

OXFORD

collaborative remembering

Theories, Research,
and Applications

EDITED BY


Michelle L. Meade

Celia B. Harris

Penny Van Bergen

John Sutton

Amanda J. Barnier



Collaborative Remembering

Theories, Research, and Applications

Edited by

Michelle L. Meade

Celia B. Harris

Penny Van Bergen

John Sutton

Amanda J. Barnier

OXFORD
UNIVERSITY PRESS

OXFORD

UNIVERSITY PRESS

Great Clarendon Street, Oxford, OX2 6DP,
United Kingdom

Oxford University Press is a department of the University of Oxford.
It furthers the University's objective of excellence in research, scholarship,
and education by publishing worldwide. Oxford is a registered trade mark of
Oxford University Press in the UK and in certain other countries

© Oxford University Press 2018

The moral rights of the authors have been asserted

First Edition published in 2018

Impression: 1

All rights reserved. No part of this publication may be reproduced, stored in
a retrieval system, or transmitted, in any form or by any means, without the
prior permission in writing of Oxford University Press, or as expressly permitted
by law; by licence or under terms agreed with the appropriate reprographics
rights organization. Enquiries concerning reproduction outside the scope of the
above should be sent to the Rights Department, Oxford University Press, at the
address above

You must not circulate this work in any other form
and you must impose this same condition on any acquirer

Published in the United States of America by Oxford University Press
198 Madison Avenue, New York, NY 10016, United States of America

British Library Cataloguing in Publication Data

Data available

Library of Congress Control Number: 2017948691

ISBN 978-0-19-873786-5

Printed and bound by
CPI Group (UK) Ltd, Croydon, CR0 4YY

Oxford University Press makes no representation, express or implied, that the
drug dosages in this book are correct. Readers must therefore always check
the product information and clinical procedures with the most up-to-date
published product information and data sheets provided by the manufacturers
and the most recent codes of conduct and safety regulations. The authors and
the publishers do not accept responsibility or legal liability for any errors in the
text or for the misuse or misapplication of material in this work. Except where
otherwise stated, drug dosages and recommendations are for the non-pregnant
adult who is not breast-feeding

Links to third party websites are provided by Oxford in good faith and
for information only. Oxford disclaims any responsibility for the materials
contained in any third party website referenced in this work.

Contents

Contributors *ix*

Digital media accompanying the book *xiii*

Part I **Introduction**

- 1 Collaborative Remembering: Background and Approaches** 3
*Michelle L. Meade, Celia B. Harris, Penny Van Bergen, John Sutton,
and Amanda J. Barnier*

Part II **Approaches to Studying Collaborative Remembering**

Developmental Perspectives

- 2 Socializing Early Skills for Remembering Through
Parent–Child Conversations During and After Events** 19
Catherine A. Haden, Maria Marcus, and Erin Jant
- 3 Developing Social Functions of Autobiographical Memory
within Family Storytelling** 38
Robyn Fivush, Widaad Zaman, and Natalie Merrill

Cognitive Psychological Perspectives

- 4 Collaborative Inhibition in Group Recall: Cognitive Principles
and Implications** 55
Suparna Rajaram
- 5 Social Aspects of Forgetting** 76
William Hirst and Jeremy Yamashiro
- 6 Memory Conformity Following Collaborative Remembering** 100
Fiona Gabbert and Rebecca Wheeler

Social Psychological Perspectives

- 7 The Socially Shared Nature of Memory: From Joint Encoding
to Communication** 113
Gerald Echterhoff and René Kopietz

Lifespan Perspectives

- 8 Collaborative Remembering and Reminiscence in Older Adults** 135
Linda A. Henkel and Alison Kris
- 9 Memories and Identities in Conversation with Dementia** 157
Nicole Müller and Zaneta Mok

Discursive and Interactional Perspectives

- 10 Multimodal Processes of Joint Remembering in Complex Collaborative Activities 177
Lucas M. Bietti and Michael J. Baker
- 11 Contextualizing Autobiographical Remembering: An Expanded View of Memory 197
Steven D. Brown and Paula Reavey
- 12 Collaborative Processes in Neuropsychological Interviews 216
Chris McVittie and Andy McKinlay

Philosophical Perspectives

- 13 Collaborative Memory Knowledge: A Distributed Reliabilist Perspective 231
Kourken Michaelian and Santiago Arango-Muñoz
- 14 Group-level Cognizing, Collaborative Remembering, and Individuals 248
Robert A. Wilson

Function, Culture, and Identity Perspectives

- 15 Remembering Good and Bad Times Together: Functions of Collaborative Remembering 261
Monisha Pasupathi and Cecilia Wainryb
- 16 Collective Memory: How Groups Remember Their Past 280
Magdalena Abel, Sharda Umanath, James V. Wertsch, and Henry L. Roediger, III
- 17 Culture in Collaborative Remembering 297
Qi Wang

Part III Applications of Collaborative Memory

- 18 Encouraging Collaborative Remembering Between Young Children and Their Caregivers 317
Elaine Reese
- 19 Parent–Child Construction of Personal Memories in Reminiscing Conversations: Implications for the Development and Treatment of Childhood Psychopathology 334
Karen Salmon
- 20 Forensic Applications of Social Memory Research 353
Helen Paterson and Lauren Monds
- 21 Digital Media and the Precarity of Memory 371
Andrew Hoskins
- 22 Design Applications for Social Remembering 386
Elise van den Hoven, Mendel Broekhuijsen, and Ine Mols
- 23 Applications of Collaborative Memory: Patterns of Success and Failure in Individuals with Hippocampal Amnesia 404
Rupa Gupta Gordon, Melissa C. Duff, and Neal J. Cohen

- 24** Collaborative Memory Interventions for Age-Related and Alzheimer's Disease-Related Memory Decline 422
Helena M. Blumen
- 25** Collaborative Remembering in Dementia: A Focus on Joint Activities 436
Lars-Christer Hydén and Mattias Forsblad

Part IV **Conclusion**

- 26** Concluding Remarks: Common Themes and Future Directions 459
*Michelle L. Meade, Celia B. Harris, Penny Van Bergen, John Sutton,
and Amanda J. Barnier*

Author Index 467

Subject Index 483

Contributors

Magdalena Abel, Department of Experimental Psychology, Regensburg University, Germany

Santiago Arango-Muñoz, Department of Philosophy, University of Antioquia, Medellín, Colombia

Michael J. Baker, Department of Economic and Social Sciences, CNRS-Telecom Paristech, i3 Interdisciplinary Institute on Innovation, Paris, France

Amanda J. Barnier, Department of Cognitive Science, Macquarie University, Sydney, Australia

Lucas M. Bietti, Institute of Work and Organizational Psychology, University of Neuchâtel, Neuchâtel, Switzerland

Helena M. Blumen, Department of Medicine and Department of Neurology, Albert Einstein College of Medicine, Bronx, NY, USA

Mendel Broekhuijsen, Department of Industrial Design, Eindhoven University of Technology (TU/e), Netherlands; and School of Software, Faculty of Engineering and Information Technology (FEIT), University of Technology Sydney (UTS), Australia

Steven D. Brown, School of Business, University of Leicester, UK

Neal J. Cohen, Beckman Institute, University of Illinois at Urbana-Champaign, Champaign County, IL, USA

Melissa C. Duff, Department of Hearing and Speech Sciences, Vanderbilt University Medical Center, Nashville, USA

Gerald Echterhoff, Department of Psychology, Westfälische Wilhelms-Universität Münster (WWU Münster), Germany

Robyn Fivush, Department of Psychology, Emory University, Atlanta, GA, USA

Mattias Forsblad, Department of Behavioural Sciences and Learning, Division of Psychology, Linköping University, Sweden

