of a crowd being perfectly accurate (as in the above simulation) might be a somewhat strict conceptualization rarely found outside the laboratory. What if the crowd is less wise than implemented in the present simulation? Note first, that even the accuracy achievable through optimal Bayesian updating is limited by the accuracy of the crowd. As the crowd becomes less wise, the accuracy of both optimal dynamic updating as well as static low weights deteriorates. Naturally, the here advocated advantage of lower weights as compared to higher weights in sequential updating depends on the degree to which averaging of several judgments outperforms choosing a random judgment. Specifically, while lower weights closely resemble true averaging, higher weights lead to an unduly high impact of the last judgment encountered. Thus, if averaging and choosing yielded comparable accuracy, the here documented advantage of lower weights would vanish. However, averaging has been

shown to outperform choosing in a wide variety of environments (Davis-Stober et al., 2014). For every two judgments that enclose or “bracket” the truth, averaging yields a higher expected accuracy than random choosing (Larrick & Soll, 2006). Consequently, the likelihood of crowd judgments bracketing the truth (as an indicator of the wisdom of the crowd) moderates the advantage of lower as compared to higher weights. Note however, that even without bracketing, random choosing generally does not outperform averaging.

The current chapter delivers compelling evidence that the ecological perspective advocated in this thesis has immense implications for existing assessments of the adaptivity of advice taking. Specifically, theoretical arguments and empirical data suggest that in a sequential advice taking process, underweighting of advice in comparison to the own judgment outperforms weighting the current judgment and the advice equally and approximates the accuracy of optimal Bayesian updating. Consequently, the habitual underweighting of advice which individuals display in many advice taking studies might be highly adaptive outside the constrained research paradigm. The next chapters complement this focus on normative considerations by introducing an information integration framework of advice taking and investigating individuals’ actual advice taking behavior in a paradigm that allows for information search.

3. Expanded information asymmetry account

The current thesis proposes an information integration account of advice taking. This account allows to explain advice taking, information search and their relation and can integrate the most central determinants of advice taking as identified by existing research (e.g., confidence; Gino et al., 2012; distance of advice; Moussaïd et al., 2013; Schultze et al., 2015). To this end, the information asymmetry account (Yaniv, 2004a, 2004b) is expanded by the opportunity to integrate additional information. As in the original information asymmetry account, it is assumed that an individual's judgment results from an underlying distribution of information. However, in line with the ecological perspective advocated here, this information sample is not fixed but can be expanded by any additional information received.

The expanded information asymmetry account has two core assumptions. First, it is assumed that any additional information received (e.g., advice) is integrated into the individual's current information sample. Second, individuals are sensitive to the reliability of this integrated information sample, which depends on the amount and consistency of information, as has recently been proposed by research on subjective confidence (Koriat, 2012a).

Note that this does imply that an individual's judgment and advice received are not treated differently per se. Rather, it is assumed that an individual's judgment usually results from an underlying distribution of information (e.g., reasons), while advice classically has been operationalized as a single numerical estimate. Thus, in the absence of any additional information, advice usually exerts a relatively low influence.

The assumption that individuals are sensitive to the reliability (i.e., amount and consistency) of their current information sample necessitates a distinction between two types of information individuals could receive. While the amount of information increases regardless of the type of information, new information can either decrease the consistency of the overall sample, when it is dissimilar from the current judgment, or it can increase the

consistency (or keep it stable), when it is similar to the current judgment. Information that increases the consistency of the integrated information sample (or keeps consistency stable while increasing the amount) should lead to an increase in confidence (Koriat, 2012a). On the other hand, information that decreases the consistency should lead to a shift of the judgment.

This account can thus explain the qualitatively different patterns in advice utilization depending on advice distance. Close advice is more likely to increase the consistency of the integrated information sample (or keep it stable while increasing the amount) and should thus increase confidence, while distant advice decreases the consistency and thus leads to a stronger shift in judgment (Moussaïd et al., 2013; Schultze et al., 2015).

Figure 5 illustrates the integration of two different types of additional information into the current information sample via a simulation. To visualize the amount and consistency of a distribution, the density of that distribution is plotted. The upper panel displays the density of three distributions, the initial information sample ($N = 10$ pieces of information, $M = 500$, $SD = 25$; outline only) and a distribution of close advice ($N = 20$, $M = 530$, $SD = 25$; golden distribution) and distant advice ($N = 20$, $M = 750$, $SD = 25$; red distribution). The lower two panels display the effect of integrating different amounts of information from the close or distant advice distribution into the initial information sample. The resulting integrated information samples are displayed for consecutively adding up to 20 pieces of advice from the close (golden) or distant (red) advice distribution. The vertical lines indicate the means of the respective integrated information samples (dashed for integrating close advice and solid for integrating distant advice). Whereas integration of distant advice greatly decreases the density of the integrated information sample this is not the case for the integration of close advice.

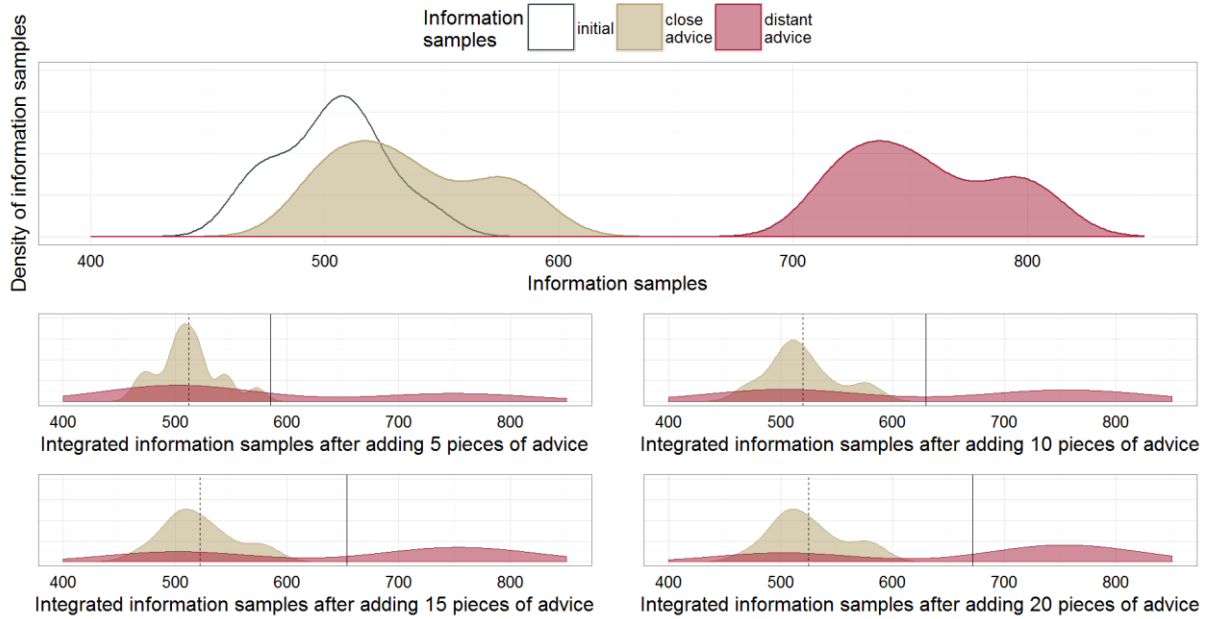


Figure 5. Upper panel: Density of three distributions, the initial information sample ($N = 10$ pieces of information, $M = 500$, $SD = 25$; outline only) and a distribution of close advice ($N = 20$, $M = 530$, $SD = 25$; golden distribution) and distant advice ($N = 20$, $M = 750$, $SD = 25$; red distribution). *Lower panels:* Density of the integrated information sample after integrating different amounts of information from the close advice distribution (golden) or distant advice distribution (red). The vertical lines indicate the means of the respective integrated information samples (dashed for close advice and solid for distant advice).

This account also predicts previously uninvestigated effects of additional information on advice taking. Assuming that additional information is integrated into the current information sample, it follows that the more additional information is integrated, the greater its impact. That is, if three instead of just one close piece of advice were encountered, the increase in confidence should be greater. Likewise, if three instead of one piece of distant advice were encountered, judgment revision should be greater (see lower panel of Figure 5). This effect of the amount of additional information should be moderated by its consistency (Koriat, 2012a). For example, if three advisors closely agreed, their position should exert a stronger impact on the integrated information sample.

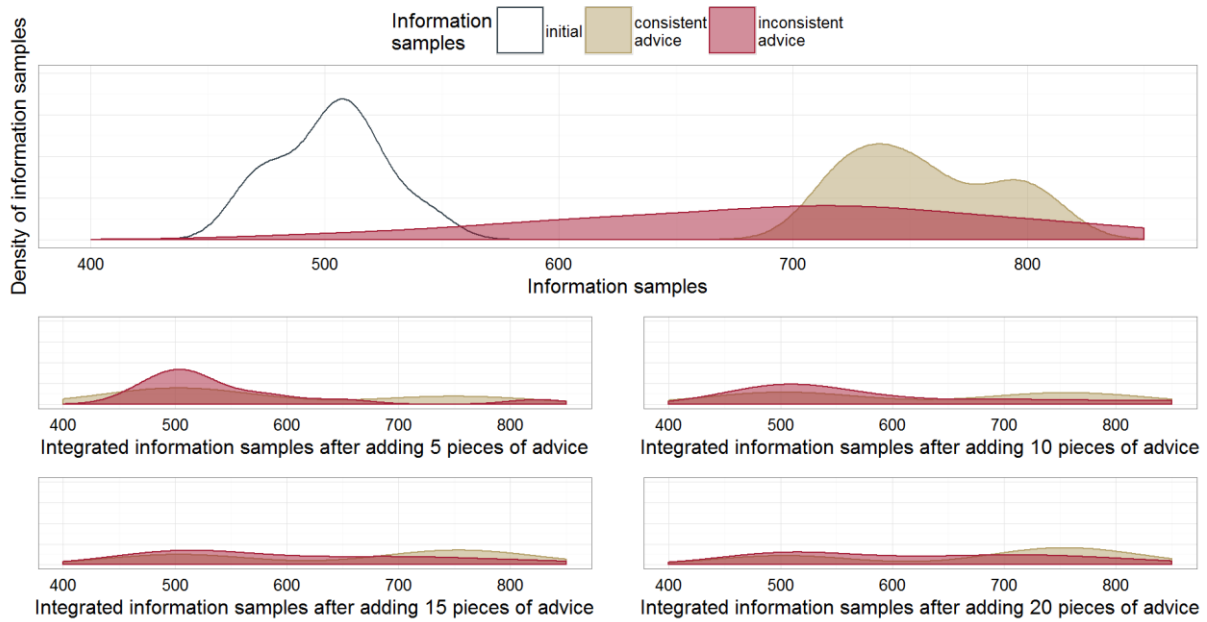


Figure 6. Upper panel: Density of three distributions, the initial information sample ($N = 10$ pieces of information, $M = 500$, $SD = 25$; outline only) and a distribution of advice that is as consistent as the initial information sample (consistent advice; $N = 20$, $M = 750$, $SD = 25$; golden distribution) and advice that is less consistent than the initial information sample (inconsistent advice; $N = 20$, $M = 750$, $SD = 150$; red distribution). *Lower panels:* Density of the integrated information sample after integrating different amounts of information from the consistent advice distribution (golden) or inconsistent advice distribution (red).

Figure 6 illustrates the integration of additional information with two levels of consistency into the current information sample via a simulation. The upper panel displays the density of three distributions, the initial information sample ($N = 10$ pieces of information, $M = 500$, $SD = 25$; outline only) and a distribution of advice that is as consistent as the initial information sample (consistent advice; $N = 20$, $M = 750$, $SD = 25$; golden distribution) and advice that is less consistent than the current information sample (inconsistent advice; $N = 20$, $M = 750$, $SD = 150$; red distribution). The lower two panels display the effect of integrating different amounts of information from the consistent advice distribution (golden) or inconsistent advice distribution (red) on the density of the integrated information sample. Integration of consistent advice leads to a greater shift of density towards the initial advice distribution.

Likewise, the amount and consistency of the current information sample should affect the relative impact of newly encountered information. If the current information sample is small, additional information should have a higher impact than when the current information sample is large (Weber-Fechner law). This could explain effects of a decreased recognizability of the target stimulus on advice weighting (Gino & Moore, 2007; Yaniv & Choshen-Hillel, 2012). In this case, individuals' initial information sample is likely small, thus increasing the impact of additional information (i.e., advice).

Because subjective confidence is assumed to resemble the reliability of the current information sample (Koriat, 2012a) this also provides a link to confidence as a determinant and mediator of advice taking. The smaller or less consistent the current information sample, the lower the associated level of confidence and in turn the higher the impact of additional information (i.e., advice; e.g., Gino et al., 2012; Gino & Moore, 2007; Moussaïd et al., 2013).

The expanded information asymmetry account assumes that individuals will seek additional information when an assessment of their current information sample yields low reliability, that is, when it is small or inconsistent. This should be the case when an individual's confidence is low (Koriat, 2012a). Likewise, this should be the case after integrating a piece of dissimilar advice (Figure 5).

Predictions of the absolute amount of information search in advice taking are difficult, as this research is the first to ever investigate information sampling in the advice taking paradigm. The absolute level of information search likely depends on a trade-off between the costs associated with seeking additional information and the benefits of this additional information (e.g., Bonaccio & Dalal, 2006; Lee, Newell, & Vandekerckhove, 2014). Given that advice taking research usually employs judgment tasks that individuals feel uncertain about (e.g., Gino, 2008) and that over 90% of individuals chose to receive advice when this was not mandatory (Gino et al., 2012; Gino & Moore, 2007), it is expected that individuals will generally seek some additional information when given the opportunity.

The expanded information asymmetry account builds on existing accounts of advice taking (Yaniv, 2004a, 2004b) and subjective confidence (Koriat, 2012a) and expands them by the opportunity to integrate additional information. It is the first account that conceptualizes advice taking as an interaction between the information in individuals' minds and the information ecology. It allows to integrate important existing findings on advice taking and can thus serve as a starting point for the development of a comprehensive theoretical account of advice taking. The next chapters derive and test hypotheses regarding the precursors and consequences of information search in advice taking to assess the validity of this theoretical framework.

4. Exploring the information ecology: Investigating information search in advice taking

The appeal of advice taking as a strategy to improve judgment accuracy rests on the two assumptions that averaging several judgments increases accuracy (e.g., Davis-Stober et al., 2014; Einhorn et al., 1977; Galton, 1907) and that many decision situations are inherently interactive (Larrick et al., 2012; Snizek & Buckley, 1995).

However, research on advice taking has hardly ever satisfied the assumption that the investigated decision situation is interactive. Specifically, advice taking research has almost exclusively focused on the weighting of advice once it is received (e.g., Harvey & Fischer, 1997; Mannes, 2009). While two studies investigated whether advice was solicited at all (Gino et al., 2012; Gino & Moore, 2007) no research has yet investigated the process of advice seeking (Bonaccio & Dalal, 2006; Gino et al., 2012). The existing research on advice taking thus never really allowed individuals to interact with their information ecology. Consequently, this research does not allow any inference with regard to individuals' advice taking behavior in a decision situation that truly is interactive.

To overcome this constraint, the current chapter establishes an advice taking paradigm that allows individuals to sample any number of other people's judgments before revising their judgment. While individuals generally could seek additional information in the form of both advice or underlying reasons, the current thesis focuses on the sampling of additional advice for three reasons. First, in everyday life the judgments of other people can be observed more frequently than the reasons underlying a judgment. Second, when judgments are based on intuition, concrete reasons are often not available or constructed after the fact (Nisbett & Wilson, 1977). Third, it is assumed that a judgment is evaluated in accordance with the amount and consistency of information as in an assessment of reliability, a principle that should be applicable to the number and consistency of advisory judgments as well as the number and consistency of supporting reasons (Koriat, 2012a).

The current thesis is the first to investigate advice taking from an ecological perspective, covering both advice seeking and weighting. This approach is necessary for the proper evaluation of the adaptivity of advice taking. Specifically, given that the information in individuals' minds is assumed to be skewed towards supporting the own judgment (Yaniv, 2004a, 2004b), it might be considered rational for individuals to underweight advice received, if this advice constitutes a single piece of information (Ravazzolo & Røisland, 2011). However, individuals should adapt their advice weighting to the information provided by the information ecology. If, for example, several advisors closely agreed on a position, this increased amount of information in support of their position should increase reliance on their position.

This approach also constitutes a necessary first step to the development of a comprehensive theoretical account. Specifically, the amount and consistency of individuals' internal information samples can only be approximated, rendering it difficult to test whether the information individuals have to begin with determines advice weighting (for an attempt see Soll & Mannes, 2011). However, the expanded information asymmetry account predicts that additional information received is integrated into this information sample. The amount (and consistency) of additional information should thus directly influence advice weighting. This assumption can easily be tested. The goal of the current chapter is thus twofold. First, by investigating advice taking in a paradigm that allows for advice seeking, this research is the first to investigate whether advice taking behavior is sensitive to the information provided by the information ecology. Second, it serves as an initial test of the here proposed theoretical framework of advice taking.

4.1 Overview of experiments and hypotheses

This chapter presents three experiments that followed a general procedure. As in the classical approach to advice taking, participants were presented with an estimation task, in this case the caloric content of various foods per serving. After giving an initial estimate,

participants received an initial piece of advice, supposedly sampled from a distribution of 100 participants in an earlier study. This piece of advice was either similar to or dissimilar from their initial estimate. All three experiments implemented this factor in a within-participants design. Experiment 1 resembled the classical research approach without a sampling phase and served as an initial test of the weighting functions of the operationalization of close and distant advice. In Experiments 2 and 3, participants could sample additional pieces of advice at will before giving their final estimate. While Experiment 2 implemented advice that was consistent among advisors, Experiment 3 utilized ecologically valid and thereby considerably less consistent advice.