Fiona Gabbert, Department of Psychology, Goldsmiths University of London, UK

Rupa Gupta Gordon, Department of Psychology, Augustana College, Rock Island, IL, USA

Catherine A. Haden, Department of Psychology, Loyola University Chicago, IL, USA

Celia B. Harris, Department of Cognitive Science, Macquarie University, Sydney, Australia

Linda A. Henkel, Department of Psychology, Fairfield University, CT, USA

William Hirst, Department of Psychology, New School for Social Research, New York, NY, USA

Andrew Hoskins, College of Social Sciences, University of Glasgow, UK

Elise van den Hoven, School of Software, Faculty of Engineering and Information Technology (FEIT), University of Technology Sydney (UTS), Australia; Department of Industrial Design, Eindhoven University of Technology (TU/e), Netherlands; Jordanstone College of Art and

Collaborative Memory Knowledge: A Distributed Reliabilist Perspective

Kourken Michaelian and Santiago Arango-Muñoz

Collaborative remembering, in which two or more individuals cooperate to remember together, is an ordinary occurrence. We will argue that, ordinary though it may be, it challenges traditional understandings of *remembering* as a cognitive process unfolding within a single subject, as well as traditional understandings of *memory knowledge* as a justified memory belief held within the mind of a single subject. Collaborative memory has come to be a major area of research in psychology, but it has so far not been investigated in epistemology. In this chapter, we attempt an initial exploration of the epistemological implications of collaborative memory research, taking as our starting point the “extended knowledge” debate which has resulted from the recent encounter between extracranialist theories of cognition and externalist theories of knowledge (Carter *et al.*, 2014; Carter *et al.*, forthcoming).

Various forms of socially and technologically augmented memory have played important roles in the extended knowledge debate, but the debate has so far not taken collaborative memory, in particular, into account. We will argue that the findings of research on collaborative memory support a novel externalist theory of knowledge: distributed reliabilism. In addition to differing from approaches such as Giere’s (2004) treatment of distributed cognition and Thagard’s (1997) treatment of collaborative knowledge, distributed reliabilism departs in two important respects from both traditional reliabilism (Goldman, 2012) and updated theories such as extended (Goldberg, 2010) and social reliabilism (Goldman, 2014). First, it acknowledges that belief-forming processes may extend extracranially to include processing performed both by other subjects and by technological artifacts. Second, it acknowledges that distributed sociotechnical systems themselves may be knowing subjects. Overall, then, the main aim of the chapter is to draw out the philosophical implications of psychological research on collaborative memory. But our argument will also suggest that it may be useful to broaden the standard conception of collaborative memory to include not only the sorts of direct interactions among subjects that have been the focus of psychological research so far but also a range of more indirect, technology-supported and—mediated interactions, and it thus has implications for psychology as well.

The Extended Knowledge Debate

We will hone in gradually on collaborative memory. We begin, in this section, with a review of the basic outlines of the extended knowledge debate, introducing a distinction between extended and distributed memory systems. In the next section, we will look at the epistemology of extended memory. In the last section, we will look at the epistemology of distributed memory; it is here that collaborative memory research and the distributed reliabilist epistemology that it supports take center stage.

Externalist Theories of Knowledge

What is *knowledge*? The question is controversial, but there is a rough consensus among epistemologists that knowledge has three ingredients, the first two being belief and truth. Beyond this consensus, there is little agreement: it is clear that knowing that *P* requires more than merely having a true belief that *P*, but it is far from obvious what else might be required. Much of the action in epistemology therefore concerns the third ingredient of knowledge.¹

In generic terms, the third ingredient of knowledge is usually referred to as “justification.” The available theories of justification fall into two broad categories. Internalist theories are so called because they maintain that justification is internal to the knowing subject, in the sense that, if a given factor affects the justificatory status of a belief, the subject can, at least in principle, become aware of the presence or absence of that factor by means of introspection alone. So, for example, coherentists argue that the justificatory status of a belief is determined by its place in the subject’s overall web of beliefs. This is an internalist theory because one is supposed to be able to look within oneself and grasp the relationships among his beliefs.

The *externalist* theories on which we will focus here, in contrast, maintain that justification need not be internal. Among the most influential forms of externalism is reliabilism, which is motivated in part by the worry that internalist standards for knowledge are too high—so high, in fact, that subjects virtually never meet them. It is, for example, doubtful that we are actually capable of grasping the relationships among all of our beliefs, as coherentism requires. Reliabilists argue that the factor that determines whether a given belief is justified is simply the reliability of the process that produced it. For example, if a subject’s memory system tends to produce more true propositions than false propositions, then the process of forming beliefs by accepting what he seems to remember is reliable for that subject, and his memory beliefs can be said to be justified and, when they are true, to amount to knowledge. This is an externalist theory because we cannot determine the reliability of our belief-producing processes simply by looking within ourselves.

The form of reliabilism just described is known as *process* reliabilism. Process reliabilism continues to be influential, but a rival approach known as *virtue* reliabilism has increasingly eclipsed it in recent years. Just as process reliabilism is motivated in part by the worry that internalist standards for knowledge are too high, virtue reliabilism is motivated in part by the worry that process reliabilist standards for knowledge are too low. Consider the case of a subject who remembers his appointments by consulting an electronic calendar. Though the calendar appears to work normally, in fact it frequently malfunctions, indicating appointment times at random. Unbeknownst to the subject, however, his hard-working secretary systematically reschedules his appointments so that they occur at the times indicated by the calendar. The subject’s beliefs are formed by a reliable process (he is always on time for his appointments!), but intuitively he does not *know* the times of his meetings.

One natural thought about why this subject lacks knowledge is that the reliability of the relevant belief-producing process has nothing to do with his own cognitive abilities. This thought—known as the ability intuition (Pritchard, 2010)—is what motivates virtue reliabilists, who argue that knowledge can be analyzed, at least in part, in terms of the degree to which the formation of a true belief by the subject is due to his cognitive abilities or agency. The version of virtue reliabilism which has received the most attention in the context of the extended knowledge debate is what we will refer to as *weak* virtue reliabilism. As Pritchard (2010) formulates it, weak virtue reliabilism says that, if subject *S* knows proposition *P*, then “*S*’s true belief that *P* is the result of a reliable

¹ We set aside issues concerning the Gettier problem and epistemic luck.

belief-forming process which is appropriately integrated within S's cognitive character such that his cognitive success is to a significant degree creditable to his cognitive agency."

Extracranialist Theories of Cognition

The extended knowledge debate concerns the compatibility of epistemologies such as weak virtue reliabilism with theories of distributed or extended cognition. *Distributed* cognition refers to the tradition of research in cognitive science which focuses primarily on cognition in complex socio-technical systems involving multiple human and technological components (Hutchins, 1995; 1996). *Extended* cognition, in turn, refers to the current in philosophy of mind which focuses primarily on cognition in systems centered on a single human subject augmented by external technological or, in some cases, social resources (Clark & Chalmers, 1998; Clark, 2008). For present purposes, we have no stake in the debate over the relative merits of the distributed and extended frameworks; what matters is the core *extracranialist* claim, common to both frameworks and rejected by intracranialists (Adams & Aizawa, 2008), that cognitive processes are best viewed as (sometimes) exceeding the bounds of the individual brain (see also van den Hoven *et al.*; Wilson, Chapters 14 and 22).

Cases of externally augmented memory are at the heart of arguments for extracranialism. Extended cognition theorists tend to emphasize relatively simple cases in which external technological resources are recruited to supplement individual memory. Clark and Chalmers (1998), for example, ask us to consider Otto, a hypothetical Alzheimer's patient, who, to compensate for his failing memory, regularly consults a notebook in which he records information that he expects to need, arguing that, if we want to understand Otto's behavior, we must treat him and his notebook as constituting a single extended memory system. In addition to this sort of technological extension, Clark and Chalmers consider the possibility of social extension, but, in both cases, the unit of analysis for the extended cognition theorist remains the technologically or socially augmented individual.