To assess whether the expanded information asymmetry account as a framework is suited to explain information search and its consequences in advice taking, the following hypotheses regarding advice seeking and weighting will be tested in the three upcoming experiments:

H1: The frequency of advice sampling will be higher when advice is distant from as compared to close to participants' initial estimates.

H2a: Advice weighting will be higher when advice is distant from as compared to close to participants' initial estimates.

H2b: Advice weighting will be higher the more pieces of advice participants sampled.

H3: Confidence will be lower when advice is distant from as compared to close to participants' initial estimates.

4.2 Data analysis

As dependent measures, the size of the self-determined samples of advisory estimates, the integration of advice when forming a final judgment and confidence regarding this final judgment were analyzed. As a measure of advice integration, the weight of advice was calculated (WOA; Harvey & Fischer, 1997; Yaniv, 2004a), defined as

$$WOA_{ij} = (F_{ij} - I_{ij}) / (A_{ij} - I_{ij}),$$

where I , F , and A indicate the initial estimate, final estimate, and advice, respectively, on a given trial j in a given participant i . To arrive at an estimate for A , the mean of all pieces of advice received was calculated. The WOA thus reflects the degree to which participants move towards the advice (i.e., a WOA of 1 indicates adoption of the advice; a WOA of 0 indicates adherence to the initial estimate; a WOA of 0.5 indicates an equal-weighting strategy). The WOA thus allows for an assessment of (increasing or decreasing) receptivity towards advice as a function of amount sampled or other variables.

The WOA is highly sensitive to outliers. Outliers primarily result from trials where the distance between initial estimate and advice is very small, so that even small alterations of the judgment may lead to values outside the range of 0 and 1. Many researchers thus truncate the WOA, setting values smaller than 0 to 0 and values greater than 1 to 1 (e.g. Gino, 2008; Schultze et al., 2015; Soll & Larrick, 2009), arguing that this practice is unproblematic if it concerns less than 5% of all trials. As participants in the expanded paradigm (Experiment 2 and 3) received multiple pieces of advice that scattered, the WOA fell outside the range of 0 and 1 on more than 5% of all trials. For the analyses reported, statistical criteria were applied to identify and remove outliers on a trial-by-trial basis (Tukey, 1977) rather than altering values of WOA. Additional analyses applying the truncation practice yielded similar results, but led to a higher number of alterations in Experiment 2 and 3. Note that for the influence of sampling on advice weighting, the qualitative pattern of results is unaltered even when no trials are excluded. However, the effect of advice distance on WOA reverses (Experiment 1 and 3) in the absence of any exclusion.

For all experiments in the current thesis, multilevel model analyses were performed for all dependent measures to assess relationships on a trial-by-trial basis (Judd, Westfall, & Kenny, 2012). Analyses are based on 20 data points per participant. Following the recommendations by Judd and colleagues (2012), all models contained random intercepts for

participants as well as items, which were fully crossed by design. The effects of the hypothesized predictors were always fixed. The best-fitting model was determined by sequentially including the hypothesized parameters and their interactions, one at a time and inspecting increase in model fit using the Kenward-Roger approximation (Kenward & Roger, 1997) for degrees of freedom. Once the best-fitting model was determined, those interactions that did not reach a t -value ≥ 2 were excluded (Baayen, Davidson, & Bates, 2008), as long as this did not decrease model fit significantly. The resulting p -values as well as log-likelihood values and approximated Bayes factors (Masson, 2011) for each model are given in the Appendix. For the sake of brevity, only the parameter estimates of the best-fitting model are reported.

To summarize, the expanded information asymmetry account predicts that sampling behavior will increase with the distance of advice from the participant's initial estimate (H 1). This assumption will be tested in Experiments 2 and 3, by looking for a main effect of advice distance on sample size in the multilevel model. A larger sample size for distant as compared to close advice is predicted. The expanded information asymmetry account also predicts advice weighting to increase when advice is distant and when participants sampled more advice. Specifically, it predicts two main effects of advice distance (H2a; Experiments 1-3) and amount of information sampled (H2b; Experiments 2 and 3) on the WOA in a multilevel model, such that the WOA will be larger with distant (rather than close) advice and larger (rather than smaller) samples of advisory estimates. The expanded information asymmetry account also predicts that close advice leads to higher confidence than distant advice (H3). This assumption will be tested in Experiments 2 and 3, by looking for a main effect of advice distance on confidence in the multilevel model. A lower confidence for distant as compared to close advice is predicted.

4.3 A-priori power analyses

A-priori power analyses were conducted to determine required sample sizes in all three experiments (Faul, Erdfelder, Lang, & Buchner, 2007). As there is no clear guideline for power analyses regarding multilevel modelling utilized in the present research, these calculations were based on repeated measures ANOVA designs. Detecting a medium-sized effect ($f = .25$) with sufficient power ($\beta = .95$) in Experiments 1 and 2 required collecting data of at least 20 participants. A smaller effect ($f = .15$) was assumed in Experiment 3 due to the increased variance in advice, requiring at least 54 participants. As the experiments were part of multi-experiment sessions, sample size was increased in accordance with other studies that required a larger sample size.

4.4 Experiment 1

The first experiment served as a conceptual replication of advice taking procedures that do not include a sampling phase using the same materials and basic manipulation as in the next two experiments. The expanded information asymmetry account assumes a qualitatively different pattern in advice utilization depending on the distance of advice from the participant's initial estimate. Specifically, advice that is close to the participant's initial estimate is more likely to increase the consistency of the integrated information sample or keep this consistency stable while increasing the amount of information. Thus, such advice should lead to an increase in confidence rather than judgment revision (e.g., Moussaïd et al., 2013; Schultze et al., 2015). On the other hand, advice that is distant from the participant's initial estimate likely decreases the consistency of the integrated information sample. Such advice should thus lead to judgment revision rather than an increase in confidence (e.g., Moussaïd et al., 2013; Schultze et al., 2015). Consequently, the following experiments implemented two advice distance conditions. The close condition was intended to confirm participants' initial judgments leading to little judgment revision. In the distant condition,

advice was intended to increase participants' tendency to revise their judgments, as in the intermediate distances in Moussaïd et al.'s (2013) and Schultze et al.'s (2015) work.

4.4.1 Method

Participants. The sample consisted of 35 University of Tübingen students of different subjects (10 males; $M_{age} = 23.00$ years, $SD_{age} = 6.18$).³ Participants took part in a one-hour session comprised of several social cognitive experiments and were compensated with either course credit or 7 € and a chocolate bar. Additionally, participants received a small bonus payment depending on task performance. Participants were recruited via e-mail and online social networks.

Design. The experiment implemented a 2 (distance of advice: close vs. distant) \times 2 (judgment phase: initial vs. final judgment) within-participants design.

Materials and procedure. After participants signed a consent form, they were asked to estimate the caloric content of various dishes per serving. Names, pictures and caloric content of dishes were retrieved from a web page of a German nutrition magazine (www.essen-und-trinken.de). 30 dishes were pretested using an online sample of 21 participants. Out of those 30 dishes, those 20 dishes were selected for which the average estimate most closely resembled the true caloric content of the dish (e.g., fish pasta). This procedure assured that the population had some knowledge about the caloric contents of those dishes (e.g., Gino, 2008).

The order of dishes was randomized for each participant, as was the order of the 10 trials per distance condition across the 20 trials.

For each trial, participants first saw a picture of the dish along with a descriptive label and a response box to give their initial estimate. If participants submitted an estimate lower than 50 or higher than 1250 calories per serving, the software prompted them to give a plausible estimate. Next, they were shown another estimate for that same dish. These

³ Due to a programming error, data of five additional participants were not recorded.

estimates were allegedly drawn from a pool of estimates of 100 participants in a preceding experiment. Depending on distance condition, the advice was simulated as a pseudo-normal distribution centering on a mean relatively close to the participant's initial estimate (on close trials) or centering on a mean relatively distant from the participant's initial estimate (on distant trials). Specifically, the mean of the distribution was simulated as $i \pm .5 \times t$ (on distant trials) and $i \pm .05 \times t$ (on close trials) where i is the participant's initial estimate and t is the true value of the dish. If the participant underestimated in comparison to the true value, the proportion was added to the initial estimate; if the participant overestimated in comparison to the true value, the proportion was subtracted. Consequently, the advice would generally point in the direction of the true value. However, if the initial estimate was very close to the true value, the advice could also lead participants away from the true value, by causing them to adjust too much. The standard deviation of the distribution was equal for close and distant advice. Random noise was added to obscure the artificial nature of advice, drawn from a uniform distribution in the range of [-8, +8] calories.

After receiving advice, participants gave a final estimate for the given dish. For each of these final estimates that fell in a range extending 10% on either side of the true caloric content of the dish, participants would receive a bonus payment of 0.10 €. During the experiment, participants did not receive feedback on their accuracy or their cumulative bonus. The bonus was revealed only at the end of the experiment. In total, participants could collect a bonus of 2 €. Finally, participants were probed for suspicion, debriefed, paid, and thanked.

4.4.2 Results

Distance of advice. As a manipulation check, the absolute difference between participants' initial estimates and the initial piece of advice on the level of trials was assessed. Whereas participants' initial estimate and the initial piece of advice differed by only $M = 30.80$ ($SD = 21.19$) calories on close trials, they differed by $M = 241.85$ ($SD = 83.44$) calories on distant trials.

The normalized distance in line with Moussaïd et al. (2013) was also calculated:

$$\Delta E_{ij} = |I_{ij} - A_{ij}| / I_{ij},$$

where I and A denote the initial estimate of the participant and the advice, respectively, on a given trial j in participant i . Moussaïd et al. (2013) define values below 0.3 as “similar” advice, which they found to be related to an updating of confidence rather than an updating of one’s judgment. Values between 0.3 and 1.1 were defined as an “intermediate distance”, for which strong judgment updating was observed. Values above 1.1 were classified as “very distant” and were empirically related to diminishing influence on participants’ judgments. Employing the average advice on each trial as the standard of reference, the average ΔE amounted to $M = .07$ ($SD = .06$) in the close condition and to $M = .53$ ($SD = .28$) in the distant condition. The operationalizations of the close and distant conditions thus resemble the “similar” and “intermediate” distance conditions of Moussaïd et al. (2013).

As a result of the distance manipulation, participants’ initial estimate and the advice enclosed the true value in 79% of the distant trials, but only 17% of the close trials.

Weighting of advice. Trials with a $WOA < -0.42$ and $WOA > 0.67$ were excluded (Tukey, 1977). In total, 49 of 700 trials were excluded (7.00%).

The expanded information asymmetry account predicts that advice weighting is greater when advice is distant from as compared to close to a participant’s initial estimate.

The following model was fit for advice weighting:

$$WOA_{ij} = \alpha_0 + b_1 Distance_j + \varepsilon,$$

where index i refers to participants and index j to questions. For this and all following models close advice was coded as 0 and distant advice as 1. The effect of *Distance* thus indicates the consequences of increasing advice distance from close to distant advice.

Replicating the results of other authors, advice distance significantly increased participants’ WOA , $b_1 = .12$, $se = .01$, $t = 10.65$ (Figure 7), although fit indices indicated worse fit for this

model over the null model in this specific data set. In all subsequent experiments, including distance as a predictor of *WOA* improved model fit (Appendix).

4.4.3 Discussion

Participants weighted the advice stronger when it was distant from rather than close to their initial estimates. These results validate the manipulation of advice distance replicating the finding of greater weighting for advice of intermediate rather than small distance (e.g., Moussaïd et al., 2013; Schultze et al., 2015).




Figure 7. Weight of advice (WOA) for close and distant trials of Experiments 1-3. Error bars represent standard errors of the means.

4.5 Experiment 2

The second experiment introduced a sampling phase that allowed participants to sample as many pieces of advice as desired. This experiment allowed to assess whether distant advice indeed leads to an increased frequency of advice sampling. Furthermore, this experiment allowed to investigate whether advice weighting is sensitive to the number of pieces of advice received. In order to investigate whether close advice as operationalized here actually led to higher levels of confidence, a confidence measure was introduced after each

final estimate. That was done for only half of the participants to assess potential effects of reflecting on one's confidence.

4.5.1 Method

Participants. The sample consisted of 44 University of Heidelberg students (14 males; $M_{age} = 26.48$ years, $SD_{age} = 8.81$). The study was run as part of an experimental session of approximately 45 minutes. Participants received either course credit or a monetary compensation of 6 €. Additionally, participants received a performance-contingent bonus as in Experiment 1.

Design. This experiment implemented a 2 (distance of advice: close vs. distant) \times 2 (judgment phase: initial vs. final judgment) within-participants design. Additionally, half of the participants [$N = 22$] were asked to provide confidence ratings for their final estimates.

Materials and procedure. Up to eight participants took part at the same time, each seated in a separate cubicle. Materials and procedure were identical to Experiment 1 with two exceptions. First, after receiving a close or a distant piece of advice on the caloric content of a given dish, participants were allowed to sample as many other estimates as desired. The sampling phase was however capped at 20 pieces of advice. Advice was consistently drawn from either the close or distant distribution depending on distance condition. Second, participants in the confidence ratings condition also rated their confidence on a 7-point scale anchored at -3 = "very unconfident" and 3 = "very confident" after each trial while participants in the no confidence ratings condition proceeded without any further questions. As the presence or absence of confidence ratings did not exert effects on any dependent measure, this factor will not be discussed any further.

4.5.2 Results

Distance of advice. As a manipulation check, the absolute difference between participants' initial estimates and the average advice for each trial was assessed. Whereas

participants' initial estimate and the advice differed by only $M = 26.34$ ($SD = 15.75$) calories on close trials, they differed by $M = 246.18$ ($SD = 77.57$) calories on distant trials.

The ΔE to the average advice amounted to $M = .06$ ($SD = .06$) in the close condition and to $M = .54$ ($SD = .26$) in the distant condition. Hence, a “similar” and “intermediate” distance condition in terms of Moussaïd et al. (2013) were realized.

As a result, participants' initial estimate was enclosed by the range of advice for 51% of close trials but for none of the distant trials. Initial estimate and range of advice enclosed the true value on 22% of the close and 84% of the distant trials.

Sampling of advice. As this is the first research to investigate advice seeking, the absolute number of pieces of advice sampled was assessed in order to back up the assumption that participants are generally willing to seek additional information in the advice taking paradigm. In support of this assumption, participants sampled more pieces of advice than the one piece classically offered. The average of $M = 8.05$ ($SD = 7.00$) advisory estimates sampled by participants across distance conditions was significantly larger than one, $t(43) = 7.27, p < .001, d = 1.10$.

Further, the expanded information asymmetry account predicts an increase in sampling frequency when advice is distant. For advice sampling, the model with the best fit was the following:

$$\text{Sample size}_{ij} = a_0 + b_1 \text{Distance}_j + \varepsilon.$$

Advice distance significantly increased sample size, $b_1 = 1.53, se = .20, t = 7.82$. That is, receiving advice that was distant from their initial estimate led participants to consider about 1.5 pieces of advice more than when receiving advice that was close to their initial estimate (Figure 8).




Figure 8. Sample size of advisory estimates for close and distant trials of Experiments 1-3. Error bars represent standard errors of the means.

Weighting of advice. Trials with $WOA < -0.90$ and $WOA > 1.40$ were excluded, amounting to 41 of 880 trials (4.66%; Tukey, 1977). Consistent with Experiment 1, distant advice was expected to lead to greater reliance on advice. Testing the assumption of sensitivity to the sampled information, increased sampling was expected to increase reliance on advice. The following model had the best fit:

$$WOA_{ij} = \alpha_0 + b_1 Distance_j + b_2 Sample\ size_{ij} + \varepsilon.$$

Replicating the results of Experiment 1, advice distance increased participants' WOA, $b_1 = .23$, $se = .02$, $t = 11.21$ (Figure 7). Sampling of additional advice increased the WOA as indicated by the effect of sample size, $b_2 = .015$, $se = .003$, $t = 5.66$.⁴

The analysis of advice weighting strongly supports the assumption of increased reliance on an advisory opinion when it is supported by additionally sampled, consistent pieces of advice. To complement the traditional WOA analysis, it was also assessed how the weighting strategy employed by participants fit with the normative rule of equal weights

⁴ For all models containing sample size as a predictor, sample size was recoded by subtracting 1 to account for the fact that each participant received at least one piece of advice on each trial.

averaging (Mannes, 2009). Therefore the normative WOA participants should have assigned to the mean advice on a given trial was calculated as $N_{advice} / (N_{advice} + 1)$, where N_{advice} is the number of advisory estimates received on a given trial (Mannes, 2009), and subtracted from the actual WOA participants assigned to receive a ΔWOA . For instance, if on a given trial a judge received nine pieces of advice, from an equal weighting perspective, the normative WOA should be 0.9. If the actual WOA on that trial is 0.6, the ΔWOA would amount to -0.3, indicating that the judge underweighted the advice. A positive score in this case implies overweighting of advice in comparison to the normative rule of equal weights averaging. The best fitting model for this deviation score was the following:

$$\Delta WOA_{ij} = \alpha_0 + b_1 Distance_j + b_2 Sample\ size_{ij} + b_3 Distance_j \times Sample\ size_{ij} + \varepsilon.$$

Replicating numerous studies on failure to follow the normative rule of equal weights averaging (e.g., Mannes, 2009), participants fall short of this rule on both, close and distant trials (Figure 9). However, assigned weights more closely matched the normative rule when advice was distant rather than close, $b_1 = .17$, $se = .03$, $t = 5.81$. Although sampling increased the weight participants placed on an advisory opinion, it did so suboptimally, as the deviation from the normative rule increased with increasing sample size, $b_2 = -.012$, $se = .003$, $t = -3.81$. This effect was more pronounced for sampling on close as compared to distant trials, as indicated by the interaction, $b_3 = .008$, $se = .003$, $t = 2.78$.




Figure 9. Deviation of participants' weight of advice (WOA) from equal weights averaging for close and distant trials of Experiments 1-3. Error bars represent standard errors of the means.

Confidence. In line with other authors (Moussaïd et al., 2013; Schultze et al., 2015), it was assumed that receiving close as compared to distant pieces of advice would lead to higher levels of confidence. The following model had the best fit:

$$Confidence_{ij} = \alpha_0 + b_1 Distance_j + \varepsilon.$$

Indeed, increasing distance of advice decreased confidence, $b_1 = -.63$, $se = .07$, $t = -9.28$ (Figure 10). Additional sampling of consistent advice did not affect confidence as indicated by the missing increase in model fit when including sample size as an additional predictor (Appendix).




Figure 10. Confidence after receiving advice for close and distant trials of Experiments 1-3. Error bars represent standard errors of the means.

4.5.3 Discussion

Allowing participants to sample additional pieces of advice, the results of Experiment 2 testify to participants' willingness to obtain additional information. Further, the effect of advice distance on sample size demonstrates that receiving dissimilar as compared to similar pieces of advice increased participants' sampling frequency. In line with the expanded information asymmetry account, participants were sensitive to their information samples, shifting more strongly towards the advisory opinion the higher the number of consistent pieces of advice supporting this opinion they sampled.

Further validating the distance manipulation as providing information supporting or diverging from participants' initial judgments, confidence was significantly lower on distant as compared to close trials.

4.6 Experiment 3

The third experiment was intended to investigate the effect of advice that is less consistent than in the previous experiments. Specifically, the advice participants sampled in Experiment 2 varied only slightly across advisors on a given trial. To increase the variance of

the advice received and also increase the ecological validity of advisory estimates, the procedure for computing the advice was altered in this experiment.

4.6.1 Method

Participants. The sample consisted of 58 University of Tübingen students (13 males; $M_{age} = 21.10$ years, $SD_{age} = 3.34$)⁵, who took part in exchange for course credit or a monetary compensation of 7 € and a chocolate bar. The experiment was part of a session that lasted about 60 minutes. Additionally, participants received a performance-contingent bonus as in the previous experiments.

Design. This experiment implemented a 2 (distance of advice: close vs. distant) \times 2 (judgment phase: initial vs. final judgment) within-participants design.

Materials and procedure. Materials and procedure were identical to Experiment 2, with two exceptions. First, the generation of advice differed in the following way. For each dish, the distribution of advice consisted of 100 values drawn at random from a normal distribution with the mean and standard deviation of the initial estimates from participants in Experiment 2. Depending on distance condition (close vs. distant advice) participants received an initial piece of advice that was either drawn from the same decile as their initial estimate (on close trials) or from the fifth next decile, above or below their initial estimate, depending on which decile their initial estimate fell into (on distant trials). For instance, if a participant's initial estimate fell into the 6th decile, she would receive an initial piece of advice from the 6th decile (*close trials*) or from the 1st decile (*distant trials*). All additional pieces of advice received during the sampling phase were randomly drawn from the distribution of advice for the respective dish, independent of distance condition. The mean standard deviation of the distributions of advice experienced by participants was $M = 163.61$ with a range of [125.08; 214.29]. Thus, the advice received was more inconsistent than in Experiment 2, where the

⁵ Due to a programming error, data of two additional participants were not recorded.

average standard deviation of the sampled advice was $M = 24.93$ with a range of [6.36; 46.67]. As a second deviation from Experiment 2, all participants provided confidence ratings for their final estimates.

4.6.2 Results

Distance of advice. As a manipulation check, the absolute difference between participants' initial estimates and the initial piece of advice for each trial was assessed. Whereas participants' initial estimate and the initial piece of advice differed by only $M = 24.88$ ($SD = 29.28$) calories on close trials, they differed by $M = 274.05$ ($SD = 76.82$) calories on distant trials. In contrast, the distance between participants' initial estimate and the average advice received, if participants chose to sample, converged as only the distance of the first piece of advice was manipulated (close: $M = 98.95$, $SD = 85.11$; distant: $M = 148.53$, $SD = 104.71$).

Likewise, the average ΔE between participants' initial estimate and the first piece of advice amounted to $M = .06$ ($SD = .08$) in the close condition and $M = .60$ ($SD = .32$) in the distant condition. Hence, a "similar" and "intermediate" distance condition in terms of Moussaïd et al. (2013) were realized. Averaging across all pieces of advice, ΔE becomes more similar between conditions with $M = .24$ ($SD = .32$) in the close condition and $M = .33$ ($SD = .32$) in the distant condition.

As a result of the greater variance of advice compared to the previous experiments, participants' initial estimate was included in the range of advice for 80% of close trials and for 77% of distant trials. Participants' initial estimate and the sampled advice enclosed the true value on 83% of close and 95% of distant trials.

Sampling of advice. Participants again sampled significantly more advisory estimates than the one piece received in previous research, $M = 10.26$ ($SD = 5.98$), $t(57) = 10.52$, $p < .001$, $d = 1.77$. The following model had the best fit for sample size:

$$\text{Sample size}_{ij} = a_0 + b_1 \text{Distance}_j + \varepsilon.$$

Although only the first piece of advice differed as a function of the distance manipulation, advice distance increased sample size, $b_1 = .87$, $se = .18$, $t = 4.83$ (Figure 8).

Weighting of advice. Trials with $WOA < -0.65$ and $WOA > 1.03$ were excluded, amounting to 119 of 1160 trials (10.26%; Tukey, 1977). The following model had the best fit for WOA:

$$WOA_{ij} = \alpha_0 + b_1 Distance_j + b_2 Sample\ size_{ij} + \varepsilon.$$

Increasing advice distance increased participants' WOA, $b_1 = .06$, $se = .02$, $t = 3.64$, (Figure 7). Albeit being less consistent among advisors than in Experiment 2, sampling of advice again increased the WOA, $b_2 = .007$, $se = .002$, $t = 3.40$.

The traditional WOA analysis was again complemented by investigating participants' behavior in comparison to the normative rule of equal weights averaging (see Experiment 2). The model with the best fit was:

$$\Delta WOA_{ij} = \alpha_0 + b_1 Distance_j + b_2 Sample\ size_{ij} + \varepsilon.$$

As in Experiment 2, participants fell short of the normative rule on both close and distant trials (Figure 9), but came closer when advice was distant, $b_1 = .05$, $se = .02$, $t = 3.23$. Again, while sampling increased the WOA, it did so suboptimally, $b_2 = -.009$, $se = .002$, $t = -4.19$.

Confidence. For confidence, the following model had the best fit:

$$Confidence_{ij} = \alpha_0 + b_1 Distance_j + b_2 Sample\ size_{ij} + \varepsilon.$$

Replicating the findings of Experiment 2, advice distance decreased confidence, $b_1 = -.26$, $se = .05$, $t = -5.24$ (Figure 10). Additional sampling of advice further reduced confidence, $b_2 = -.035$, $se = .008$, $t = -4.46$.

4.6.3 Discussion

Experiment 3 replicated the findings of Experiment 2 under ecologically more valid conditions. Even though the distance of advice was manipulated via a single instance of advice only, initially receiving one dissimilar as compared to similar piece of advice led

participants to sample more advice and to weight advice more strongly. Despite being less consistent among advisors than in Experiment 2, sampling of additional advice increased reliance on advice. The effect of advice distance on confidence was also replicated. Moreover, in contrast to the previous experiment, sampling of advice generally reduced confidence, as advice was less consistent.

4.7 Interim discussion

The above experiments are the first to investigate advice taking from an ecological perspective, covering both advice seeking and weighting. In accord with the ecological perspective advocated in the current thesis, the experiments demonstrate that individuals are generally willing to seek additional information (here: additional advice) in the advice taking paradigm and that this additional information affects advice utilization.

The experiments provide initial support to the expanded information asymmetry account. As predicted, information sampling was more pronounced when advice was distant from participants' initial estimates (H1). This advice likely decreased the consistency of participants' information samples. In further support of this assumption, distant advice also led to greater judgment revision as compared to close advice (H2a), whereas confidence was greater after receiving close advice (H3). Most importantly, participants' advice weighting was sensitive to the information samples created through seeking advice. The more pieces of advice a participant sampled, the higher the reliance on this advice during judgment revision (H2b). This strongly supports the assumption that participants were sensitive to the amount of information received.

A post-hoc comparison of Experiment 2 (coded as 0) and 3 (coded as 1) can serve to investigate the effect of consistency of information. When advice was less consistent among advisors, additional sampling increased advice weighting to a smaller extent, $b_{Sample\ size * Experiment} = -.008$, $se = .003$, $t = -2.51$. Furthermore, participants' sampling frequency on close trials was higher when advice was less consistent, $b_{Experiment} = 2.54$, $se = 1.16$, $t = 2.18$;

$b_{\text{Experiment} \times \text{Distance}} = -0.67$, $se = 0.267$, $t = -2.49$ (Figure 8). Finally, the integration of additional information that was less consistent led to a general decrease in confidence, as indicated by the negative effect of sampling on confidence in Experiment 3. It can thus be assumed that participants were not only sensitive to the size of their integrated information samples but also to their consistency.

The current chapter also complements the previous chapter's assessment of the adaptivity of advice taking. In line with the assumption that information search renders advice taking a sequential process with several instances of updating, participants sequentially sampled on average 9.30 ($SD = 6.53$; across Experiment 2 and 3) pieces of advice for a given judgment task. Thus, a single instance of updating does not appear to capture advice taking in an interactive setting. At a first glance, the findings of the experiments and the simulation might appear to be at odds. Specifically, the experimental data revealed an increase in advice weighting after sampling additional advice. At the same time, the simulation concludes that lower weighting of advice yields higher accuracy. In light of this conclusion, the increase in advice weighting after sampling additional advice might appear to be maladaptive. However, there is a procedural difference between the experimental and simulation design that helps reconcile this conflict. The simulation investigated several steps of sequential updating in light of single pieces of advice. In contrast, in the experiments, participants reported only one updated judgment, after sampling a self-determined amount of advice. Note that consecutively assigning a low weight to each single piece of advice yields an increase in the weight on the mean of these pieces of advice with an increase in the number of pieces of advice. Thus, optimal Bayesian integration demands a higher weight to be placed on a mean that represents more pieces of advice. Therefore, the increased weight placed on the mean advice after sampling of additional pieces of advice does not constitute a violation of the low weighting of advice that yielded high accuracy in the simulation. While advice discounting appears a (rational) prior in many advice taking situations, participants were willing to explore the

information ecology and increase reliance on the advice when it was supported by additional information. Underweighting of advice is thus not an unalienable feature of advice taking but, at least partly, an artefact of the classical research paradigm.

The next chapter will focus on participants' confidence before and after receiving advice to provide converging evidence for the expanded information asymmetry account. At the same time, it will be assessed whether the adaptivity of advice taking as documented here generalizes to environments where information search is costly.

5. Integrating the information ecology: Converging evidence and potential boundaries

One important determinant of individuals' advice weighting is the confidence in their own judgment (Gino et al., 2012). Individuals weight advice more strongly the less confident they are (e.g., Gino et al., 2012; Gino & Moore, 2007; Moussaïd et al., 2013) and confidence was suggested to mediate the effect of other moderators (e.g., anxiety; Gino et al., 2012). Furthermore, confidence can also be the outcome of advice taking. Specifically, individuals increase their confidence after receiving advice that is similar to their own judgment (e.g., Moussaïd et al., 2013; Schultze et al., 2015; Yaniv et al., 2009).

However, there is no existing theoretical framework of advice taking that can explain how confidence affects advice weighting and is affected by advice taking. The current thesis proposes an expanded information asymmetry account that allows to integrate the above reviewed findings regarding confidence as a determinant and outcome of advice taking. Drawing on recent research on subjective confidence (Koriat, 2012a) it is assumed that confidence reflects an individual's assessment of the reliability of its information sample which in turn depends on the amount and consistency of information. It is further assumed that additional information encountered is integrated into this information sample.

This chapter strives to provide converging evidence for the expanded information asymmetry account. To this end, individuals' confidence will be assessed before and after receiving advice. If individuals really integrate additional advice received into their information samples, confidence should change as a function of the type and amount of information encountered. Specifically, confidence should increase more strongly after receiving advice that is similar to as compared to dissimilar from the own judgment. This increase in confidence should be greater the more pieces of advice individuals received.

Note that it is difficult to predict how confidence changes when integrating advice that is dissimilar from the own judgment. An initial piece of dissimilar advice decreases the

consistency of the integrated information sample. However, depending on the original size of that information sample, several pieces of advice that deviate from the own judgment but are highly consistent among each other could lead to an overall increase in the consistency of the integrated information sample. While similar advice should linearly increase an individual's confidence in his or her judgment, the relation between integration of dissimilar advice and confidence could be u-shaped. Critically, the change in slope from decreasing to increasing confidence should be determined by the size and consistency of the individual's initial information sample. As these initial information samples cannot be measured, it is difficult to predict whether individuals' confidence will decrease or increase as a function of integrating dissimilar advice. Given that the judgment tasks employed were selected so that individuals have only vague knowledge (e.g., Gino, 2008) it could be expected that even dissimilar advice leads to an increase in confidence.

While the exact amount and consistency of individuals' integrated information samples cannot be measured, smaller or less consistent information samples should be related to lower confidence. The smaller or less consistent an individual's information sample is, the more strongly it should be affected by additional information (Weber-Fechner Law). Consequently, the lower an individual's confidence, the higher the expected impact of additional information on advice weighting and confidence change.

Furthermore, the expanded information asymmetry account predicts that individuals are more likely to seek additional information when their integrated information sample is less reliable. While the experiments in the previous chapter provided support for this assumption by documenting an increase in advice sampling for distant as compared to close advice, the following experiments allow to assess the hypothesis that lower initial confidence predicts increased advice sampling.

The current chapter focuses on confidence as a determinant and outcome of advice seeking and taking. This constitutes an additional test of the assumptions of the here proposed

expanded information asymmetry account. Furthermore, it illustrates how the expanded information asymmetry account can integrate findings from existing research on advice taking.

5.1 Overview of experiments and hypotheses

The two experiments reported below followed the same general procedure as the experiments of the previous chapter. In both experiments, participants estimated the airline distance between several pairs of European cities. After giving an initial estimate, participants received one piece of advice. Across trials (i.e., within participants) this advice was either similar to or dissimilar from participants' initial estimates. Next, participants could sample additional advisory estimates before giving a final estimate. While sampling in Experiment 1 was free, Experiment 2 implemented costs of sampling as a between-participants factor.

To further assess the validity of the expanded information asymmetry account, the hypotheses investigated in the previous chapter will be tested again. Given this chapter's focus on confidence, these hypotheses will be complemented by additional hypotheses (bold) regarding the role of confidence as a determinant and outcome of advice seeking and weighting.

H1a: The frequency of advice sampling will be higher when advice is distant from as compared to close to participants' initial estimates.

H1b: The frequency of advice sampling will be higher the lower participants' initial confidence.

H2a: Advice weighting will be higher when advice is distant from as compared to close to participants' initial estimates.

H2b: Advice weighting will be higher the more pieces of advice participants sampled.

H2c: Advice weighting will be higher the lower participants' initial confidence.

H3a: Confidence will increase more when advice is close to as compared to distant from participants' initial estimates.

H3b: Confidence will increase more the more piece of advice participants sampled.

H3c: Confidence will increase more the lower participants' initial confidence.

5.2 Data analysis

As dependent measures, the self-determined sample size of advisory estimates, integration of advice when forming a final judgment and change in confidence after receiving advice were analyzed. The weight of advice (WOA) was calculated as a measure of advice integration in the same way as in chapter four.

Following the reasoning presented in chapter four, statistical criteria were applied to identify and remove outliers on the level of trials (Tukey, 1977) rather than altering values of WOA. Additional analyses applying the truncation practice yielded similar results, but led to a higher number of alterations than the exclusion of outliers.

As a measure of participants' confidence, participants were asked to complement their initial and final estimates by range estimates (confidence intervals; e.g., Klayman, Soll, González-Vallejo, & Barlas, 1999). The assessment of initial and final confidence via an interval estimate has several advantages. First, it allows for a more fine-grained analysis of participants' confidence as compared to the 7-point scale used in the previous chapter. Second, assessing confidence on the very same scale as the judgment it is associated with might provide a more direct access to participants' judgment-related confidence. Specifically, it has been argued that individuals outside the laboratory often express their judgment related confidence by providing such an interval (e.g., Soll & Klayman, 2004; Yaniv & Foster, 1995, 1997). Third, given that the here proposed theoretical account assumes confidence to reflect the reliability (i.e., amount and consistency) of individuals' information samples, it might be easier for individuals to transform this reliability assessment into a range estimate than into a single point on a rating scale.

To assess changes in confidence, participants' final confidence interval was divided by their initial confidence interval. Thus, a value smaller than 1 denotes a decrease in confidence interval width or an increase in confidence, whereas a value greater than 1 denotes an increase in confidence interval width or a decrease in confidence. This proportion was also highly sensitive to outliers. Analogous to the analysis of advice weighting, outlier trials were thus excluded (Tukey, 1977).