Distributed cognition theorists, in contrast, tend to emphasize remembering in more complex sociotechnical systems. Hutchins' (1995) classic treatment of the system, consisting of the members of the human crew plus the relevant instruments, responsible for allowing an airplane cockpit to remember its speeds, is representative. Hutchins argues that the memory of the cockpit cannot be reduced to or inferred from the memory of any individual crew member. Nor can it be reduced to the memories of all human crew members. Instead, if we want to be able to understand the behavior of the cockpit, we need to treat it—including both humans and artifacts—as a single distributed memory system.

In line with both extended and distributed frameworks, we want to emphasize that an adequate understanding of memory will take both social and technological dimensions into account. There are several reasons for this emphasis. First, while theoretical arguments for distributed cognition have, to some extent, informed collaborative memory research (Michaelian & Sutton, 2013), the latter has for the most part focused on purely social systems, as opposed to hybrid sociotechnical systems, and one aim of this chapter is to suggest an expansion of the focus of collaborative memory research to include technological, as well as social dimensions of remembering (see also Hoskins; van den Hoven *et al.*, Chapters 21 and 22). Second, though there are exceptions (e.g., Palermos, 2015), most work in social epistemology has focused on remembering in social groups, ignoring the contributions of technological artifacts, and another aim of this chapter is to emphasize the relevance of technology for social epistemology. Finally, while social epistemology has arguably neglected technology, the extended knowledge debate has grown out of individual rather than social epistemology, and its primary focus has been on technologically augmented systems; a further aim of this chapter is to urge participants in the extended knowledge debate not to lose sight of the importance of the social dimension of extracranial cognition.

Extracranialism and Knowledge

We have introduced two distinctions: between internalism and externalism, and between intracranialism and extracranialism. The extended knowledge debate is about the relationship between externalism and extracranialism. It is tempting to think that these two views fit together naturally. Combining them, however, turns out to be more difficult than one might expect, and there is no guarantee that all forms of externalism are compatible with extracranialism. In particular, while all that process reliabilists care about is whether a given belief was produced by a reliable process, virtue reliabilists care not only about whether the belief was produced by a reliable process, but also about whether that process was integrated with the agent's cognitive character in such a manner that his cognitive success was attributable to his cognitive agency to a significant degree. In cases where the majority of the cognitive processing resulting in the formation of a true belief was performed by extracranial resources, it may be difficult to maintain that the process in question was integrated in the agent's character in the requisite manner.

Many of the contributions to the extended knowledge debate are thus concerned with showing that theories like weak virtue reliabilism are indeed compatible with extracranialism. Pritchard (2010) has argued that, even in simple cases, the subject's cognitive agency may make a significant contribution to the formation of a true belief. For example, due to Otto's active role in setting up, maintaining, and using his notebook, the notebook is arguably integrated into his cognitive character in a way that allows his cognitive success, when he forms a true belief by relying on the notebook, to be significantly creditable to his cognitive agency. If this is right, then when Otto forms a true belief by relying on the notebook, the belief may, given weak virtue reliabilism, amount to knowledge. This strategy is fairly plausible with respect to the Otto case, but it is unclear how far it generalizes. What the partisan of weak virtue reliabilism would need to show is that, in every case (or at least in central cases) in which the subject relies on an external resource in forming a true belief, the subject's own cognitive agency plays a significant role. The extended knowledge debate has tended to focus on hypothetical cases of external memory that are simplified to the point of being stylized at best and unrealistic at worst. We argue here that bringing the debate into closer contact with research on extended and distributed memory makes it clear that the strategy does *not*, in fact, generalize.

The Epistemology of Extended Memory

We first briefly consider extended memory, arguing that the role of the subject's agency in such systems is highly variable. This conclusion undermines weak virtue reliabilism, and, after considering and ruling out alternative forms of extended virtue reliabilism, we will argue that a form of extended process reliabilism is better able to accommodate extended memory.

External Memory

In previous work (Michaelian, 2012), we have argued that any subject capable of remembering must cope with two problems. First, the *selection* problem: when faced with a need for information, the subject must choose between retrieving it from memory and relying on some other cognitive resource. For example, he might choose between remembering the result of a calculation and performing the calculation anew. Second, the *endorsement* problem: when the subject does retrieve information from memory, he must determine whether to endorse it or, rather, to reject it. We initially conceived of these problems without reference to external memory; taking external memory into account complicates things somewhat (Arango-Muñoz, 2013). The selection problem becomes the *extended* selection problem, which arises because the subject must choose between relying on various internal and external resources. For example, he might choose

between asking someone else for the answer to a question, looking it up on the Internet, and relying on his own internal memory. Similarly, the (internal) endorsement problem is supplemented by the *external* endorsement problem, which arises because the subject must determine whether or not to endorse information retrieved from an external resource. For example, after retrieving information online, the subject must determine whether to endorse this version of the facts or to search for another one.

How do subjects solve these problems? Our suggestion is that metacognition often plays a key role. Metacognition can be defined as the capacity to monitor and control cognitive processes (Nelson & Narens, 1990), and it has been argued that doxastic (Michaelian, 2012), subpersonal (Clark, 2015), or feeling-based (Arango-Muñoz, 2013) metacognition may feed into selection and endorsement. We focus here on feeling-based metacognition.² Notable examples of metacognitive feelings include the feeling of knowing, the feeling of confidence, the feeling of error, the feeling of forgetting, and the tip-of-the-tongue phenomenon (Arango-Muñoz & Michaelian, 2014). On the one hand, these feelings influence selection by motivating the subject either to retrieve information from internal memory or to seek out an appropriate external resource. A feeling of knowing, for example, might motivate him to search his internal memory. On the other hand, they can also motivate the subject to endorse retrieved information. A feeling of confidence, for example, might motivate him to endorse information received from another subject or from a technological resource.

It is tempting to conclude that, if subjects often rely on metacognitive feelings to solve the selection and endorsement problems, their agency usually plays a significant role in the formation of beliefs based on external memory. In fact, however, the role played by agency appears to be highly variable. Consider first the extended selection problem. An internal strategy is normally quicker but costlier, since it requires working memory and attention; moreover, it is often less accurate. An external strategy is often more accurate but has sensory-motor costs and is normally less efficient; that is, it takes longer. In some cases, subjects successfully assess such trade-offs by relying on metacognitive feelings (Reder, 1987). For example, Kalnikaitė and Whittaker (2007) have shown that reliance on external memory correlates with subjective reports of the feeling of uncertainty (i.e., we resort to external memory primarily when we feel uncertain about internal memory). This result dovetails nicely with Sparrow *et al.*'s (2011) finding that subjects resort to searching for information on the Internet primarily when they do not know the answer to a question. Although the subject's agency clearly plays an important role in these cases, there are other cases where selection seems to be effectively blind and the contribution of the subject's agency is correspondingly minimal. A subject may gain knowledge by asking someone on the street for directions or the time, even if she selects her informant more or less randomly. Indeed, in some cases the subject may not, strictly speaking, *select* at all—she simply relies on a habitual source and acquires knowledge without displaying any agency.³

The role played by agency in the external endorsement problem is similarly variable. Subjects often rely on metacognitive feelings when determining whether to endorse externally retrieved information. Thus, a feeling of truth would motivate endorsement (Reber & Unkelbach, 2010), whereas a feeling of error would motivate rejection (Gangemi *et al.*, 2015). However, there are also many cases in which agency does not appear to play a significant role. In the case of testimony,

² This reflects our assessment of the relative importance of these forms of metacognition, but a more extensive treatment would explore the roles of all forms of metacognition.