As in chapter four, multilevel models were computed to assess relationships on the level of trials (Judd et al., 2012). Analyses are based on 12 (20) data points per participant for Experiment 4 (Experiment 5). Following the recommendations by Judd and colleagues (2012), all models contained random intercepts for participants as well as items, which were fully crossed by design. The effects of the hypothesized predictors were always fixed. The best-fitting model was determined by sequentially including the hypothesized parameters and their interactions, one at a time and inspecting increase in model fit using the Kenward-Roger approximation (Kenward & Roger, 1997) for degrees of freedom. Once the best-fitting model was determined, those interactions were excluded that did not reach a t -value ≥ 2 (Baayen et al., 2008) as long as this did not decrease model fit significantly. The resulting p -values as well as log-likelihood values and approximated Bayes factors (Masson, 2011) for each model are given in the Appendix. For the sake of brevity, only the parameter estimates of the respective best-fitting model are reported.

To summarize, the expanded information asymmetry account predicts sampling frequency to increase with lower initial confidence (H1b). This assumption will be tested by looking for a main effect of initial CI on advice seeking. It is expected that sampling frequency is greater the broader a participant's initial CI is. The expanded information asymmetry account also predicts confidence change to be sensitive to the type and amount of information integrated. Specifically, confidence should increase more when close as compared to distant advice is integrated (H3a). Likewise, confidence should increase more,

the more pieces of advice are integrated (H3b). These assumptions will be tested by looking for main effects of advice distance and sampling frequency on confidence change.

Furthermore, the impact of additional information should be greater, the smaller or less consistent participants' initial information sample (H2c/H3c). This assumption will be tested by looking for a main effect of initial CI on both advice weighting and confidence change. It is expected that advice weighting and confidence change are greater the broader a participant's initial CI is.

5.3 A-priori power analyses

A-priori power analyses were conducted to determine required sample sizes for both experiments (Faul et al., 2007). As there is no clear guideline for power analyses regarding multilevel modelling utilized in the present research, calculations were based on repeated measures ANOVA designs. Detecting a medium-sized effect ($f = .25$) with sufficient power ($\beta = .95$) in Experiment 4 (Experiment 5) required collecting data of at least 28 (118) participants.⁶ As the experiments were part of multi-experiment sessions that lasted up to one hour, sample size was increased in accordance with other studies that required a larger sample size.

5.4 Experiment 4

The first experiment employed free sampling of close versus distant advice. The design thus replicates that of Experiment 2 in the previous chapter. To test the hypotheses regarding the role of confidence in advice seeking and taking, participants were asked to provide confidence intervals (CIs) for their initial and final point estimates.

⁶ Experiment 5 constitutes an integrative analysis of data from two separate experiments ($N = 60$ and $N = 68$). To avoid the interpretation of random patterns caused by too little power the data was combined. The experimental design was equal with one exception. In one experiment, the advice pointed away from the true value on half of the trials and toward the true value on the other half, whereas in the other experiment, the advice always pointed toward the true value. This quality of advice did not exert a systematic influence and was thus dropped from the analyses. The pattern of results was highly similar in both data sets. Aggregation of data should thus increase the reliability of the conclusions.

5.4.1 Method

Participants. Forty-two students of the University of Tübingen (15 male; $M_{age} = 23.38$, $SD_{age} = 4.59$)⁷ took part in a session of several psychological experiments that lasted about one hour and were compensated with either 7€ or course credit and a chocolate bar.

Design. The experiment implemented a 2 (distance of advice: close vs. distant) \times 2 (judgment phase: initial vs. final judgment) within-participants design.⁸

Materials and procedure. Materials and procedure were equal to those of Experiment 2, with two exceptions. First, participants were asked to estimate the airline distance in kilometers between 12 pairs of European cities (e.g., Cambridge and Belfast: 481km). True airline distance for each pair was retrieved from Google Maps (www.google.de/maps). Every city was part of only one pair. The order of city pairs was randomized, as was the order of the 6 trials per distance condition across the 12 trials.

Second, participants were asked to complement their initial and final estimates by range estimates. Thus, there were three response boxes on the screen, one to give their initial and final estimate and two for the lower and upper bounds. Participants were asked to use the lower and upper bound to create an 80% CI. Specifically, participants were instructed to assign a lower bound for which they were 90% certain that the true value would not be lower. Likewise, for the upper bound participants were instructed to give a value for which they were 90% certain that the true value would not be higher. The bonus was determined solely by their final estimate and not affected by the associated range.

⁷ The entries by two participants were regarded implausible (>100) and thus excluded from this descriptive statistic.

⁸ This experiment originally included a third distance condition. On these trials, participants received advice that was very close to their own initial estimate. Specifically, the advice was created by adding random uniform noise in the range of [0; 8] to the participant's initial estimate. This condition was introduced in an exploratory attempt to reconcile the contradicting effects of advice distance on advice weighting reported in previous research (compare Yaniv, 2004a with Moussaïd et al., 2013; Schultze et al., 2015). Specifically, whereas Yaniv (2004a) originally reported higher weighting of close as compared to distant advice, recent work reliably documented higher weighting of distant as compared to close advice (Moussaïd et al., 2013; Schultze et al., 2015). However, as the WOA is extremely sensitive to outliers, which mainly result from trials where the distance between initial estimate and advice is very low (see Data analysis), this condition mainly produced outliers of WOA and was thus dropped from all analyses.

5.4.2 Results

Distance of advice. As a manipulation check, the absolute difference between participants' initial estimates and the average advice was assessed for each trial. Whereas participants' initial estimate and the advice differed by only $M = 30.76$ ($SD = 55.02$) calories on close trials, they differed by $M = 269.63$ ($SD = 107.53$) calories on distant trials.

The ΔE to the average advice amounted to $M = .07$ ($SD = .10$) in the close condition and to $M = .77$ ($SD = .64$) in the distant condition. Hence, a “similar” and “intermediate” distance condition in terms of Moussaïd et al. (2013) were realized.

As a result, participants' initial estimate was enclosed by the range of advice for 61% of close trials but for none of the distant trials and participants' initial CI included the initial advice on 90% of close and on 29% of distant trials. Initial estimate and range of advice enclosed the true value on 18% of the close and 65% of the distant trials.

Sampling of advice. The expanded information asymmetry account predicts advice seeking to increase when participants' integrated information sample is small or inconsistent. Consequently, both distant advice and low initial confidence should increase advice seeking. The best-fitting model for advice seeking was the following:

$$Sample\ size_{ij} = a_{0ij} + b_1 Distance_j + b_2 Initial\ CI_{ij} + b_3 Distance_j * Initial\ CI_{ij} + \varepsilon.$$

Close advice was again coded as 0 and distant advice as 1, so that *Distance* indicates the effect of increasing advice distance. Likewise, participants' initial CIs were z-transformed for all models, so that *Initial CI* indicates the effect of a CI that is one standard deviation broader than the average CI, indicating lower confidence.⁹

Replicating the results from chapter four, participants sampled more additional advisory estimates when advice was distant rather than close, $b_1 = 1.63$, $se = .28$, $t = 5.76$

⁹ For both experiments, implausible initial CIs (> 15000) were excluded as they were assumed to result from typing errors. In Experiment 4, this concerned 1 out of 504 trials (< 0.2 %) and in Experiment 5 this concerned 3 out of 2560 trials (< 0.2 %). As model comparisons only work when both models were fit to the same data, these trials could also not be included in the models whenever initial CI served as a predictor.

(Figure 11). In line with the hypotheses, advice seeking was also greater for lower initial confidence, $b_2 = .92$, $se = .33$, $t = 2.80$. However, this effect was qualified by advice distance, $b_3 = -.76$, $se = .34$, $t = -2.22$, indicating that confidence only predicted advice seeking when advice was close.




Figure 11. Sample size of advisory estimates for close and distant trials of Experiment 4. Error bars represent standard errors of the means.

Weighting of advice. Trials with $WOA > 2.33$ and < -1.4 were excluded, amounting to 23 of 503 trials (4.57 %). Analogous to the previous experiments, it was assumed that distance and sample size increase advice weighting. Furthermore, lower initial confidence should also increase advice weighting. The best-fitting model for advice weighting was the following:

$$WOA_{ij} = a_{0ij} + b_1 Distance_j + b_2 Sample\ size_{ij} + \varepsilon.$$

Distant advice led to greater judgment revision, $b_1 = .31$, $se = .03$, $t = 9.30$ (Figure 12). Furthermore, additional sampling¹⁰ increased judgment revision, $b_2 = .011$, $se = .004$, $t = 2.75$.

¹⁰ To account for the fact that every participant received at least one piece of advice, 1 was subtracted from the sample size whenever it was used as a predictor in any model. Thus, *Sample size* indicates the effect of additional sampling of advisory estimates.

Interestingly, initial confidence did not exert any significant effect on judgment revision beyond its indirect influence via advice seeking, as apparent from the missing increase in model fit when including this predictor.




Figure 12. Weight of advice (WOA) for close and distant trials of Experiment 4. Error bars represent standard errors of the means.

The traditional WOA analysis was again complemented by investigating participants' behavior in comparison to the normative rule of equal weights averaging (see Experiment 2).

The model with the best fit was:

$$\Delta WOA_{ij} = \alpha_0 + b_1 Distance_j + \varepsilon.$$

As in the previous experiments, participants fell short of the normative rule on both close and distant trials (Figure 13), but came closer when advice was distant, $b_1 = .29$, $se = .03$, $t = 8.89$. However, sample size was not systematically related to the deviation from the normative rule.




Figure 13. Deviation of participants' weight of advice (WOA) from equal weights averaging for close and distant trials of Experiment 4. Error bars represent standard errors of the means.

Confidence change. Trials with confidence change > 1.64 were excluded, amounting to 29 of 503 trials (5.77 %). Based on the assumption that participants integrate the advice received into their information sample, it was expected that close advice increases confidence more than distant advice. Likewise, a higher frequency of sampling should lead to a greater increase in confidence. Furthermore, confidence should increase more, the lower participants' initial confidence. The best-fitting model for confidence change was the following:

$$\text{Confidence change}_{ij} = a_{0ij} + b_1 \text{Distance}_j + b_2 \text{Sample size}_{ij} + b_3 \text{Initial CI}_{ij} + b_4 \text{Distance}_j * \text{Sample size}_{ij} + \varepsilon.$$

Note that values smaller than 1 denote an increase in confidence, which is greater the smaller the value. On average, participants increased their confidence after receiving advice (Figure 14). Contrary to the assumption, confidence increase was not significantly greater for close as compared to distant advice, $b_1 = .05$, $se = .04$, $t = 1.32$. While additional sampling further increased confidence, $b_2 = -.009$, $se = .003$, $t = -2.55$, this was only true for close advice, $b_4 = .010$, $se = .004$, $t = 2.78$. The influence of advice distance on confidence was thus

mediated by advice seeking. Participants with a broader initial CI increased their confidence more strongly, $b_3 = -.096$, $se = .014$, $t = -6.84$.




Figure 14. Confidence change (Final CI/Initial CI) for close and distant trials of Experiment 4. Error bars represent standard errors of the means.

5.4.3 Discussion

Experiment 4 provides further support for the expanded information asymmetry account. Besides replicating an increase in advice seeking for distant as compared to close advice, the number of advisory estimates sampled also increases for lower initial confidence, but only when advice is close. Furthermore, confidence change is sensitive to the generated information samples, with greater increases in confidence after sampling additional close pieces of advice. Interestingly, low initial confidence did not increase advice weighting beyond its indirect effect via advice seeking but it led to greater changes in confidence.

5.5 Experiment 5

While Experiment 4 provided further support for the expanded information asymmetry account, some hypotheses were not confirmed. To increase reliability of the above findings, another experiment was run. Additionally, this experiment investigates a highly interesting

context factor that is relevant to advice seeking and taking outside the laboratory, namely costs of advice. The absolute amount of advice seeking is likely determined by a trade-off between the costs of obtaining advice and the benefits associated with this advice (Bonaccio & Dalal, 2006). Indeed, participants were found to be less likely to solicit advice at all, when this was costly (Gino, 2008) and costs reduced sampling of information in related paradigms (Lee et al., 2014). While it is thus expected that costs generally decrease advice seeking, the relative increase in advice seeking after integrating distant advice and when initial confidence is low should generalize even to a setting where seeking advice is costly. At the same time, costly advice was found to lead to greater advice weighting (Gino, 2008; Patt, Bowles, & Cash, 2006; Snizek, Schrah, & Dalal, 2004). If confidence is a measure of advice taking analogous to advice weighting, costs could also increase the impact of advice on confidence.

5.5.1 Method

Participants. 128 students of the University of Tübingen (26 male; $M_{age} = 22.48$, $SD_{age} = 3.76$) took part in a session of several unrelated experiments that lasted about one hour in total. They were compensated with either 7€ or course credit and a chocolate bar. Participants were randomly assigned to either the free ($N = 63$) or the costly sampling condition ($N = 65$).

Design. The experiment implemented a 2 (distance of advice: close vs. distant) \times 2 (judgment phase: initial vs. final judgment) \times 2 (sampling: free vs. costly) mixed design. The first two factors were implemented within participants, whereas the third factor was implemented between participants.

Materials and procedure. The experiment was a replication of Experiment 4 with two exceptions. First, the number of trials per participant was increased to 20, thus every participant received close advice on 10 trials and distant advice on 10 other trials. Order of trials and city pairs remained random. Second, advice sampling was costly for half of the participants. Specifically, in the *free sampling* condition, additional advice appeared immediately after clicking the respective button. In contrast, in the *costly sampling* condition,

every additional piece of advice appeared 6 seconds after clicking the respective button. During this waiting time, all buttons were disabled. The instructions did not inform participants about this waiting time. Thus, participants only experienced during the experiment that advice was costly to obtain.

5.5.2 Results

Distance of advice. As a manipulation check, the absolute difference between participants' initial estimates and the average advice was assessed for each trial. Whereas participants' initial estimate and the advice differed by only $M = 57.62$ ($SD = 396.50$) calories on close trials, they differed by $M = 282.47$ ($SD = 393.19$)¹¹ calories on distant trials.

The ΔE to the average advice amounted to $M = .10$ ($SD = .18$) in the close condition and to $M = .75$ ($SD = 1.07$) in the distant condition. Hence, a “similar” and “intermediate” distance condition in terms of Moussaïd et al. (2013) were realized.

As a result, participants' initial estimate was enclosed by the range of advice for 43% of close trials but for none of the distant trials and participants' initial CI included the initial advice on 88% of close and on 26% of distant trials. Initial estimate and range of advice enclosed the true value on 11% of the close and 48% of the distant trials.

Sampling of advice. The best-fitting model for advice seeking was the following:

$$Sample\ size_{ij} = a_{0ij} + b_1 Distance_j + b_2 Costs_i + b_3 Initial\ CI_{ij} + b_4 Distance_j * Costs_i + \varepsilon.$$

For all models including costs, the free condition was coded as 0 and the costly condition as 1, thus *Costs* indicate the effect of changing sampling from free to costly.

Replicating the findings from Experiment 4, distant advice increased advice seeking, $b_1 = 1.31$, $se = .14$, $t = 9.31$. Interestingly, lower initial confidence also increased advice seeking, $b_3 = .20$, $se = .06$, $t = 3.48$, and this effect was not qualified by advice distance. As expected, costs decreased advice seeking, $b_2 = -6.23$, $se = .65$, $t = -9.51$ (Figure 15). Costs

¹¹ Because participants' initial estimates were not constrained to a plausible range in these experiments but the computed advice was, advice was extremely distant on some trials. For the descriptive statistics we thus excluded the trials on which the absolute distance was greater than 10000 (3 out of 2560 trials, < 0.2 %).

also attenuated the effect of advice distance on advice seeking, $b_4 = -.94$, $se = .20$, $t = -4.76$.

To test whether advice distance increased advice seeking even when advice was costly to obtain, the same model was fit separately to data from each cost condition.

When advice was free to obtain, distant advice, $b_1 = 1.31$, $se = .19$, $t = 6.88$ (Figure 15), and lower initial confidence, $b_3 = .24$, $se = .10$, $t = 2.30$, increased advice seeking.

Likewise, when advice was costly to obtain, both distant advice, $b_1 = .36$, $se = .06$, $t = 6.06$ (Figure 15), and lower initial confidence, $b_3 = .15$, $se = .04$, $t = 3.73$, increased advice seeking.

Interestingly, while the effect of advice distance on advice seeking seems to be attenuated by advice costs, the effect of initial confidence appears more pronounced. Z-transforming sample sizes within each cost condition to account for the different base lines due to costs, the interaction between costs and initial CI is positive, $b = .04$, $se = .03$, $t = 1.54$, but the model including this interaction explains the data only slightly better than the model without it, $\chi^2(1) = 2.38$, $p = .12$.




Figure 15. Sample size of advisory estimates for close and distant trials in the free and costly sampling condition of Experiment 5. Error bars represent standard errors of the means.

Weighting of advice. Trials with WOA > 2.17 and < -1.3 were excluded, amounting to 130 of 2553 trials (5.09 %). The best-fitting model for advice weighting was the following:

$$WOA_{ij} = a_{0ij} + b_1 Distance_j + b_2 Sample\ size_{ij} + b_3 Costs_i + b_4 Initial\ CI_{ij} + b_5 Sample\ size_{ij} * Costs_i + b_6 Sample\ size_{ij} * Initial\ CI_{ij} + b_7 Costs_i * Initial\ CI_{ij} + b_8 Sample\ size_{ij} * Costs_i * Initial\ CI_{ij} + \epsilon.$$

Replicating the findings from Experiment 4, both advice distance, $b_1 = .27$, $se = .02$, $t = 16.90$, and additional advice sampling, $b_2 = .015$, $se = .003$, $t = 5.08$, increased advice weighting, whereas initial CI did not, $b_4 = -.01$, $se = .02$, $t = -0.80$. Costs did not influence advice weighting per se, $b_3 = -.02$, $se = .05$, $t = -0.45$, but interacted with both, sample size, $b_5 = .042$, $se = .009$, $t = 4.38$, and initial CI, $b_7 = -.07$, $se = .04$, $t = -1.84$, whereas the latter two did not interact, $b_6 = .001$, $se = .002$, $t = 0.40$. Because the three-way interaction between sample size, costs and initial CI was also significant, $b_8 = .03$, $se = .01$, $t = 2.70$, the model was fit separately to data from each cost condition.