³ Our argument here presupposes a view of agency on which habitual action suggests a lack of agency, but there is room for opposing views (e.g., Sutton *et al.* 2011).

for example, there seems to be a default tendency to endorse information provided by others (Michaelian, 2010): intuitively, we feel that we monitor testifiers for deception, but in fact there is considerable evidence that we do not routinely monitor for dishonesty and are not good at detecting it when we do monitor (Vrij, 2008). In many technological cases, agency likewise seems to play a limited role. Sensitivity to the inaccuracy of retrieved information depends on prior knowledge (Sparrow & Chatman, 2013) and, when they lack relevant prior knowledge, subjects will be unable to meaningfully evaluate retrieved information.

Extended Reliabilism

As far as the role of agency in extended memory systems goes, then, we are faced with a continuum of cases. At one extreme, we have cases in which the subject clearly plays an active role in both selection and endorsement. In such cases, the contribution of the subject's agency presumably meets the standard of significance that proponents of weak virtue reliabilism have in mind. At the other extreme, however, we have a class of cases in which the subject's agency appears to make only a negligible contribution to solving both the selection and the endorsement problems. In such cases, it is at best unclear whether the contribution of the subject's agency meets the relevant standard of significance.

In response to this problem, we might opt to replace weak virtue reliabilism with a form of extended virtue reliabilism. Green (2012), for example, considers a view similar to the following *socially extended* virtue reliabilism: if *S* knows that *P*, then the abilities that contribute to the formation of *S*'s belief that *P* deserve significant credit for *S*'s knowing *P*, whether those abilities are contributed solely by *S* or also by other subjects. Taking into account the fact that the same metacognitive feelings may shape our reliance on both social and technological resources suggests moving to an alternative form of *socially and technologically extended* virtue reliabilism: if *S* knows that *P*, then the abilities that contribute to the formation of *S*'s belief that *P* deserve significant credit for *S*'s knowing *P*, whether those abilities are contributed solely by *S* or also by other human and nonhuman resources.

Socially and technologically extended virtue reliabilism, however, raises a problem: intuitively, credit for an accomplishment can be assigned only where there is agency, and the nonhuman components of extended cognitive systems will in general not qualify as agents in a suitably rich sense. The need to assign credit to the nonhuman components of an extended system might be avoided by modifying socially and technologically extended virtue reliabilism so that it instead refers to assigning credit to the extended system as a whole. But such a view would raise much the same problem. Credit for an accomplishment can be assigned only where there is agency, and (as we will argue) extended cognitive systems will in general not qualify as agents in a rich sense (but see Kirchhoff & Newsome, 2012). Socially and technologically extended virtue reliabilism thus appears not to be a viable option.

Given the close relationship between virtue reliabilism and process reliabilism, it is natural at this point to consider adopting some form of the latter, thereby dispensing with references to agency in our account of knowledge. Doing so will require us to sacrifice the ability intuition. But while adopting process reliabilism comes at a cost, that cost might be worth paying in order to avoid incompatibility with the reality of extended memory systems.

If we adopt a form of process reliabilism, which form of the theory should we adopt? In principle, we might simply return to an intracranialist form of process reliabilism, but this move would be implausible in light of the indispensable role of extracranial resources, including external memory, in enabling subjects to form true beliefs. That role lends support, for example, to Goldberg's (2010) *socially extended* process reliabilism, which sees epistemically relevant belief-forming processes as extending into other human subjects but not into nonhuman technological resources.

There does not, however, appear to be a principled reason for distinguishing between processing performed by humans and processing performed by technological artifacts (Michaelian, 2014a), and it is thus more natural to move to *socially and technologically extended* process reliabilism, which sees epistemically relevant belief-forming processes as extending both into human subjects and into nonhuman technological resources. Like socially and technologically extended virtue reliabilism, socially and technologically extended process reliabilism is consistent with the fact that the same metacognitive feelings may shape our reliance on social and technological resources. It thus represents a further step in the direction of an epistemology capable of accommodating knowledge based on socially and technologically augmented memory. We will argue in the next section, however, that socially and technologically extended process reliabilism does not go far enough, as it fails to take into account the possibility that, even if they do not qualify as agents, distributed memory systems may qualify as knowing subjects.

The Epistemology of Distributed Memory

When one remembers with others, those others may play a variety of different roles. So far, we have focused on one relatively simple role: others can serve as sources of information which is incorporated into one's own memories, that is, they can serve as external memory stores. Cases in which others serve as external memory stores already involve partly extracranial processes, in the sense that the behavior of a subject who relies on an external memory store cannot be fully understood if we ignore the flow of information between him and the store. When the subject has reliable access to the store, for example, he may tend to (internally) remember how to retrieve information from the store, rather than remember the content of the information itself (Sparrow *et al.*, 2011). While remembering in extended memory systems is a partly extracranial process, such systems remain subject-centered, in the sense that, though the subject's agency may sometimes play a negligible role in directing the overall memory process, he remains the ultimate locus of control in the system. We turn now to the more complex interactions among subjects (and artifacts) that give rise to distributed (rather than extended) memory systems. In many such systems, including those studied in collaborative memory research, there is no longer a clear locus of control, and we will argue that, even if one component of a given system plays a particularly central role (e.g., that of a dominant narrator; Cuc *et al.*, 2006), it is often the system as a whole that is appropriately viewed as the remembering entity.

Distributed Memory

Collaborative memory research covers a broad spectrum of social influences on memory (Rajaram & Pereira-Pasarin, 2010). At one extreme are cases (e.g., socially shared retrieval-induced forgetting) in which remembering, though subject to important social influences, can still be understood in intracranial terms. Cases at this extreme align with what Wilson (2005) refers to as the *social manifestation* thesis, according to which remembering is a social process in the sense that an individual's remembering sometimes depends for its occurrence on the presence of other people. Given our focus here, we will, strictly speaking, be concerned not with the social manifestation thesis but with what we can refer to as the *social and technological manifestation* thesis, according to which remembering is a social and technological process in the sense that an individual's remembering sometimes depends for its occurrence or features on the presence of other people or technological resources. Partway along the spectrum are cases (e.g., collaborative inhibition) which can only be adequately described by taking the group of remembering subjects into account but in which it may not be necessary to treat the group itself as a remembering agent. At the other extreme are cases (e.g., transactive memory) in which it appears to be necessary to

treat the group itself as constituting a remembering entity. Cases at this extreme align with what Wilson refers to as the *group mind* thesis, according to which remembering is a social process in the sense that groups themselves may sometimes remember. Given our focus here, we will, again, be concerned not with the group mind thesis but with what we can refer to as the *sociotechnical system mind* thesis, according to which remembering is a social and technological process in the sense that sociotechnical systems themselves may sometimes remember.⁴

We begin at one extreme of the spectrum. Research on the social contagion of memory has firmly established that exposure to information provided by others can reshape one's memory (Roediger *et al.*, 2001; see also Gabbert & Wheeler; Paterson & Monds, Chapters 6 and 20). While the social contagion paradigm views social influence on memory as a unidirectional process, with influence flowing from speaker to audience, research on conversational remembering more broadly emphasizes that social influence is in fact a bidirectional process, allowing conversational remembering to reshape the memories of both audience and speaker. This bidirectional process is illustrated by research on retrieval-induced forgetting (Hirst and Echterhoff, 2008; see also Hirst & Yamashiro, Chapter 5). Within-individual retrieval-induced forgetting (WI-RIF) occurs when retrieval of an item by a subject both strengthens his memory for the retrieved item and causes forgetting of other, related items. Socially shared retrieval-induced forgetting (SS-RIF) occurs when retrieval of an item by a speaker similarly—but more surprisingly—also causes forgetting of other, related items in his listeners (Cuc *et al.*, 2007; Stone *et al.*, 2012). This effect apparently occurs because listeners concurrently retrieve the same information as the speaker. In the context of conversational remembering, in other words, WI-RIF gives rise to SS-RIF. Conversational remembering thus tends to cause agents to converge on common memory representations, giving rise to a simple form of shared memory (Fagin *et al.*, 2013; Stone *et al.*, 2013).

In line with our suggestion that the domain of collaborative memory research be broadened to include the effects of interactions with technologies as well as people, we note that external memory technologies might in principle give rise to effects analogous to SS-RIF. Consider a small-scale distributed memory system consisting of a human subject interacting with a web browser configured to prompt him with possible completions (drawn from his browsing history) when he begins to type a web address in the navigation bar. Being prompted in this manner might well result in strengthened memory for certain addresses and forgetting of other, related addresses. Moreover, if the web browser is designed to respond adaptively to the user, it, too, might be subject to a form of retrieval-induced forgetting. A sufficiently sophisticated web browser will modify future suggested completions on the basis of past accepted completions, strengthening its “memory” for accepted completions and eventually “forgetting” rejected completions. Thus, such a distributed memory system might display patterns of shared memory analogous to those observed in SS-RIF in groups of human subjects.

Similar effects might also occur in larger-scale systems involving multiple human and technological components. Consider a system consisting of multiple users interacting via a social networking website. On such websites, the prominence with which a given item is displayed to a given user is often in part a function of the amount of attention that it has received from other users in his social network. As some items are displayed more prominently and others less prominently

⁴ By offering the examples of collaborative inhibition and transactive memory, we do not mean to suggest that, as we move along the spectrum, the effects of collaboration are increasingly beneficial. It is consistent with our view that (as a matter of contingent fact) more collaborative forms of memory are less beneficial and that less collaborative forms of memory are more beneficial. Our concern here is not with the benefits provided by collaboration but with the extent to which forms of collaborative memory satisfy the group mind thesis.

or not at all, the user's memory may be reshaped through retrieval-induced forgetting, with the former becoming easier for him to recall and the latter being forgotten. This, in turn, may feed back into his own activity on the website, determining how much attention he pays to various items. Since his activity influences the items that other users in his social network see, they may be subject to a similar effect, leading to a form of widely distributed SS-RIF.

Setting aside speculation on SS-RIF in hybrid sociotechnical systems, the key point is that, in groups—such as married couples—characterized by frequent, ongoing conversational remembering, SS-RIF tends to give rise to shared memory representations. There is therefore a clear sense in which SS-RIF research is concerned with distributed memory systems. This form of distributed memory, however, can be accounted for by the social (or the social and technological) manifestation thesis. The systems in question can be understood entirely in terms of individual-level mechanisms such as covert retrieval. There is no sense in which the system itself remembers, above and beyond the sense in which the individuals who constitute it remember.

While research on SS-RIF is concerned with what individuals remember in social contexts, research on collaborative recall (Basden *et al.*, 1997; Weldon & Bellinger, 1997) is in an important sense concerned with what social groups themselves remember. Collaborative recall thus brings us to the midpoint of the spectrum between mere social manifestation and genuine group minds. Research on collaborative recall has consistently identified two superficially opposed effects (Barnier & Sutton, 2008; Betts & Hinsz, 2010; Weldon, 2000; see also Blumen; Henkel & Kris; Rajaram, Chapters 4, 8, and 24). On the one hand, a form of what Larson (2009) refers to as “weak synergy” occurs when the quantity of information recalled by a group remembering together is greater than that recalled by any of its members individually. This presumably occurs in part because group members remember nonoverlapping sets of items. On the other hand, collaborative inhibition occurs when the quantity of information recalled by a group remembering together is less than that recalled by a nominal group (i.e., than the sum of the nonredundant items individually recalled by the same number of people, working independently). Collaborative inhibition occurs because individuals recall less when remembering in groups than when remembering individually. A number of possible mechanisms may be at work here (Rajaram & Pereira-Pasarin, 2010; Thompson, 2008), but the key mechanism appears to be retrieval disruption, in which the retrieval strategies employed by different group members interfere with each other (Basden *et al.*, 1997). The retrieval disruption hypothesis is supported by the observation that, in conditions in which group members are likely to use similar retrieval strategies, collaborative inhibition can be overcome or even reversed, allowing the real group to recall more than the nominal group. Such “collaborative facilitation” has been found with groups of experts in a given domain (Meade *et al.*, 2009), as well as with married couples (Harris *et al.*, 2014). Overall, the potential benefits of collaborative recall are likely to be maximized, and its costs minimized, in stable, continuing groups (Harris *et al.*, 2011; 2014).

It is unclear whether research on collaborative recall supports a move from the social (and technological) manifestation thesis to the group mind (or sociotechnical system mind) thesis. On the one hand, collaborative inhibition cannot be adequately described without reference to what a real—as opposed to a nominal—group remembers: if our focus is on remembering as a strictly intracranial process or on groups as mere aggregates of noninteracting individuals, we will miss the effect entirely. On the other hand, the collaborative recall paradigm may not capture a genuinely emergent form of group-level memory, since, in the cases at issue, the reference to the group does no real work: nominal groups remember more than individuals, but they do so in a purely aggregative sense (Pavitt, 2003; Theiner *et al.*, 2010). Cases of collaborative inhibition may be different, since the effect occurs only in real groups, causing groups to remember less than they otherwise might. Nevertheless, the retrieval disruption hypothesis identifies an individual-level

mechanism as being responsible for the occurrence of the effect, so even here it is unclear whether we are entitled to view remembering as a group-level process.

The finding that collaborative facilitation sometimes occurs in stable, continuing groups, however, does bring us a step closer to a view of remembering as a group-level process. Group-level memory is the focus of research on transactive memory (Hollingshead *et al.*, 2011; Wegner, 1987; Wegner *et al.*, 1991), which refers to the sharing across group members of responsibility for the different stages and aspects of the memory process. Broadly speaking, a transactive memory system (TMS) consists of two components (Theiner, 2013). Its representational component includes both the first-order (declarative and procedural) memories of its members and their meta-level knowledge of each other's memories. Its procedural component includes the various (implicit and explicit) communication processes through which group members assign responsibility for and coordinate performance of the memory process. TMSs often perform better than individuals do on their own, at least with respect to tasks that lend themselves to a division of cognitive labor (e.g., Liang *et al.*, 1995; see Ren & Argote, 2011 for a review). Theiner (2013) points out that the members of a TMS have nonoverlapping memories, so they are not interchangeable: that, if enough members of a TMS are removed, the system will fail, again due to members' nonoverlapping knowledge; that a TMS's history makes a difference to its performance, so disassembly and reassembly may affect its ability to remember; and that, since a TMS is welded together in part by its members' knowledge of what each member of the system is responsible for knowing, cooperative and inhibitory actions among its members are critical to its function. Appealing to Wimsatt's (1986) notion of emergence, Theiner concludes that TMSs constitute a form of emergent group-level memory, in the sense that the group has a memory capacity of its own, over and above those of its individual members.

Most empirical research in this area focuses on the quantitative effects of transactive memory, but TMSs may manifest qualitative forms of emergence as well. Focusing on married couples, who are particularly likely to form transactive memory systems, Harris *et al.* (2014; cf. Sutton *et al.*, 2010) point to several additional forms of emergence: emergence of new information (information that neither individual is capable of remembering on his own); emergence of increased "quality" (greater emotional richness and episodic detail); and emergence of new forms of understanding (e.g., recontextualizing the significance of an event) (Barnier *et al.*, 2014; Harris *et al.*, 2011). Each of these cases offers an example of the form of organization-dependence highlighted by Theiner in his argument: the married couples in question perform as integrated systems, as opposed to mere aggregates of individuals. Thus there is a strong case to be made for viewing TMSs as remembering entities in their own right, bringing us to the other end of the spectrum between mere social manifestation and genuine group minds.