When advice was free to obtain, both distant advice, $b_1 = .31$, $se = .02$, $t = 13.08$ (Figure 16), and additional sampling, $b_2 = .014$, $se = .003$, $t = 4.70$, increased advice weighting. Initial CI was not related to advice weighting, $b_4 = -.01$, $se = .02$, $t = -0.77$, and did not interact with sample size, $b_6 = .001$, $se = .002$, $t = 0.36$.

When advice was costly to obtain, both distant advice, $b_1 = .23$, $se = .02$, $t = 10.81$ (Figure 16), and additional sampling, $b_2 = .059$, $se = .009$, $t = 6.52$, increased advice weighting. Interestingly, if participants did not sample additional advice, lower initial confidence decreased advice weighting, $b_4 = -.08$, $se = .03$, $t = -2.45$, but the lower the initial confidence, the stronger additional sampling increased advice weighting, $b_6 = .03$, $se = .01$, $t = 2.78$.




Figure 16. Weight of advice (WOA) for close and distant trials in the free and costly sampling condition of Experiment 5. Error bars represent standard errors of the means.

The traditional WOA analysis was again complemented by investigating participants' behavior in comparison to the normative rule of equal weights averaging (see Experiment 2). The model with the best fit was:

$$\Delta WOA_{ij} = a_{0ij} + b_1 Distance_j + b_2 Sample\ size_{ij} + b_3 Costs_i + b_4 Initial\ CI_{ij} + b_5 Sample\ size_{ij} * Costs_i + b_6 Sample\ size_{ij} * Initial\ CI_{ij} + b_7 Costs_i * Initial\ CI_{ij} + b_8 Sample\ size_{ij} * Costs_i * Initial\ CI_{ij} + \epsilon.$$

Replicating the previous experiments, participants weighted advice more in accord with the normative rule when it was distant, $b_1 = .26$, $se = .02$, $t = 16.27$, additional advice sampling, $b_2 = -.001$, $se = .003$, $t = -0.24$ and participants' initial CI, $b_4 = -.01$, $se = .02$, $t = -0.61$, however, did not systematically affect the deviation from the normative rule. Costs decreased the deviation from the normative rule, $b_3 = .11$, $se = .05$, $t = 2.19$, and interacted with initial CI, $b_7 = -.075$, $se = .037$, $t = -2.02$, but not with sample size, $b_5 = .004$, $se = .010$, $t = 0.37$. Sample size and initial CI did not interact, $b_6 = .000$, $se = .002$, $t = 0.11$. Because the three-way interaction between sample size, costs and initial CI was significant, $b_8 = .03$, $se = .01$, $t = 2.63$, the model was fit separately to data from each cost condition.

When advice was free to obtain, participants' advice weighting more closely matched the normative rule when advice was distant, $b_1 = .31$, $se = .02$, $t = 12.86$ (Figure 17), but additional sampling did not have any influence, $b_2 = -.002$, $se = .003$, $t = -0.58$. Initial CI was not related to the deviation, $b_4 = -.01$, $se = .02$, $t = -0.59$, and did not interact with sample size, $b_6 = .00$, $se = .002$, $t = 0.11$.

When advice was costly to obtain, participants' advice weighting more closely matched the normative rule when advice was distant, $b_1 = .22$, $se = .03$, $t = 10.12$ (Figure 17). Additional sampling was again unrelated, $b_2 = .006$, $se = .009$, $t = 0.63$. Interestingly, if participants did not sample additional advice, lower initial confidence increased the deviation from the normative rule, $b_4 = -.08$, $se = .03$, $t = -2.48$, but the lower the initial confidence, the stronger additional sampling decreased this deviation, $b_6 = .03$, $se = .01$, $t = 2.65$.




Figure 17. Deviation of participants' weight of advice (WOA) from equal weights averaging for close and distant trials in the free and costly sampling condition of Experiment 5. Error bars represent standard errors of the means.

Confidence change. Trials with confidence change > 1.71 were excluded, amounting to 126 of 2553 trials (4.94 %). The best-fitting model for confidence change was the following:

$$\begin{aligned}
\text{Confidence change}_{ij} = & a_{0ij} + b_1\text{Distance}_j + b_2\text{Sample size}_{ij} + b_3\text{Costs}_i + b_4\text{Initial CI}_{ij} + \\
& b_5\text{Distance}_j * \text{Sample size}_{ij} + b_6\text{Distance}_j * \text{Costs}_i + b_7\text{Distance}_j * \text{Initial CI}_{ij} + b_8\text{Sample} \\
& \text{size}_{ij} * \text{Costs}_i + b_9\text{Sample size}_{ij} * \text{Initial CI}_{ij} + b_{10} \text{Costs}_i * \text{Initial CI}_{ij} + b_{11}\text{Distance}_j * \text{Sample} \\
& \text{size}_{ij} * \text{Costs}_i + b_{12}\text{Distance}_j * \text{Sample size}_{ij} * \text{Initial CI}_{ij} + b_{13}\text{Distance}_j * \text{Costs}_i * \text{Initial CI}_{ij} + \\
& b_{14}\text{Sample size}_{ij} * \text{Costs}_i * \text{Initial CI}_{ij} + b_{15}\text{Distance}_j * \text{Sample size}_{ij} * \text{Costs}_i * \text{Initial CI}_{ij} + \varepsilon.
\end{aligned}$$

On average participants increased their confidence after receiving advice (Figure 18). Replicating the findings from Experiment 4, increase in confidence was not greater for close as compared to distant advice, $b_1 = .05$, $se = .03$, $t = 1.82$. However, additional sampling increased confidence more when advice was close, $b_2 = -.008$, $se = .003$, $t = -3.06$, as indicated by the interaction between advice distance and sample size, $b_5 = .005$, $se = .003$, $t = 1.92$. Costs did not influence confidence increases per se, $b_3 = .02$, $se = .04$, $t = 0.39$, but interacted with sample size, $b_8 = -.030$, $se = .009$, $t = -3.41$. While advice distance and costs did not interact, $b_6 = .004$, $se = .036$, $t = 0.10$, the three-way interaction between distance, sample size and costs was marginal, $b_{11} = .02$, $se = .01$, $t = 1.88$. Again, a broader initial CI led to a higher increase in confidence, $b_4 = -.04$, $se = .02$, $t = -2.62$ and increased the impact of additional sampling, $b_9 = -.005$, $se = .002$, $t = -2.54$. Initial CI did not interact with distance, $b_7 = .02$, $se = .02$, $t = 1.00$, and costs, $b_{10} = -.02$, $se = .03$, $t = -0.80$. The three-way interactions between distance, sample size and initial CI, $b_{12} = -.002$, $se = .003$, $t = -0.85$, between distance, costs and initial CI, $b_{13} = -.008$, $se = .046$, $t = -0.18$, and between sample size, costs and initial CI did not reach significance, $b_{14} = .01$, $se = .01$, $t = 1.40$. However, the four-way interaction between distance, sample size, costs and initial CI was significant, $b_{15} = -.06$, $se = .02$, $t = -3.04$. Hence, the model was fit separately to data from each cost condition.

When advice was free to obtain, advice distance marginally reduced the increase in confidence, $b_1 = .05$, $se = .03$, $t = 1.79$ (Figure 18). While sampling of close advice increased confidence, $b_2 = -.008$, $se = .003$, $t = -3.17$, sampling on distant trials did so to a lesser degree, $b_5 = .005$, $se = .003$, $t = 1.86$. Again, lower initial confidence led to a greater increase in

confidence, $b_4 = -.04$, $se = .02$, $t = -2.51$, and increased the impact of additional sampling, $b_9 = -.005$, $se = .002$, $t = -2.44$. The interaction between distance and initial CI, $b_7 = .02$, $se = .02$, $t = 0.95$, and the three-way interaction between distance, sample size and initial CI were not significant, $b_{12} = -.002$, $se = .003$, $t = -0.79$.

When advice was costly to obtain, advice distance also reduced the increase in confidence, $b_1 = .05$, $se = .02$, $t = 2.50$ (Figure 18). While sampling of close advice increased confidence, $b_2 = -.037$, $se = .008$, $t = -4.56$, sampling on distant trials again did so to a lesser degree, $b_5 = .023$, $se = .009$, $t = 2.63$. Lower initial confidence again led to a greater increase in confidence, $b_4 = -.06$, $se = .03$, $t = -2.49$. While the interactions between sampling and initial CI, $b_9 = .008$, $se = .009$, $t = 0.89$, and between distance and initial CI, $b_7 = .01$, $se = .04$, $t = 0.36$, were not significant, the three-way interaction between distance, sample size and initial CI was significant, $b_{12} = -.06$, $se = .02$, $t = -3.35$. Hence, the sampling of additional distant advice increased confidence only when participants' confidence was low at the beginning of a trial.




Figure 18. Confidence change (Final CI/Initial CI) for close and distant trials in the free and costly sampling condition of Experiment 5. Error bars represent standard errors of the means.

5.5.3 Discussion

Experiment 5 replicated all findings from Experiment 4. Crucially, lower confidence was again related to more sampling of additional advice and more sampling of close advice again led to a greater increase in confidence. While costs decreased advice seeking, the reported increase in impact of costly advice on judgment revision was replicated (e.g., Gino, 2008). Additionally, this increased impact of costly advice generalizes to increase in confidence as a measure of advice taking. Interestingly, the missing direct influence of initial confidence on advice weighting was also replicated in the free sampling condition. In contrast, when advice was costly to obtain, initial confidence influenced advice weighting beyond advice seeking by increasing the impact of additionally sampled advice. As in Experiment 4, lower initial confidence led to a stronger increase in confidence. In contrast to Experiment 4, lower initial confidence also increased the positive effect of additional sampling on confidence.

5.6 Interim discussion

The current chapter investigates participants' confidence before and after taking advice to provide converging evidence for the expanded information asymmetry account. Research on subjective confidence assumes confidence to result from participants' evaluations of the reliability of their information samples which in turn is determined by the amount and consistency of information (Koriat, 2012a). Consequently, measuring confidence change as a result of the advice received allows to test the assumption that this advice is integrated into participants' information samples. In line with this assumption, confidence increased more when participants sampled more advice and this effect was more pronounced when advice was close to participants' initial estimates.

The two experiments also shed more light on the determinants of information search in advice taking. In line with the expanded information asymmetry account, participants sampled more advice when their initial confidence was low and when advice was distant. Both relate

to a lower reliability of the integrated information sample. While low initial confidence should reflect a small or inconsistent information sample (Koriat, 2012a), integration of an initial piece of distant advice should decrease the consistency of the information sample, regardless of its original size or consistency. In line with the assumption that overall sampling frequency is determined by a cost-benefit trade-off (Bonaccio & Dalal, 2006), costs decreased sampling frequency.

Costs increased the impact of the sampled advice on both advice weighting and confidence change. Previous research suggested a sunk-cost fallacy to explain the increased impact of costly advice on judgment revision (Gino, 2008). However, in Experiment 5 costs not only increased the impact of additional information received, but also the impact of initial confidence on advice weighting. Costs also slightly increased the impact of initial confidence on advice seeking. It thus appears that costs did not only increase the impact of external information, but increased participants' sensitivity to the information samples in general – within their minds and in the information ecology. Similar assumptions have been made regarding the influence of benign or malign environments on top-down or bottom-up processing (Fiedler & Hütter, 2014).

5.7 Limitations

The expanded information asymmetry account allows to derive a number of hypotheses regarding the interaction of individuals' initial information samples with the additional information sampled from the ecology. While the majority of hypotheses were supported by experimental data, some received only mixed support or would have to be rejected given the current data.

The information asymmetry account assumes that the reliability of individuals' information samples is reflected by their subjective confidence (Koriat, 2012a). It is assumed that lower reliability leads to more sampling of additional information. While this hypothesis was supported in Experiment 5, sampling in Experiment 4 was only predicted by participants'

initial confidence when advice was close. Potentially, the reliability of individuals' information samples only predicts additional sampling when this sampling increases the reliability. Opposing this assumption, distant advice in the same experiment also led to an increase in confidence, which according to the proposed account indicates an increase in the reliability of participants' information samples.

On a related note, it is assumed that additional information that is integrated into the current information sample has a greater impact the smaller or less consistent this information sample is (Weber-Fechner Law). Consequently, lower initial confidence should increase the impact of additional advice on both judgment revision and change in confidence. While the experiments document and increased change in confidence when initial confidence was lower, the relation between initial confidence and advice weighting was only significant in the costly sampling condition of Experiment 5.

Furthermore, the same amount of additional information did not consistently affect advice utilization to the same degree. Specifically, a closer look at data from Experiment 5 shows that costs greatly reduced participants' sample size (Figure 15), but did not decrease the weight of advice (Figure 16) or the increase in confidence (Figure 18). The results remain unaltered when the absolute width of participants' final CI is modelled instead of confidence change. This increased impact of costly advice cannot result from a simple integration of the advice into the current information sample in the absence of any other influence factors. One would at least have to assume an increased impact of any kind of information in a malign environment (Fiedler & Hütter, 2014). However, this effect shows that an account that assumes advice utilization to result from pure information integration cannot accommodate all determinants.

While some of the above limitations might just reflect randomness in experimental data (e.g., the moderation of the relation between initial confidence and advice seeking by advice distance in Experiment 4), others highlight the necessity to include additional

processes into the so far purely informational account (e.g., the influence of costs). The uncertain influence of individuals' initial information samples as indicated by confidence might reflect a methodological problem. Specifically, subjective confidence only approximates the reliability of individuals' information samples (Koriat, 2012a). Subjective confidence, however, can be influenced by many more determinants. In fact, in some studies the increased weighting of advice resulted from decreased confidence that was not task related (e.g., Gino et al., 2012) and thus should not have reflected individuals' information samples. To truly measure the influence of individuals' internal information samples, future research could manipulate these information samples to gain experimental control (see General discussion).

6. General discussion

The current thesis aims to advance research on advice taking by identifying and addressing three core problems: the uncertain adaptivity of advice taking behavior; the lack of a unifying theoretical framework; and the conceptual constraint of the existing research. To this end, an expanded paradigm is introduced which allows to incorporate information search in advice taking. A theoretical framework is proposed, that explains advice taking, information search and their relation and integrates the most important identified determinants of advice taking. Finally, the adaptivity of advice taking behavior is re-evaluated from this ecological perspective. In the following, each of these contributions shall be put to careful scrutiny. Specifically, it will be assessed how far these contributions advance existing research on advice taking, which new questions arise from them and how these questions can be addressed by future research.

6.1 Information search in advice taking

The current thesis is the first to investigate information search in advice taking. By investigating the sequential process of seeking and integrating advice the current research answers a repeated call from other scientists (Bonaccio & Dalal, 2006; Gino et al., 2012). While previous research only investigated whether advice was solicited at all (Gino et al., 2012; Gino & Moore, 2007) the current research allows individuals to seek out several additional advisors after receiving an initial piece of advice.

This thesis advances advice taking research for three reasons: First, it constitutes the first investigation of advice taking in a situation that allows individuals to interact with their information ecology. Second, it allows to assess whether advice seeking and taking is sensitive to features of the information ecology, such as the distance or consistency of advice. Third, it provides a conceptual link between research on advice taking and related research that investigates information search in judgment and decision making.

One central claim of proponents of advice taking research is that many decision situations are interactive (Sniezek & Buckley, 1995). Consequently, advice taking constitutes a way of improving judgment accuracy that builds on existing ecological structures. In stark contrast to this claim, no research on advice taking has yet satisfied this assumption of an interactive situation (Bonaccio & Dalal, 2006; Gino et al., 2012). This thesis is thus the first to present experiments that investigate advice taking in a situation that can truly be described as interactive.

The present experiments demonstrate that individuals are generally willing to seek additional information (here: additional advice) in the advice taking paradigm. Moreover, this information search is sensitive to features of the information ecology. Individuals sampled more additional advice when this advice diverged more strongly from their own initial judgment, when they were less confident and when advice was overall less consistent. Furthermore, advice utilization was highly sensitive to the information thus sampled. Individuals increased their reliance on advice, for both judgment revision and confidence change, when they sampled more pieces of advice. This effect of the amount of information was moderated by its consistency. Specifically, additional sampling increased reliance on advice more, when advice was more consistent among advisors.

The present research complements findings from a related research area that investigates sampling behavior prior to the choosing of options. In research on optional stopping, for instance, participants sample (costly) offers until they stop sampling and keep the latest offer (e.g., Bearden & Rapoport, 2005; Rapoport & Tversky, 1966; Seale & Rapoport, 1997). In research investigating decisions from experience, participants are usually presented with two risky lotteries and explore their outcomes before deciding which one of these lotteries to play for actual profit (e.g., Hertwig, Barron, Weber, & Erev, 2004; Hertwig & Erev, 2009). Although there is some debate about the optimal sample size (e.g., Hertwig & Pleskac, 2010), this research generally demonstrates that participants sample substantially,

investing money and/or time. Moreover, as in the present research the amount of sampling is adaptive, being sensitive to the variability of the sample (e.g., Lejarraga, Hertwig, & Gonzalez, 2012; Mitra, Reiss, & Capella, 1999; Pachur & Scheibehenne, 2012), the importance of the decision (Hau, Pleskac, Kiefer, & Hertwig, 2008), and competition for the choice options (Phillips et al., 2014).

However, advice taking research differs greatly from the above reviewed literatures in the nature of the information presented. Whereas research on decisions from experience and optional stopping presents factual information (i.e., individuals' own observations of outcomes), research on advice taking presents social information (i.e., other individuals' judgments). This difference has two important implications. First, it gives rise to different underlying motives. Both paradigms usually reward accuracy. However, the social nature of advice adds to this accuracy motive a motive to be affirmed by others (Kahneman & Tversky, 1973; Schulz-Hardt, Frey, Lüthgens, & Moscovici, 2000) and a motive to reciprocate or cooperate (Goldsmith & Fitch, 1997; Harvey & Fischer, 1997; Homans, 1958). Second, due to the social nature of the information source, the assessment of reliability of advice should be influenced greatly by various kinds of social information (e.g., attractiveness; Chaiken, 1979; e.g., group membership; Tajfel, 1982). Interestingly, initial experiments investigating the integration of information that was framed as either social or non-social suggest that individuals are highly sensitive to this framing (Collins, Percy, Smith, & Kruschke, 2011). Specifically, whereas redundant information within each framing (i.e., social or non-social) was disregarded, redundant information across framings (i.e., social and non-social) was weighted more strongly.

Observing others' actions constitutes one of the most central processes of human learning (Bandura, 1977) and might even have been the cause for the emergence of human intelligence (Herrmann et al., 2007; Moll & Tomasello, 2007). Consequently, understanding how individuals seek and utilize information provided by others constitutes a highly important

research goal. Expanding research on advice taking by the opportunity to seek additional information, the current thesis provides a necessary contribution to this endeavor.

6.1.1 Limitations and future directions

The current thesis introduces information search into the advice taking paradigm thus providing individuals with the opportunity to interact with their information ecology. While this constitutes an important advancement of the ecological validity of advice taking research, there are a number of features of the ecology in which advice taking most likely takes place that have not been investigated in the here documented experiments.

The current experiments did not provide individuals with any information regarding their advisors. However, outside of such deprived laboratory settings it appears highly unlikely that individuals receive advice without any information regarding its source. While a number of such characteristics of the advisor have been shown to influence advice weighting (e.g., conflict of interests; Gino et al., 2012; expertise; Harvey & Fischer, 1997), no research has yet investigated how these characteristics relate to the seeking of additional advice. Furthermore, some important characteristics of the advisor or the advisee have not yet been investigated by advice taking research. Arguably the most important characteristic of any individual is its group membership (at least from a social psychologist's perspective). Indeed, individuals are highly sensitive to others' group membership and often favor in-group members (for reviews see Hewstone, Rubin, & Willis, 2002; Tajfel, 1982). For example, in-group members are trusted more (Insko, Schopler, Hoyle, Dardis, & Graetz, 1990) and are protected more against norm violations (Bernhard, Fischbacher, & Fehr, 2006). Consequently, one could expect advice from in-group members to be weighted more strongly than advice from out-group members. At the same time however, groups are often highly homogeneous (McPherson, Smith-Lovin, & Cook, 2001) and tend to think along common lines (Janis, 1972), potentially limiting the degree to which judgment biases from individuals within a group are uncorrelated. Because this uncorrelatedness constitutes an important precondition to

maximal gains from aggregation of judgments (Einhorn et al., 1977; Larrick & Soll, 2006), utilization of advice from out-group members could yield greater accuracy. This conflict between in-group favoritism and normative considerations in the area of advice renders group membership an intriguing determinant of advice taking. Likewise, it might be interesting whether advice taking is sensitive to social exclusion (ostracism; Baumeister & Leary, 1995; Over & Carpenter, 2009; Williams, Cheung, & Choi, 2000). Individuals experiencing ostracism have been found to conform more to group norms (Williams et al., 2000) and imitate others more strongly (Over & Carpenter, 2009). Potentially, experiencing ostracism also increases individuals' motives to cooperate or reciprocate, which in turn might increase their willingness to accept advice from others.

A related limitation to the current research concerns the choice of the information source. Specifically, the present experiments only investigated the amount of information sought. Outside the laboratory, individuals most likely have the opportunity to choose among several information sources. The here documented results suggest that individuals are more likely to seek additional advice and revise their judgment in light of diverging advice. However, given the opportunity to decide whom to ask for advice, individuals might appear less adaptive. For example, individuals are more likely to turn to close and similar others for social comparisons (Suls, Martin, & Wheeler, 2002) and even assume these others to be representative of the actual population distribution (Galesic, Olsson, & Rieskamp, 2012). Likewise, individuals are more likely to bet on judgments that have been revised in light of advice that was sampled dependently rather than independently from their own judgment (Yaniv et al., 2009). Whereas such a choice of advisor based on similarity might be beneficial for matters of taste (Yaniv, Choshen-Hillel, & Milyavsky, 2011), accuracy should benefit most from consulting others with maximally different opinions (Larrick & Soll, 2006). Consequently, future research should complement this thesis' focus on the amount of information sought by an investigation of the source of information consulted.

While the above two paragraphs illustrate that this thesis only constitutes a starting point towards the investigation of advice taking from an ecological perspective, the sampling paradigm proposed can easily be expanded by additional factors, be it information about the source of information or the opportunity to consult different such sources.

6.2 Expanded information asymmetry account revisited

The current thesis proposes an information integration account of advice taking. The expanded information asymmetry account is the first account that conceptualizes advice taking as an interaction between individuals' minds and the information ecology. Furthermore, it provides a theoretical link to related literatures on subjective confidence (Koriat, 2012a) and sampling approaches to judgment and decision making (Denrell, 2005; Fiedler, 2000; Hertwig & Erev, 2009; Newell, 2005).

There is no theory of advice taking, let alone advice seeking (Bonaccio & Dalal, 2006). Previous theoretical approaches were mainly designed to explain isolated phenomena of advice taking research. For example, the original information asymmetry account (Yaniv, 2004a, 2004b) was introduced to explain underweighting of advice and the (then assumed) decrease in advice weighting when advice is distant from the own judgment. It was assumed that individuals' information samples are skewed towards supporting the own judgment. While this approach is suited to explain the general underweighting of advice (e.g., Yaniv & Kleinberger, 2000), it does not allow for the expansion of these information samples, for example by exploring the information ecology. Recently, advice taking has been compared to a basic stimulus response model (Schultze et al., 2015). This approach was introduced to explain the change in advice weighting depending on the distance of advice from the own judgment. It was assumed that integration of advice depends on the strength of the stimulus (i.e., advice), which is determined for example by its distance from the own judgment. The strength of the stimulus might also be greater when several pieces of advice converge. This account might thus also explain how advice weighting changes as a function of the amount

and consistency of advice received. However, this account completely neglects the information samples in individuals' minds. For example, it does not account for the influence of individuals' knowledge (e.g., Harvey & Fischer, 1997) or confidence (e.g., Gino et al., 2012) on advice taking. It has long been emphasized that behavior is best understood as an interaction between the mind and the information ecology (Brunswik, 1955; Simon, 1956). While such interactions are still overlooked too often (Fiedler, 2000; Le Mens & Denrell, 2011), they have recently greatly advanced research in related paradigms (e.g., Denrell, 2005). The expanded information asymmetry account constitutes the first theoretical approach to advice taking that can account for such interactions.

Additionally, the expanded information asymmetry account resonates well with theoretical models proposed in related literatures. Recent research on subjective confidence assumes that individuals' confidence reflects the reliability of their information samples, which in turn is determined by the amount and consistency of the information (Koriat, 2012a). The current thesis expands this account by allowing for the integration of additional information into these information samples. Likewise, a variety of research employs sampling approaches to judgment and decision making to account for interactions between individuals' minds and the information ecology (e.g., Denrell, 2005; Fiedler, 2000; Hertwig & Erev, 2009; Newell, 2005). Just as the theoretical account proposed here, this research often assumes a continuous integration of additional information into an individuals' current information sample (e.g., Denrell, 2005; Newell, 2005).

6.2.1 Limitations and future directions

The current thesis provides evidence that the expanded information asymmetry account is suited to explain the precursors and consequences of information search in advice taking and can integrate existing findings. However, there are several ways in which the account can be expanded and hypotheses could be tested more rigorously. The most intriguing ones shall be outlined below.

The current work identifies several determinants of individuals' sampling frequency in the advice taking paradigm. In line with the assumption that individuals are sensitive to the reliability of their information samples, factors that likely decreased this reliability (e.g., distant or inconsistent advice) or indicated a low reliability (low confidence) were related to greater sampling of additional advice. However, no predictions were tested with regard to the truncation of sampling behavior. Following the above reasoning, it could be assumed that individuals truncate sampling once the reliability of their integrated information sample reaches a certain threshold (relative or absolute). Indeed, confidence has been successfully modelled as a stopping criterion for information search in research on decisions from experience (Lee et al., 2014). A comparison of participants' confidence after receiving close and distant advice appears to oppose this assumption. Specifically, both participants' absolute confidence (Experiment 2 and 3) and increase in confidence (Experiment 4 and 5) were lower on distant as compared to close trials. However, the experiments were not designed to investigate confidence as a truncation criterion for advice sampling. Sampling was capped at a maximum of 20 pieces of advice, so that individuals were not entirely free in determining their sample sizes. Confidence was only investigated before and after advice sampling. There is thus no data revealing how confidence changed throughout the sampling process. Potentially, individuals truncated sampling as soon as additional information did not sufficiently change the reliability of their information sample any more. Future research should thus investigate unlimited sampling of advice and track confidence changes throughout the sampling process (e.g., using graphical displays of intervals).

Such an investigation of confidence throughout the sampling process could also serve to shed light on the relation between integration of distant advice and confidence. Whereas an initial piece of distant advice most likely decreases the consistency of the integrated information sample, distant advice generally led to an increase in confidence in the above experiments (although this increase was smaller than for close advice). It was argued that the

relation between integration of distant advice and confidence could be u-shaped. Tracking confidence throughout the sampling process allows to test this assumption.

A related consideration concerns the here proposed differences between close and distant advice. The distance levels investigated in the above experiments were based on findings by other authors (Moussaïd et al., 2013; Schultze et al., 2015) and replicated the reported pattern in advice utilization. These authors account for this qualitative difference by assuming an area of confirmation around an individual's judgment. Advice that falls inside this area of confirmation increases confidence but does not lead to judgment revision. Advice that falls outside this area leads to judgment revision. However, these authors propose models which define the area of confirmation based on empirical data aggregated across participants. Tracing this pattern back to increases or decreases of the consistency of individuals' information samples affords more precise predictions. Specifically, whether advice primarily increases confidence or leads to judgment revision should vary as a function of an individual's specific information sample with regard to a specific judgment task and the specific piece of advice received. While the information samples in individuals' minds are difficult to measure, range estimates might serve as a proxy. One could argue that advice that falls within an individuals' range estimate likely increases consistency while advice that falls outside the range decreases consistency. Indeed, close advice as operationalized here was highly correlated with the initial advice falling inside individuals' initial range estimates ($r = .62$, for Experiments 4 and 5). Future research could test this assumption more thoroughly by manipulating advice distance continuously and pitting absolute distance against advice falling in- versus outside of the range estimate as predictors for confidence increase and judgment revision.

A different approach to individuals' initial information samples can afford even more precise predictions. Whereas initial confidence only approximates individuals' information samples (Koriat, 2012a), these information samples could be directly manipulated

experimentally. For example, one could design a learning paradigm in which individuals face tasks that they have no prior knowledge about (e.g., Harvey & Fischer, 1997). One could then present individuals with some information regarding the specific item before presenting advice. This approach allows to thoroughly test the here proposed assumption that initial information and advice received impact advice taking depending on their amount and consistency. For example, it would be expected that the less information individuals were presented on a given item, the greater the impact of additionally sampled advice. Initial experiments support this assumption. The less training individuals received to predict cue-criterion relations and the less strong this relation was, the higher they weighted the advice received (Harvey & Fischer, 1997). Future research could investigate the integration of additionally sampled advice with varying consistency after different levels of training.

Such an approach could also deliver the data necessary to transform the here proposed conceptual model into a computational model that affords even more precise predictions about advice seeking and taking. Specifically, some processes need to be investigated more closely for the successful development of such a computational model. For example, it is unclear how individuals produce a point estimate from a distribution of information. Related research suggests that individuals often focus on the mean of distributions, even when this is not warranted (Lindskog, Winman, & Juslin, 2013). However, the current research focuses on the amount and consistency of information, which might also be assessed in form of the relative density (Figure 5). Manipulating the information individuals have and then measuring the point estimate they make of it could provide an initial opportunity to disentangle these two (and many more possible) opportunities.

The above paragraphs highlight the fact that the here proposed theoretical framework is only a starting point towards a comprehensive theory of advice taking. However, they also illustrate that the here proposed account can easily be expanded. More importantly, the

expanded information asymmetry account allows to derive a range of specific hypotheses that can be tested experimentally.

6.3 Adaptive advice taking

The current thesis investigates advice taking from an ecological perspective. This constitutes a necessary first step towards assessing the adaptivity of advice taking as it might occur outside the laboratory. The results help to reconcile the immense adaptive potential of advice taking with the presumed maladaptivity of individuals' advice utilization.

The ecological perspective advocated in this thesis claims that advice taking outside the laboratory is an interactive process that affords individuals with the opportunity to actively search for information. Specifically, advice taking here was conceptualized as a sequential process that can be truncated at different points in time, that is, after seeking different amounts of advice.

This ecological perspective has two important implications for the adaptivity of advice taking. First, previous research assumed the underweighting of advice in relation to the own judgment to be maladaptive (e.g., Mannes, 2009). However, the theoretical arguments and simulation data presented suggest a markedly different conclusion. When advice taking is conceptualized as a sequential process with several steps of updating, weighting advice by 0.5 performs poorly. In stark contrast, placing a low weight on the advice yields accuracy that is comparable to the optimal integration via Bayesian updating. What has been assumed a shortcoming in the literature might thus be highly adaptive outside of the constrained research paradigm.

Second, previous research assumed underweighting of advice to result from skewed samples of information that relate more to the own judgment than to the advice (Yaniv, 2004a, 2004b). However, the present experiments demonstrate that individuals do not stick to these information samples but expand them, when given the opportunity. The additional information thus acquired strongly affects individuals' advice utilization. While skewed

samples of information might thus be a prior in many advice taking situations, they are not an unalienable feature of advice taking. Given the opportunity, individuals explore and adapt to the information ecology.

6.3.1 Limitations and future directions

The current thesis provides compelling arguments that advice taking outside the laboratory might very well be much more adaptive than existing research suggests. Nevertheless, some aspects of the adaptivity of advice taking outside the laboratory deserve further attention.

The simulation design assumes individuals' integration of advice to constitute a lasting change of their judgment. Theoretical arguments render this assumption likely. Advice probably constitutes a form of informational influence (Deutsch & Gerard, 1955), which is more likely to lead to conversion rather than compliance (Moscovici, 1980). However, this assumption should be tested empirically. For example, one could assess the time stability of individuals' updated judgments. Likewise, one could assess whether the change in judgment due to advice on a specific item generalizes to related items. In line with the here proposed assumption that individuals' point estimates result from an underlying distribution, research found that individuals do not merely reproduce their estimate when asked on two occasions (Vul & Pashler, 2008). Still, the distribution should change after integrating advice and thus the influence of advice should be measurable even after some delay.

The current thesis suggests habitual low weighting of advice to be effective for two reasons. First, individuals accumulate information over several instances of updating. Second, perfect Bayesian weighting requires a counter of the integrated information, which individuals might not possess. While the accumulation of information is a necessary feature of a sequential process, the assumption that Bayesian updating is not manageable might be overly pessimistic. Specifically, the proposed theoretical account assumes that individuals' confidence mirrors the amount and consistency of their information samples. Higher

confidence could thus indicate that individuals already integrated several pieces of information. Consequently, individuals could use their confidence to assign approximate Bayesian weights. Indeed, individuals have been found to weight advice depending on their confidence (e.g., Gino et al., 2012; Gino & Moore, 2007; Moussaïd et al., 2013). What has not been assessed so far is whether such a weighting yields higher accuracy than equal weighting independent of individuals' confidence.

On a related note, it has been suggested that individuals intuitively rather choose among a set of judgments than integrate them into an aggregate (e.g., Mannes, Soll, & Larrick, 2014; Soll & Larrick, 2009). Furthermore, recent research revealed that choosing the response option associated with the higher level of confidence yields high accuracy (Koriat, 2012b). However, categorical judgments do not allow for aggregation and the associated reduction in error that makes averaging such a successful strategy (e.g., Davis-Stober et al., 2014; Einhorn et al., 1977). Consequently, it appears an intriguing avenue for future research to investigate whether confidence-based choice of a judgment (or advisor) also yields such high accuracy for continuous judgments.

6.4 Ecological validity of advice taking: A disclaimer

The current thesis investigates advice taking from an ecological perspective. The core assumption of this ecological perspective is that the interactive nature of advice taking outside the laboratory affords individuals with the opportunity to seek additional information. Incorporating this opportunity into advice taking research, the current thesis increases the ecological validity of advice taking research. However, there are other ways in which advice taking outside the laboratory likely differs from advice taking as studied in the laboratory. One such difference that is highly important is the nature of the task studied. Advice taking research exclusively focused on continuous judgments (e.g., distance between two cities; Schultze et al., 2015). In contrast, advice outside the laboratory often constitutes a recommendation of one among several response options. Be it the question which new mobile

phone to purchase, which discipline to study or which president to elect, many important decisions are categorical rather than continuous.