Before turning to the epistemology of emergent group-level memory, we briefly consider ways in which TMSs might incorporate technological resources, thus supporting the sociotechnical system mind thesis. It is useful here to bear in mind the distinction between a TMS's first-order memory and its meta-level knowledge. As far as first-order knowledge is concerned, we have already seen that technological resources often serve as external memory stores and are subject to much the same sort of metacognitive monitoring as both other subjects and internal memory; hence there would seem to be no barrier to incorporating such resources into TMSs. Indeed, from an extracranialist perspective, this is to be expected, and the focus of most TMS research on purely social systems, as opposed to mixed sociotechnical systems, is likely misplaced. The possibility of hybrid sociotechnical TMSs receives further support from one of the few studies investigating such systems. Treating families coping with a member's memory impairment as distributed cognitive systems, Wu *et al.* (2008) found that such groups tend to rely heavily on technological supports (ranging from calendars to personal digital assistants), with nonimpaired family

members serving not only as sources of first-order information but also carrying out metacognitive functions by developing and maintaining the technological supports.

As far as such metamemory functions are concerned, it is useful to distinguish between small groups, such as families, and larger groups, such as business organizations. In small groups, it is possible for individuals to use their onboard metamemory capacities to keep track of which component—whether human or technological—is responsible for knowing what. Interestingly, this opens up the possibility of a division of cognitive labor at the metamemory level as well as the first-order level, with different members of the TMS having different responsibilities with respect to the technological components of the TMS. For example, Harris *et al.* (unpublished data) have found that, in married couples, one spouse often has primary responsibility for the maintenance and use of external memory stores. Such divisions of metacognitive labor are likely to be even more vital to the functioning of large-scale TMSs, since transactive memory does not scale up well, due to increased demands on members' metamemory capacities with increased group size (Moreland, 2006; Ren & Argote, 2011; Theiner, 2013). One proposed means of enabling transactive memory in larger groups involves a division of metacognitive labor between human and technological components, allowing the technological components to take over crucial metamemory functions, such as keeping track of which component is responsible for knowing a given piece of information (Nevo & Wand, 2005; Theiner, 2013).

The general moral is that the role of artifacts should not be left out of our accounts of collaborative remembering. Just as we miss many of the patterns and mechanisms responsible for the successes and failures of human remembering if we treat it as a purely intracranial process, ignoring the vital roles played by other people, we miss many of the important patterns and mechanisms if we treat remembering as a purely social process, ignoring the vital roles played by technological artifacts (see also Hoskins; van den Hoven *et al.*, Chapters 21 and 22).

Distributed Reliabilism

In the previous section, we suggested that the workings of extended memory systems, in which the subject's agency often plays a secondary role, support moving from virtue reliabilism to a form of socially and technologically extended process reliabilism. In this section, we suggest that the workings of distributed memory systems support moving further, to a form of socially and technologically *distributed* process reliabilism.

We have argued that, while some of the memory phenomena investigated under the heading of collaborative memory, such as SS-RIF, support the relatively weak social and technological manifestation thesis, others, such as transactive memory, support the much stronger sociotechnical system mind thesis. A similarly nonreductionist view of transactive memory was voiced by Wegner (1987), who viewed a TMS as “a knowledge-acquiring, knowledge-holding, and knowledge-using system that is greater than the sum of its individual member systems.” More recently, Barnier *et al.* (2008) have argued along the same lines that the transactive memory framework “captures a truly shared memory system,” in the sense that it “predicts that memories recalled by the individuals in the transactive systems will be more than the sum of individual memory.” Transactive memory thus represents a potential step toward distributed process reliabilism.

Attributions not only of memory but of any cognitive capacity to social groups or sociotechnical systems have, of course, been subject to numerous objections (see Wilson, 2005). While we cannot deal with all such objections here, we can respond to one that is particularly relevant to the question of collaborative memory knowledge. Intuitively, treating distributed memory systems as remembering entities presupposes treating them as cognitive agents, but there are real concerns about the legitimacy of attributions of agency to groups or to sociotechnical systems. Indeed, we

have conceded that such attributions are to be rejected. Hence, it seems that we must reject the claim that distributed memory systems are themselves remembering entities.

In principle, we might respond to this objection by arguing that being a remembering entity does not in fact presuppose being a cognitive agent. But a more promising move is to grant that being a remembering entity presupposes being a cognitive agent but to argue that the sense of cognitive agency at issue here does not pose a problem for our account. Wilson (2005) usefully distinguishes between two related notions of agency. *Functional* agency is the minimal notion: it requires only that the system has control over what lies within its boundaries and autonomy with respect to what lies without its boundaries. *Cognitive* agency is a somewhat more demanding but still fairly minimal notion: it requires functional agency, plus cognitive capacities. TMSs have the right sort of internal control and external autonomy to qualify as functional agents. Since they inherit the cognitive capacities of their constituent members, moreover, they would appear to qualify as cognitive agents.

Why, then, have we granted that such systems should not be counted as agents? We granted this point in response to a form of virtue reliabilism that would have required us to assign credit or responsibility for cognitive success and failure to extended memory systems. Such assignments presuppose a richer notion of agency: what we might refer to as *responsible* cognitive agency, where responsible cognitive agency requires cognitive agency, plus responsibility. There is no clear sense in which a TMS, for example, might be assigned responsibility for its cognitive success and failures, and it is in this sense that we have suggested that extended and distributed memory systems do not qualify as cognitive agents: they may be cognitive agents simpliciter, but they are not responsible cognitive agents.⁵ Since process reliabilism does not require assignments of credit or responsibility, cognitive agency simpliciter is sufficient for *socially and technologically distributed* process reliabilism, which simply extends process reliabilism by allowing distributed cognitive systems to count as knowing subjects. Our position, then, is that some extended and distributed memory systems qualify as cognitive agents in a sense rich enough to support distributed reliabilism, but weak enough to be theoretically unproblematic.

In the broader debate over group knowledge, socially and technologically distributed process reliabilism aligns with nonsummativism in the debate between summativists and nonsummativists (Quinton, 1975). Summativists accept that groups can have beliefs only in a derivative sense; for example, a group *G* might be said to believe that *P* when all or most of the members of *G* believe that *P*. Nonsummativists maintain that groups can have beliefs that diverge from those of their members; *G* might in principle believe that *P*, even if most of the members of *G* do not believe that *P*.

Nonsummativism can be motivated by noting that certain voting procedures, for example, can result in a group adopting a position that is held by few or none of its members. As collaborative memory research shows, a divergence between the beliefs of a group and those of its members can also arise more organically in the process of remembering together, with the emergence of a shared memory that deviates from the memories held by individual group members, for example through the recombination and recontextualization of individual memories.

If distributed cognitive systems are capable of (nonsummatively) holding beliefs, the reliability of the distributed processes responsible for generating those beliefs can in principle be determined. Some might deny that an evaluation of the reliability of a distributed cognitive process

⁵ There may be room for views on which groups can be assigned responsibility for their cognitive successes and failures. However, while a considerable amount of work has been done on the concept of collective moral responsibility, to our knowledge little work has been done on the concept of collective epistemic responsibility.

suffices to ground an evaluation of the epistemic status of the resulting belief, on the grounds that distributed cognitive systems are not appropriate subjects of epistemic evaluation. Giere (2004), for example, endorses distributed *cognition* but rejects distributed *knowledge*. He rejects distributed knowledge, however, because he rejects distributed *agency*, and, as we have seen, a notion of agency weak enough to allow us to count distributed cognitive systems as agents is available. If distributed cognitive systems can be viewed as cognitive agents, at least in a minimal sense, it is meaningful to view them as knowers, again perhaps in a minimal sense.