The current thesis focuses on continuous judgments as well. This renders it possible to expand existing theoretical accounts (e.g., information asymmetry; Yaniv, 2004a, 2004b) and allows to use established measures (e.g., weight of advice; Harvey & Fischer, 1997; Yaniv, 2004a) which increases the comparability to existing research. Given that this thesis' primary goal was to emphasize that an ecological perspective has important implications for advice taking as investigated previously, this tight link to existing research is highly important.

Nevertheless, the current work can be expanded to account for categorical judgments. The theoretical framework proposed focuses on the amount and consistency of individuals' information samples. These qualities are equally valid for continuous and categorical judgments. For example, an individuals' internal information sample might consist of a number of reasons in support of an option A and an option B. While the absolute number of reasons resembles the amount of information, the ratio of reasons relating to either option could resemble the consistency of the information sample. In this way, additional information integrated into this information sample affects both the amount and consistency of information just as assumed for continuous judgments here. Additional information that supports the current choice increases the consistency (i.e., ratio in favor of current choice) of the internal information sample and should thus lead to an increase in confidence rather than a change in choice. Additional information that proposes the alternative choice decreases the consistency and should thus rather lead to a change in choice. In fact, the assumption that amount and consistency of information determine subjective confidence was developed on the basis of two-alternative forced-choice tasks (Koriat, 2012a). The fact that the determinants of information search in advice taking as identified here converge with the many findings regarding information search in research focusing on choice alternatives (e.g., variability of

the sample; Lejarraga et al., 2012; Mitra et al., 1999; Pachur & Scheibehenne, 2012) supports this analogy.

A crucial constraint to the generalization of advice taking research to categorical decisions lies in the dependent measure of advice taking. The weight of advice provides a straightforward measure of integration of advice into the initial judgment. In comparison, preference reversal after receiving an opposing recommendation constitutes a rather crude measure of advice taking. Potentially, drift diffusion models (Ratcliff, 1978) could deliver a more fine-grained measure of advice taking concerning categorical decisions. For example, the starting point of the diffusion process might resemble an individuals' current state of knowledge concerning either choice option. Consequently, a shift of the starting point over repeated choices could indicate integration of additional information. Such a shift of the starting point is not necessarily reflected in a change of choice but rather in a change of response latency. Previous research suggests that the specific parameters of the diffusion model are indeed sensitive to various manipulations (Voss, Rothermund, & Voss, 2004). Future research could thus use diffusion model parameters as a fine-grained measure of advice taking for categorical judgments.

A potential difference between continuous and categorical judgments concerns the benefits of aggregation. Research on continuous judgments argues that the combination of several judgments via averaging cancels out random error and opposing biases (Davis-Stober et al., 2014; Einhorn et al., 1977). Consequently, integrating continuous advice into the own judgment yields an increase in expected accuracy. In contrast, categorical judgments do not allow for a combination of judgments but only for a choice among judgments. Condorcet's theorem (Condorcet, 1785; Sorokin, West, & Robinson, 1998) holds that the likelihood of a majority judgment being correct increases rapidly with an increase in size of this majority. However, this is only true when the probability of a single judgment being correct is greater than 0.5. It is an open empirical question whether this assumption is true for the categorical

judgments individuals face in their daily lives and whether there are cues indicating such benign environments (Hertwig, 2012). Consequently, it remains to be investigated whether taking advice on categorical judgments increases accuracy.

While the expansion of advice taking research to categorical decision tasks is beyond the scope of the current thesis, the proposed framework constitutes a starting point towards this endeavor.

6.5 Conclusion: Implications of an ecological perspective

Behavior is best understood as an interaction between the mind and the environment (Simon, 1956). It follows that the environment in which a behavior is investigated affects this very behavior. To draw inference with regard to a behavior outside the laboratory, the design of laboratory experiments thus has to be considered carefully.

The current thesis argues that existing research on advice taking is too constrained conceptually. Specifically, it is assumed that the interactive nature of advice taking outside the laboratory affords individuals with the opportunity to seek additional information. A new paradigm is proposed that allows to investigate information search in advice taking. A theoretical framework is proposed that explains interactions between individuals' minds and the information ecology in advice taking. Finally, the adaptivity of advice taking is re-evaluated from this ecological perspective.

The results highlight that this ecological perspective has important empirical, theoretical, and normative implications. Individuals explore the information ecology and adapt their utilization of advice to the information thus received. Consequently, theoretical frameworks of advice taking have to allow for such an interaction between individuals' minds and the information ecology. Finally, current normative considerations have to be changed dramatically when conceptualizing advice taking in line with the here advocated ecological perspective.

Claiming that this thesis returned advice taking to the wild would be too bold even for a concluding statement. Advice taking research still has to go a long way before it can inform social or political decision making. However, the current thesis does show in which direction to proceed.

References

- Baayen, R. H., Davidson, D. J., & Bates, D. M. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*, *59*(4), 390–412. <http://doi.org/10.1016/j.jml.2007.12.005>
- Bandura, A. (1977). *Social learning theory*. Oxford, UK: Prentice-Hall.
- Baumeister, R. F., & Leary, M. R. (1995). The need to belong: Desire for interpersonal attachments as a fundamental human motivation. *Psychological Bulletin*, *117*(3), 497–529. <http://doi.org/10.1037/0033-2909.117.3.497>
- Bearden, J. N., & Rapoport, A. (2005). Operations research in experimental psychology. In *INFORMS Tutorials in Operations Research* (pp. 213–236). CHAP, Hanover, MD: INFORMS. <http://doi.org/doi:10.1287/educ.1053.0011>
- Bernhard, H., Fischbacher, U., & Fehr, E. (2006). Parochial altruism in humans. *Nature*, *442*(7105), 912–5. <http://doi.org/10.1038/nature04981>
- Bonaccio, S., & Dalal, R. S. (2006). Advice taking and decision-making: An integrative literature review, and implications for the organizational sciences. *Organizational Behavior and Human Decision Processes*, *101*(2), 127–151. <http://doi.org/10.1016/j.obhdp.2006.07.001>
- Brunswik, E. (1955). Representative design and probabilistic theory in a functional psychology. *Psychological Review*, *62*(3), 193–217. <http://doi.org/10.1037/h0047470>
- Busemeyer, J. R., & Townsend, J. T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review*, *100*(3), 432–459. <http://doi.org/10.1037/0033-295X.100.3.432>

- Chaiken, S. (1979). Communicator physical attractiveness and persuasion. *Journal of Personality and Social Psychology*, 37(8), 1387–1397. <http://doi.org/10.1037/0022-3514.37.8.1387>
- Clemen, R. (1989). Combining forecasts: A review and annotated bibliography. *International Journal of Forecasting*, 5(4), 559–583. [http://doi.org/10.1016/0169-2070\(89\)90012-5](http://doi.org/10.1016/0169-2070(89)90012-5)
- Collins, E. C., Percy, E. J., Smith, E. R., & Kruschke, J. K. (2011). Integrating advice and experience: Learning and decision making with social and nonsocial cues. *Journal of Personality and Social Psychology*, 100(6), 967–82. <http://doi.org/10.1037/a0022982>
- Condorcet, M. D. (1785). *Essay on the Application of Analysis to the Probability of Majority Decisions*. Paris: Imprimerie Royale.
- Davis-Stober, C., Budescu, D., Dana, J., & Broomell, S. (2014). When is a crowd wise? *Decision*, 1(2), 79–101. <http://doi.org/10.1037/dec0000004>
- Denrell, J. (2005). Why most people disapprove of me: Experience sampling in impression formation. *Psychological Review*, 112(4), 951–978. <http://doi.org/10.1037/0033-295X.112.4.951>
- Deutsch, M., & Gerard, H. (1955). A study of normative and informational social influences upon individual judgment. *The Journal of Abnormal and Social Psychology*, 51(3), 629–636. <http://doi.org/10.1037/h0046408>
- Einhorn, H., Hogarth, R., & Klempner, E. (1977). Quality of group judgment. *Psychological Bulletin*, 84(1), 158–172. <http://doi.org/10.1037/0033-2909.84.1.158>
- Faul, F., Erdfelder, E., Lang, A.-G., & Buchner, A. (2007). G*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, 39(2), 175–191. <http://doi.org/10.3758/BF03193146>

- Fiedler, K. (2000). Beware of samples! A cognitive-ecological sampling approach to judgment biases. *Psychological Review*, *107*(4), 659–676. <http://doi.org/10.1037/0033-295X.107.4.659>
- Fiedler, K., & Hütter, M. (2014). The limits of automaticity. In J. Sherman, B. Gawronski, & Y. Trope (Eds.), *Dual-process theories of the social mind* (pp. 497–513). New York, NY: Guilford Press.
- Galesic, M., Olsson, H., & Rieskamp, J. (2012). Social sampling explains apparent biases in judgments of social environments. *Psychological Science*, *23*(12), 1515–1523. <http://doi.org/10.1177/0956797612445313>
- Galton, F. (1907). Vox populi. *Nature*, *75*(7), 450–451. <http://doi.org/10.1038/075450a0>
- Gino, F. (2008). Do we listen to advice just because we paid for it? The impact of advice cost on its use. *Organizational Behavior and Human Decision Processes*, *107*(2), 234–245. <http://doi.org/10.1016/j.obhdp.2008.03.001>
- Gino, F., Brooks, A. W., & Schweitzer, M. E. (2012). Anxiety, advice, and the ability to discern: Feeling anxious motivates individuals to seek and use advice. *Journal of Personality and Social Psychology*, *102*(3), 497–512. <http://doi.org/10.1037/a0026413>
- Gino, F., & Moore, D. (2007). Effects of task difficulty on use of advice. *Journal of Behavioral Decision Making*, *20*(1), 21–35. <http://doi.org/10.1002/bdm>
- Goldsmith, D., & Fitch, K. (1997). The normative context of advice as social support. *Human Communication Research*, *23*(4), 454–476. <http://doi.org/10.1111/j.1468-2958.1997.tb00406.x>
- Harvey, N., & Fischer, I. (1997). Taking advice: Accepting help, improving judgment, and sharing responsibility. *Organizational Behavior and Human Decision Processes*, *70*(2), 117–133. <http://doi.org/10.1006/obhd.1997.2697>

- Hau, R., Pleskac, T. J., Kiefer, J., & Hertwig, R. (2008). The description-experience gap in risky choice: The role of sample size and experienced probabilities. *Journal of Behavioral Decision Making*, *21*(5), 493–518. <http://doi.org/10.1002/bdm>
- Herrmann, E., Call, J., Hernández-Lloreda, M. V., Hare, B., & Tomasello, M. (2007). Humans have evolved specialized skills of social cognition: The cultural intelligence hypothesis. *Science*, *317*(5843), 1360–1366. <http://doi.org/10.1126/science.1146282>
- Hertwig, R. (2012). Tapping into the wisdom of the crowd - with confidence. *Science*, *336*(6079), 303–304. <http://doi.org/10.1126/science.1221403>
- Hertwig, R. (2015). Decisions from experience. In G. Keren & G. Wu (Eds.), *Blackwell handbook of judgment and decision making* (pp. 239–267). Oxford, UK: Blackwell.
- Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, *15*(8), 534–539. <http://doi.org/10.1111/j.0956-7976.2004.00715.x>
- Hertwig, R., & Erev, I. (2009). The description-experience gap in risky choice. *Trends in Cognitive Sciences*, *13*(12), 517–523. <http://doi.org/10.1016/j.tics.2009.09.004>
- Hertwig, R., & Pleskac, T. J. (2010). Decisions from experience: Why small samples? *Cognition*, *115*(2), 225–237. <http://doi.org/10.1016/j.cognition.2009.12.009>
- Herzog, S. M., & Hertwig, R. (2009). The wisdom of many in one mind: Improving individual judgments with dialectical bootstrapping. *Psychological Science*, *20*(2), 231–238. <http://doi.org/10.1111/j.1467-9280.2009.02271.x>
- Hewstone, M., Rubin, M., & Willis, H. (2002). Intergroup bias. *Annual Review of Psychology*, *53*(1), 575–604. <http://doi.org/10.1146/annurev.psych.53.100901.135109>

- Homans, G. C. (1958). Social behavior as exchange. *American Journal of Sociology*, *63*(6), 597–606. <http://doi.org/10.1086/222355>
- Insko, C. A., Schopler, J., Hoyle, R. H., Dardis, G. J., & Graetz, K. A. (1990). Individual-group discontinuity as a function of fear and greed. *Journal of Personality and Social Psychology*, *58*(1), 68–79. <http://doi.org/10.1037/0022-3514.58.1.68>
- Janis, I. L. (1972). *Victims of groupthink: A psychological study of foreign-policy decisions and fiascoes*. Oxford, UK: Houghton Mifflin.
- Judd, C. M., Westfall, J., & Kenny, D. a. (2012). Treating stimuli as a random factor in social psychology: A new and comprehensive solution to a pervasive but largely ignored problem. *Journal of Personality and Social Psychology*, *103*(1), 54–69. <http://doi.org/10.1037/a0028347>
- Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review*, *80*(4), 237–251. <http://doi.org/10.1037/h0034747>
- Kennedy, J., Kleinmuntz, D., & Peecher, M. (1997). Determinants of the justifiability of performance in ill-structured audit tasks. *Journal of Accounting Research*, *35*(1997), 105–123. <http://doi.org/10.2307/2491456>
- Kenward, M. G., & Roger, J. H. (1997). Small sample inference for fixed effects from restricted maximum likelihood. *Biometrics*, *53*(3), 983–997. <http://doi.org/10.2307/2533558>
- Klayman, J., Soll, J. B., González-Vallejo, C., & Barlas, S. (1999). Overconfidence: It depends on how, what, and whom you ask. *Organizational Behavior and Human Decision Processes*, *79*(3), 216–247. <http://doi.org/http://dx.doi.org/10.1006/obhd.1999.2847>

- Koriat, A. (2012a). The self-consistency model of subjective confidence. *Psychological Review*, *119*(1), 80–113. <http://doi.org/10.1037/a0025648>
- Koriat, A. (2012b). When are two heads better than one and why? *Science*, *336*(6079), 360–362. <http://doi.org/10.1126/science.1216549>
- Larrick, R. P., Mannes, A. E., & Soll, J. B. (2012). The social psychology of the wisdom of crowds. In J. I. Krueger (Ed.), *Frontiers in social psychology: Social judgment and decision making* (pp. 227–242). New York, NY: Psychology Press.
- Larrick, R. P., & Soll, J. B. (2006). Intuitions about combining opinions: Misappreciation of the averaging principle. *Management Science*, *52*(1), 111–127. <http://doi.org/10.1287/mnsc>.
- Le Mens, G., & Denrell, J. (2011). Rational learning and information sampling: On the “naivety” assumption in sampling explanations of judgment biases. *Psychological Review*, *118*(2), 379–392. <http://doi.org/10.1037/a0023010>
- Lee, M., Newell, B., & Vandekerckhove, J. (2014). Modeling the adaptation of search termination in human decision making. *Decision*, *1*(4), 223–251. <http://doi.org/10.1037/dec0000019>
- Lejarraga, T., Hertwig, R., & Gonzalez, C. (2012). How choice ecology influences search in decisions from experience. *Cognition*, *124*(3), 334–342. <http://doi.org/10.1016/j.cognition.2012.06.002>
- Lindskog, M., Winman, A., & Juslin, P. (2013). Naïve point estimation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *39*(3), 782–800. <http://doi.org/10.1037/a0029670>

- Mannes, A. E. (2009). Are we wise about the wisdom of crowds? The use of group judgments in belief revision. *Management Science*, *55*(8), 1267–1279.
<http://doi.org/10.1287/mnsc.1090.1031>
- Mannes, A. E., Soll, J. B., & Larrick, R. P. (2014). The wisdom of select crowds. *Journal of Personality and Social Psychology*, *107*(2), 276–299.
- Masson, M. E. J. (2011). A tutorial on a practical Bayesian alternative to null-hypothesis significance testing. *Behavior Research Methods*, *43*(3), 679–690.
<http://doi.org/10.3758/s13428-010-0049-5>
- McPherson, M., Smith-Lovin, L., & Cook, J. (2001). Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, *27*, 415–444.
<http://doi.org/10.1146/annurev.soc.27.1.415>
- Minson, J. A., Liberman, V., & Ross, L. (2011). Two to tango: Effects of collaboration and disagreement on dyadic judgment. *Personality & Social Psychology Bulletin*, *37*(10), 1325–1338. <http://doi.org/10.1177/0146167211410436>
- Mitra, K., Reiss, M. C., & Capella, L. M. (1999). An examination of perceived risk, information search and behavioral intentions in search, experience and credence services. *Journal of Services Marketing*, *13*(3), 208–228.
<http://doi.org/10.1108/08876049910273763>
- Moll, H., & Tomasello, M. (2007). Cooperation and human cognition: The Vygotskian intelligence hypothesis. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *362*(1480), 639–648.
<http://doi.org/10.1098/rstb.2006.2000>
- Moscovici, S. (1980). Toward a theory of conversion behavior. *Advances in Experimental Social Psychology*, *13*, 209–239. [http://doi.org/10.1016/S0065-2601\(08\)60133-1](http://doi.org/10.1016/S0065-2601(08)60133-1)

- Moussaïd, M., Kämmer, J. E., Analytis, P. P., & Neth, H. (2013). Social influence and the collective dynamics of opinion formation. *PloS One*, *8*(11), 1–8.
<http://doi.org/10.1371/journal.pone.0078433>
- Newell, B. R. (2005). Re-visions of rationality? *Trends in Cognitive Sciences*, *9*(1), 11–15.
<http://doi.org/10.1016/j.tics.2004.11.005>
- Nisbett, R. E., & Wilson, T. D. (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review*, *84*(3), 231–259. <http://doi.org/10.1037/0033-295X.84.3.231>
- Over, H., & Carpenter, M. (2009). Priming third-party ostracism increases affiliative imitation in children. *Developmental Science*, *12*(3), 1–9. <http://doi.org/10.1111/j.1467-7687.2008.00820.x>
- Pachur, T., & Scheibehenne, B. (2012). Constructing preference from experience: The endowment effect reflected in external information search. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *38*(4), 1108–1116.
<http://doi.org/10.1037/a0027637>
- Patt, A. G., Bowles, H. R., & Cash, D. W. (2006). Mechanisms for enhancing the credibility of an adviser: Prepayment and aligned incentives. *Journal of Behavioral Decision Making*, *19*(4), 347–359. <http://doi.org/10.1002/bdm.532>
- Phillips, N. D., Hertwig, R., Kareev, Y., & Avraami, J. (2014). Rivals in the dark: How competition influences search in decisions under uncertainty. *Cognition*, *133*(1), 104–119. <http://doi.org/10.1016/j.cognition.2014.06.006>
- Rand, D. G. (2016). Cooperation, fast and slow: Meta-analytic evidence for a theory of social heuristics and self-interested deliberation. *Psychological Science*, *27*(9), 1192–1206.
<http://doi.org/10.1177/0956797616654455>

- Rand, D. G., Greene, J. D., & Nowak, M. A. (2012). Spontaneous giving and calculated greed. *Nature*, *489*(7416), 427–430. <http://doi.org/10.1038/nature11467>
- Rapoport, A., & Tversky, A. (1966). Cost and accessibility of offers as determinants of optional stopping. *Psychonomic Science*, *4*(1), 145–146. <http://doi.org/10.3758/BF03342220>
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, *85*(2), 59–108.
- Ravazzolo, F., & Røisland, øistein. (2011). Why do people place lower weight on advice far from their own initial opinion? *Economics Letters*, *112*(1), 63–66. <http://doi.org/10.1016/j.econlet.2011.03.032>
- Schultze, T., Rakotoarisoa, A.-F., & Schulz-Hardt, S. (2015). Effects of distance between initial estimates and advice on advice utilization. *Judgment and Decision Making*, *10*(2), 144–171. Retrieved from <http://journal.sjdm.org/14/141112a/jdm141112a.html>
- Schulz-Hardt, S., Frey, D., Lüthgens, C., & Moscovici, S. (2000). Biased information search in group decision making. *Journal of Personality and Social Psychology*, *78*(4), 655–669. <http://doi.org/10.1037/0022-3514.78.4.655>
- Seale, D. A., & Rapoport, A. (1997). Sequential decision making with relative ranks: An experimental investigation of the “secretary problem.” *Organizational Behavior and Human Decision Processes*, *69*(3), 221–236. <http://doi.org/10.1006/obhd.1997.2683>
- Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, *63*(2), 129–138. <http://doi.org/10.1037/h0042769>
- Simonsohn, U. (in preparation). *Two lines: The first valid test of u-shaped relationships*.

- Snizek, J., & Buckley, T. (1995). Cueing and cognitive conflict in judge-advisor decision making. *Organizational Behavior and Human Decision Processes*, 62(2), 159–174.
<http://doi.org/10.1006/obhd.1995.1040>
- Snizek, J., Schrah, G., & Dalal, R. (2004). Improving judgement with prepaid expert advice. *Journal of Behavioral Decision Making*, 17(3), 173–190. <http://doi.org/10.1002/bdm.468>
- Soll, J. B., & Klayman, J. (2004). Overconfidence in interval estimates. *Journal of Experimental Psychology: Learning, Memory, and Cognition.*, 30(2), 299–314.
<http://doi.org/10.1037/0278-7393.30.2.299>
- Soll, J. B., & Larrick, R. P. (2009). Strategies for revising judgment: How (and how well) people use others' opinions. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, 35(3), 780–805. <http://doi.org/10.1037/a0015145>
- Soll, J. B., & Mannes, A. E. (2011). Judgmental aggregation strategies depend on whether the self is involved. *International Journal of Forecasting*, 27(1), 81–102.
<http://doi.org/10.1016/j.ijforecast.2010.05.003>
- Sorkin, R. D., West, R., & Robinson, D. E. (1998). Group performance depends on the majority rule. *Psychological Science*, 9(6), 456–463. <http://doi.org/10.1111/1467-9280.00085>
- Suls, J., Martin, R., & Wheeler, L. (2002). Social comparison: Why, with whom, and with what effect? *Current Directions in Psychological Science*, 11(5), 159–163.
<http://doi.org/10.1111/1467-8721.00191>
- Tajfel, H. (1982). Social psychology of intergroup relations. *Annual Review of Psychology*, 33(1), 1–39. <http://doi.org/http://dx.doi.org/10.1146/annurev.ps.33.020182.000245>
- Tukey, J. W. (1977). *Exploratory data analysis*. Reading, MA: Addison-Wesley.

- Tversky, A., & Koehler, D. (1994). Support theory: A nonextensional representation of subjective probability. *Psychological Review*, *101*(4), 547–567.
<http://doi.org/10.1037/0033-295X.101.4.547>
- Voss, A., Rothermund, K., & Voss, J. (2004). Interpreting the parameters of the diffusion model: An empirical validation. *Memory & Cognition*, *32*(7), 1206–1220.
<http://doi.org/10.3758/BF03196893>
- Vul, E., & Pashler, H. (2008). Measuring the crowd within: Probabilistic representations within individuals. *Psychological Science*, *19*(7), 645–647.
<http://doi.org/10.1111/j.1467-9280.2008.02136.x>
- Williams, K. D., Cheung, C. K. T., & Choi, W. (2000). Cyberostracism: Effects of being ignored over the internet. *Journal of Personality and Social Psychology*, *79*(5), 748–762.
<http://doi.org/10.1037/0022-3514.79.5.748>
- Yaniv, I. (2004a). Receiving other people's advice: Influence and benefit. *Organizational Behavior and Human Decision Processes*, *93*(1), 1–13.
<http://doi.org/10.1016/j.obhdp.2003.08.002>
- Yaniv, I. (2004b). The benefit of additional opinions. *Current Directions in Psychological Science*, *13*(2), 75–79. <http://doi.org/10.1111/j.0963-7214.2004.00278.x>
- Yaniv, I., & Choshen-Hillel, S. (2012). Exploiting the wisdom of others to make better decisions: Suspending judgment reduces egocentrism and increases accuracy. *Journal of Behavioral Decision Making*, *25*(5), 427–434. <http://doi.org/10.1002/bdm>
- Yaniv, I., Choshen-Hillel, S., & Milyavsky, M. (2009). Spurious consensus and opinion revision: Why might people be more confident in their less accurate judgments? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *35*(2), 558–63.
<http://doi.org/10.1037/a0014589>

- Yaniv, I., Choshen-Hillel, S., & Milyavsky, M. (2011). Receiving advice on matters of taste: Similarity, majority influence, and taste discrimination. *Organizational Behavior and Human Decision Processes*, *115*(1), 111–120.
<http://doi.org/10.1016/j.obhdp.2010.11.006>
- Yaniv, I., & Foster, D. (1995). Graininess of judgment under uncertainty: An accuracy-informativeness trade-off. *Journal of Experimental Psychology: General*, *124*(4), 424–432. <http://doi.org/10.1037/0096-3445.124.4.424>
- Yaniv, I., & Foster, D. (1997). Precision and accuracy of judgmental estimation. *Journal of Behavioral Decision Making*, *10*(1), 21–32. [http://doi.org/10.1002/\(SICI\)1099-0771\(199703\)10:1<21::AID-BDM243>3.0.CO;2-G](http://doi.org/10.1002/(SICI)1099-0771(199703)10:1<21::AID-BDM243>3.0.CO;2-G)
- Yaniv, I., & Kleinberger, E. (2000). Advice taking in decision making: Egocentric discounting and reputation formation. *Organizational Behavior and Human Decision Processes*, *83*(2), 260–281. <http://doi.org/10.1006/obhd.2000.2909>
- Yaniv, I., & Milyavsky, M. (2007). Using advice from multiple sources to revise and improve judgments. *Organizational Behavior and Human Decision Processes*, *103*(1), 104–120.
<http://doi.org/10.1016/j.obhdp.2006.05.006>

Appendix

Sequential pair-wise model comparisons

Experiment 1

DV	Predictors	Log Likelihood	Df (denominator)	p	approx. Bayes factor
WOA	Rnd. Intercepts (Participant + Item)	248.57			
	+ Distance	300.8	616.45	< .001	1

Experiment 2

DV	Predictors	Log Likelihood	Df (denominator)	p	approx. Bayes factor
Sample size	Rnd. Intercepts (Participant + Item)	-2317.2			
	+ Distance	-2287.7	832.74	< .001	1
WOA	Rnd. Intercepts (Participant + Item)	-275.57			
	+ Distance	-203.42	789.55	< .001	1
	+ Sample size	-188.28	197.92	< .001	.999
	+ Distance * Sample size	-186.84	798.58	.089	.127
Δ WOA	Rnd. Intercepts (Participant + Item)	-255.73			
	+ Distance	-205.96	789.88	< .001	1
	+ Sample size	-202.31	194.12	.008	.570
	+ Distance * Sample size	-198.48	798.21	.006	.613
Confidence	Rnd. Intercepts (Participant + Item)	-552.38			
	+ Distance	-513.25	416.52	< .001	1
	+ Sample size	-513.25	429.79	.999	.045
	+ Distance * Sample size	-511.47	416.23	.061	.222

Experiment 3

DV	Predictors	Log Likelihood	Df (denominator)	p	approx. Bayes factor
Sample size	Rnd. Intercepts (Participant + Item)	-3067.8			
	+ Distance	-3056.2	1099.53	< .001	.999
WOA	Rnd. Intercepts (Participant + Item)	-151.25			
	+ Distance	-143.36	981.51	< .001	.988
	+ Sample size	-137.63	218.68	< .001	.905
	+ Distance * Sample size	-137.31	992.40	.428	.040
Δ WOA	Rnd. Intercepts (Participant + Item)	-166.29			
	+ Distance	-162.38	980.50	.005	.608
	+ Sample size	-154.22	249.76	< .001	.990
	+ Distance * Sample size	-154.21	990.55	.898	.030
Confidence	Rnd. Intercepts (Participant + Item)	-1567.6			
	+ Distance	-1550.6	1091.48	< .001	.999
	+ Sample size	-1540.7	1119.16	< .001	.998
	+ Distance * Sample size	-1540.7	1087.72	.841	.029

Experiment 4

DV	Predictors	Log Likelihood	Df (denominator)	p	approx. Bayes factor
Sample size	Rnd. Intercepts (Participant + Item)	-1392.3			
	+ Distance	-1376.6	458.63	< .001	1
	+ Initial CI	-1375.2	453.70	.093	.16
	+ Distance * Initial CI	-1372.7	461.15	.027	.343
WOA	Rnd. Intercepts (Participant + Item)	-276.3			
	+ Distance	-231.9	431.78	<.001	1
	+ Sample size	-228.2	202.50	.007	.665
Δ WOA	Rnd. Intercepts (Participant + Item)	-265.26			
	+ Distance	-228.73	431.52	<.001	1
Confidence change	Rnd. Intercepts (Participant + Item)	-108.07			
	+ Distance	-93.80	424.40	<.001	1
	+ Sample size	-92.14	194.14	.075	.200
	+ Initial CI	-70.61	460.96	<.001	1
	+ Distance * Sample size	-66.74	421.36	.006	.690

Experiment 5

DV	Predictors	Log Likelihood	Df (denominator)	p	approx. Bayes factor
Sample size	Rnd. Intercepts (Participant + Item)	-6280.45			
	+ Distance	-6246.69	2422.24	<.001	1
	+ Costs	-6206.98	126.00	<.001	1
	+ Initial CI	-6201.06	2459.28	<.001	.880
	+ Distance * Costs	-6189.75	2420.76	<.001	.999
WOA	Rnd. Intercepts (Participant + Item)	-1442.09			
	+ Distance	-1286.96	2290.62	<.001	1
	+ Sample size	-1266.66	733.63	<.001	1
	+ Costs	-1265.17	167.70	.088	.083
	+ Initial CI	-1265.15	2389.25	.833	.020
	+ Distance * Sample size	-1264.99	2307.55	.568	.023
	+ Distance * Costs	-1263.03	2297.76	.049	.125
	+ Distance * Initial CI	-1262.82	2346.90	.516	.025
	+ Sample size * Costs	-1252.05	1554.99	<.001	.999
	+ Sample size * Initial CI	-1251.93	2399.47	.614	.023
	+ Costs * Initial CI	-1251.54	2397.31	.384	.029
	+ Dist. * Sampl. * Costs	-1251.42	2310.42	.623	.022
	+ Dist. * Sampl. * In. CI	-1251.12	2302.29	.440	.027
	+ Dist. * Costs * In. CI	-1251.10	2363.05	.834	.020
	+ Sampl. * Costs * In. CI	-1247.06	2398.30	.004	.536
	Parsimonious model	-1252.00	2314.14	.129	1
	Δ WOA	Rnd. Intercepts (Participant + Item)	-1389.07		
+ Distance		-1261.74	2292.77	<.001	1
+ Sample size		-1260.61	676.91	.133	.059
+ Costs		-1256.83	169.58	.006	.471
+ Initial CI		-1256.67	2376.13	.571	.023
+ Distance * Sample size		-1254.95	2309.66	.064	.102
+ Distance * Costs		-1252.86	2299.59	.041	.141
+ Distance * Initial CI		-1252.47	2350.83	.379	.029
+ Sample size * Costs		-1252.25	1509.70	.503	.025
+ Sample size * Initial CI		-1252.21	2400.26	.781	.021
+ Costs * Initial CI		-1252.06	2395.23	.586	.023
+ Dist. * Sampl. * Costs		-1252.06	2311.26	.970	.020
+ Dist. * Sampl. * In. CI		-1251.71	2303.53	.406	.028
+ Dist. * Costs * In. CI		-1251.71	2364.81	.939	.020
+ Sampl. * Costs * In. CI		-1247.78	2396.54	.005	.511
Parsimonious model		-1253.00	2315.37	.104	1

Experiment 5 (continued)

DV	Predictors	Log Likelihood	Df (denominator)	p	approx. Bayes factor
Confidence change	Rnd. Intercepts (Participant + Item)	-542.86			
	+ Distance	-513.00	2272.03	<.001	1
	+ Sample size	-504.04	756.97	<.001	.994
	+ Initial CI	-459.78	2393.91	<.001	1
	+ Distance * Sample size	-457.22	2290.64	.024	.207
	+ Costs	-456.74	163.68	.331	.032
	+ Distance * Initial CI	-453.89	2329.14	.017	.261
	+ Distance * Costs	-453.15	2272.05	.226	.041
	+ Sample size * Initial CI	-445.31	2378.26	<.001	.981
	+ Sample size * Costs	-441.95	1652.18	.010	.370
	+ Costs * Initial CI	-441.13	2384.29	.203	.044
	+ Dist. * Sampl. * In. CI	-441.12	2282.28	.900	.020
	+ Dist. * Sampl. * Costs	-438.35	2286.81	.019	.246
	+ Dist. * Costs * In. CI	-430.82	2334.66	<.001	.974
	+ Sampl. * Costs * In. CI	-430.82	2381.71	.936	.020
+ Dist. * Sampl. * Costs * In. CI	-426.18	2309.53	.002	.678	

Note. Degrees of freedom are based on the Kenward-Roger approximation (Kenward & Roger, 1997). For each dependent variable (DV), fit of each model was compared to the fit of the model in the row above. Approximate Bayes factors were calculated using the BIC (Masson, 2011). The approx. Bayes factor indicates the conditional probability of the model being true given the data obtained in comparison to the model in the row above.

Author contributions

This thesis constitutes work that has been documented in three manuscripts, all of which represent joint work with Mandy Hütter. Myself and Mandy Hütter contributed to each manuscript in the following ways.

Ache, F., & Hütter, M. (submitted). The wise crowd and the real world: Low weights emulate Bayesian updating in sequential information uptake.

F. Ache developed the idea. F. Ache and M. Hütter developed the simulation design. F. Ache performed the simulation and the data analysis. F. Ache and M. Hütter interpreted the data. F. Ache drafted the manuscript, and M. Hütter provided critical revisions.

Hütter, M., & Ache, F. (2016).* Seeking advice: A sampling approach to advice taking.

Judgment and Decision Making, 11(4), 401–415.

* shared first authorship

F. Ache and M. Hütter developed the experiments. F. Ache programmed the experiments and performed the data analysis. F. Ache and M. Hütter interpreted the data. F. Ache and M. Hütter drafted the manuscript. M. Hütter developed the idea.

Ache, F., & Hütter, M. (in preparation). Sensitivity to information samples: On the role of confidence in advice seeking and weighting.

F. Ache developed the idea. F. Ache and M. Hütter designed the experiments. F. Ache programmed the experiments and performed the data analysis. F. Ache and M. Hütter interpreted the data. F. Ache drafted the manuscript, and M. Hütter provided critical revisions.

Experiment 2 (Chapter 4.5) has been published as my master thesis at the Ruprecht-Karls-Universität Heidelberg.

Ache, F. (2014). *Corroboration or information: What are we looking for in advice – Expanding advice taking to information search*. Master Thesis, University of Heidelberg.