To treat distributed cognitive systems as even minimal knowers is already to go further than most epistemologists have been willing to go, but it is not unprecedented. One view which is similar in spirit to socially and technologically distributed process reliabilism is Thagard's (1997) treatment of "collaborative knowledge." Thagard's focus is on the benefits for reliability that can be derived from collaboration in scientific inquiry. While we have touched on the costs and benefits of collaboration in remembering, our focus is different: socially and technologically distributed process reliabilism is concerned not with the effects on reliability of collaboration in remembering but rather with the nature of collaborative memory knowledge itself, which we analyze in terms of reliability. Thus, while similar in spirit, the two views are ultimately concerned with different questions.

A view which is closer to socially and technologically distributed process reliabilism is the "social process reliabilism" defended recently by Goldman (2014). Goldman distinguishes between "vertical" and "horizontal" dimensions of the justification of group beliefs. The vertical dimension sees the justificatory status of a group's beliefs in terms of their relationship to the beliefs of its members. The horizontal dimension sees the justificatory status of a group's beliefs in terms of the group-level process for producing those beliefs. Goldman's focus is largely on the vertical dimension (e.g., on the effects of different belief-aggregation rules). In contrast, since we treat the group itself as a remembering subject, our focus is on the horizontal dimension. With respect to the horizontal dimension, Goldman suggests that the process reliabilist criterion for justification is applicable at the level of the group. Hence, the key difference between our view and his is that our view is meant to cover not only systems consisting of human subjects but also hybrid systems consisting of human subjects and technological artifacts.

Concluding Remarks

While socially and technologically distributed process reliabilism goes further than most epistemologists have been willing to go, it does not—at least in the version defended here—go as far in the direction of socially and technologically distributed knowledge as one might in principle go. Our focus here has been on remembering and memory knowledge in relatively small-scale sociotechnical systems. In principle, we might attempt to extend our argument to very large-scale systems as well—moving from the sort of *collaborative* memory exemplified by TMSs to the sort of *collective* memory that operates at the level of societies and nations. This sort of move would not be unprecedented. Anastasio *et al.* (2012), for example, have argued that the very same processes of memory consolidation occur both at the level of individuals and at the level of whole societies. Nevertheless, there are good reasons to suppose that the notion of collective remembering, and hence the notion of collective memory knowledge, should not be taken literally. Regardless of whether the parallels that Anastasio *et al.* (2012) see between individual and collective consolidation in fact obtain (and there is reason to doubt that they do; Michaelian, 2014b), societies are unlikely to display sufficient control and autonomy to count as cognitive agents in even a minimal sense. Distributed reliabilism is thus most plausibly seen as applying to collaborative but not to collective memory knowledge.

Setting collective memory aside, the take-home points are the following. Epistemology risks going astray when it tries to analyze knowledge from the armchair. This risk is clearly illustrated by the extended knowledge debate, which has tended not to take relevant data from cognitive science and psychology into account and has thus produced implausible accounts of extended and distributed knowledge. Our method in this chapter was to attempt to counteract this tendency by taking relevant data into account in order to produce a more plausible account of extended and distributed knowledge, focusing on collaborative memory knowledge in particular. Extracranialist approaches to cognition have provided convincing reasons to hold that remembering, rather than unfolding entirely within the brain, crosses the boundary between the individual and external social and technological resources, and in some cases—such as transactive memory—is best viewed as a process occurring at the level of the distributed sociotechnical system as a whole. Virtue reliabilism, including extended and distributed forms of the theory, appears to be unable to account for the knowledge produced by extended and distributed remembering. In contrast, distributed reliabilism is better able to accommodate extended and distributed memory knowledge: setting aside the considerations of agency on which virtue reliabilism focuses, the distributed reliabilist argues that the beliefs produced by extended and distributed remembering are justified—and count as knowledge—simply because (and to the extent that) extended and distributed remembering are reliable. Distributed reliabilism thus provides an empirically plausible account of collaborative memory knowledge.

Acknowledgments

The authors wish to thank audiences at a 2015 seminar at the University of Otago, the 2015 Memory and Knowledge conference at the Université Pierre-Mendès-France, the 2015 Australasian Association of Philosophy meeting at Macquarie University, and the 2016 Colombian Conference of Logic, Epistemology, and Philosophy of Science. Thanks also to the editors for their comments on a rough draft. Supported by grant 16-UOO-016 to KM from the Marsden Fund, administered by the Royal Society of New Zealand.

References

- Adams, F., & Aizawa, K. (2008). *The bounds of cognition*. Oxford, UK: Wiley-Blackwell.
- Anastasio, T. J., Ehrenberger, K. A., Watson, P., & Zhang, W. (2012). *Individual and collective memory consolidation: Analogous Processes on different levels*. Cambridge, MA: MIT Press.
- Arango-Muñoz, S. (2013). Scaffolded memory and metacognitive feelings. *Review of Philosophy and Psychology*, 4, 135–52. doi:10.1007/s13164-012-0124-1
- Arango-Muñoz, S., & Michaelian, K. (2014). Epistemic feelings, epistemic emotions: Review and introduction to the focus section. *Philosophical Inquiries*, 2, 97–122.
- Barnier, A., Sutton, J., Harris, C., & Wilson, R. (2008). A conceptual and empirical framework for the social distribution of cognition: The case of memory. *Cognitive Systems Research*, 9, 33–51. doi:10.1016/j.cogsys.2007.07.002
- Barnier, A. J., Priddis, A. C., Broekhuijse, J. M., Harris, C. B., Cox, R. E., Addis, D. R., . . . Congleton, A. R. (2014). Reaping what they sow: Benefits of remembering together in intimate couples. *Journal of Applied Research in Memory and Cognition*, 3, 261–5. doi:10.1016/j.jarmac.2014.06.003
- Barnier, A. J., & Sutton, J. (2008). From individual to collective memory: Theoretical and empirical perspectives. *Memory*, 16, 177–82. doi:10.1080/09541440701828274
- Basden, B. H., Basden, D. R., Bryner, S., & Thomas, R. L. (1997). A comparison of group and individual remembering: does collaboration disrupt retrieval strategies? *Journal*

- of *Experimental Psychology: Learning, Memory, and Cognition*, **23**, 1176–91. doi:10.1037/0278-7393.23.5.1176
- Betts, K. R., & Hinsz, V. B. (2010). Collaborative group memory: Processes, performance, and techniques for improvement. *Social and Personality Psychology Compass*, **4**, 119–30. doi:10.1111/j.1751-9004.2009.00252.x
- Carter, A., Clark, A., Kallestrup, J., Palermos, O., & Pritchard, D. (Eds.) (forthcoming). *Extended epistemology*. Oxford, UK: Oxford University Press.
- Carter, A., Kallestrup, J., Palermos, O., & Pritchard, D. (Eds.) (2014). Extended knowledge. *Philosophical Issues*, **24**, 1–23. doi:10.1111/phils.12023
- Clark, A. (2008). *Supersizing the mind: Embodiment, action, and cognitive extension*. New York, NY: Oxford University Press.
- Clark, A. (2015). What “extended me” knows. *Synthese*, **192**, 3757–75. doi:10.1007/s11229-015-0719-z
- Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, **58**, 7–19. doi:10.1093/analys/58.1.7
- Cuc, A., Koppel, J., & Hirst, W. (2007). Silence is not golden: A case for socially shared retrieval-induced forgetting. *Psychological Science*, **18**, 727–33. doi:10.1111/j.1467-9280.2007.01967.x
- Cuc, A., Ozuru, Y., Manier, D., & Hirst, W. (2006). On the formation of collective memories: The role of a dominant narrator. *Memory & Cognition*, **34**, 752–62. doi:10.3758/bf03193423
- Fagin, M. M., Yamashiro, J. K., & Hirst, W. C. (2013). The adaptive function of distributed remembering: Contributions to the formation of collective memory. *Review of Philosophy and Psychology*, **4**, 91–106. doi:10.1007/s13164-012-0127-y
- Gangemi, A., Bourgeois-Gironde, A., & Mancini, F. (2015). Feelings of error in reasoning—in search of a phenomenon. *Thinking & Reasoning*, **21**, 383–6. doi:10.1080/13546783.2014.980755
- Giere, R. (2004). The problem of agency in scientific distributed cognitive systems. *Journal of Cognition and Culture*, **4**, 759–74. doi:10.1163/1568537042484887
- Goldberg, S. C. (2010). *Relying on others: An essay in epistemology*. New York, NY: Oxford University Press.
- Goldman, A. I. (2012). *Reliabilism and contemporary epistemology: Essays*. New York, NY: Oxford University Press.
- Goldman, A. I. (2014). Social process reliabilism. In: J. Lackey (Ed.), *Essays in collective epistemology* (pp. 11–41). Oxford, UK: Oxford University Press.
- Green, A. (2012). Extending the credit theory of knowledge. *Philosophical Explorations*, **15**, 121–32. doi:10.1080/13869795.2012.670720
- Harris, C. B., Barnier, A. J., Sutton, J., & Keil, P. G. (2014). Couples as socially distributed cognitive systems: Remembering in everyday social and material contexts. *Memory Studies*, **7**, 285–97. doi:10.1177/1750698014530619
- Harris, C. B., Keil, P. G., Sutton, J., Barnier, A. J., & McIlwain, D. J. F. (2011). We remember, we forget: Collaborative remembering in older couples. *Discourse Processes*, **48**, 267–303. doi:10.1080/0163853x.2010.541854
- Hirst, W., & Echterhoff, G. (2008). Creating shared memories in conversation: Toward a psychology of collective memory. *Social Research*, **75**, 183–216.
- Hollingshead, A. B., Yoon, N. G. K., & Brandon, D. P. (2011). Transactive memory theory and teams: Past, present, and future. In: E. Salas, S. M. Fiore, & M. Letsky (Eds.), *Theories of team cognition: Cross-disciplinary perspectives* (pp. 421–55). Oxford, UK: Taylor and Francis.
- Hutchins, E. (1995). How a cockpit remembers its speeds. *Cognitive Science*, **19**, 265–88. doi:10.1207/s15516709cog1903_1
- Hutchins, E. (1996). *Cognition in the wild*. Cambridge, MA: MIT Press.
- Kalnikaite, V., & Whittaker, S. (2007). Software or wetware?: Discovering when and why people use digital prosthetic memory. In: *Proceedings of the SIGCHI conference on human factors in computing systems, CHI '07* (pp. 71–80). New York, NY: ACM.

- Kirchhoff, M. D., & Newsome, W. (2012). Distributed cognitive agency in virtue epistemology. *Philosophical Explorations*, 15, 165–80. doi:10.1080/13869795.2012.670722
- Larson, J. R. (2009). *In search of synergy in small group performance*. New York, NY: Psychology Press.
- Liang, D. W., Moreland, R., & Argote, L. (1995). Group versus individual training and group performance: The mediating role of transactive memory. *Personality and Social Psychology Bulletin*, 21, 384–93. doi:10.1177/0146167295214009
- Meade, M. L., Nokes, T. J., & Morrow, D. G. (2009). Expertise promotes facilitation on a collaborative memory task. *Memory*, 17, 39–48. doi:10.1080/09658210802524240
- Michaelian, K. (2010). In defence of gullibility: The epistemology of testimony and the psychology of deception detection. *Synthese*, 176, 399–427. doi:10.1007/s11229-009-9573-1
- Michaelian, K. (2012). Metacognition and endorsement. *Mind & Language*, 27, 284–307. doi:10.1111/j.1468-0017.2012.01445.x
- Michaelian, K. (2014a). JFGI: From distributed cognition to distributed reliabilism. *Philosophical Issues*, 24, 314–46. doi:10.1111/phils.12036
- Michaelian, K. (2014b). Review essay on Anastasio *et al.*, *Individual and Collective Memory Consolidation: Analogous Processes on Different Levels*. *Memory Studies*, 7, 254–64. doi:10.1177/1750698013515365
- Michaelian, K., & Sutton, J. (2013). Distributed cognition and memory research: History and current directions. *Review of Philosophy and Psychology*, 4, 1–24. doi:10.1007/s13164-013-0131-x
- Moreland, R. L. (2006). Transactive memory: Learning who knows what in work groups and organizations. In: J. M. Levine, & R. L. Moreland (Eds.), *Small groups* (pp. 327–46). New York, NY: Psychology Press.
- Nelson, T. O., & Narens, L. (1990). Metamemory: A theoretical framework and new findings. In: G. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory*, Vol. 26 (pp. 125–73). New York, NY: Academic Press.
- Nevo, D., & Wand, Y. (2005). Organizational memory information systems: A transactive memory approach. *Decision Support Systems*, 39, 549–62. doi:10.1016/j.dss.2004.03.002
- Palermos, S. O. (2015). Active externalism, virtue reliabilism and scientific knowledge. *Synthese*, 192, 1955–2986. doi:10.1007/s11229-015-0695-3
- Pavitt, C. (2003). Colloquy: Do interacting groups perform better than aggregates of individuals? *Human Communication Research*, 29, 592–9. doi:10.1093/hcr/29.4.592
- Pritchard, D. (2010). Cognitive ability and the extended cognition thesis. *Synthese*, 175, 133–51. doi:10.1007/s11229-010-9738-y
- Quinton, A. (1975). Social objects. *Proceedings of the Aristotelian Society*, 76, 1–28. doi:10.1093/aristotelian/76.1.1
- Rajaram, S., & Pereira-Pasarin, L. P. (2010). Collaborative memory: Cognitive research and theory. *Perspectives on Psychological Science*, 5, 649–63. doi:10.1177/1745691610388763
- Reber, R., & Unkelbach, C. (2010). The epistemic status of processing fluency as source for judgments of truth. *Review of Philosophy and Psychology*, 1, 563–81.
- Reder, L. M. (1987). Strategy selection in question answering. *Cognitive Psychology*, 19, 90–138. doi:10.1007/s13164-010-0039-7
- Ren, Y., & Argote, L. (2011). Transactive memory systems 1985-2010: An integrative framework of key dimensions, antecedents, and consequences. *The Academy of Management Annals*, 5, 189–229. doi:10.1080/19416520.2011.590300
- Roediger, H. L. III, Meade, M. L., & Bergman, E. T. (2001). Social contagion of memory. *Psychonomic Bulletin & Review*, 8, 365–71. doi:10.3758/bf03196174
- Sparrow, B., & Chatman, L. (2013). Social cognition in the internet age: Same as it ever was? *Psychological Inquiry*, 24, 273–92. doi:10.1080/1047840x.2013.827079
- Sparrow, B., Liu, J., & Wegner, D. M. (2011). Google effects on memory: Cognitive consequences of having information at our fingertips. *Science*, 333, 776–8. doi:10.1126/science.1207745

- Stone, C. B., Barnier, A. J., Sutton, J., & Hirst, W. (2013). Forgetting our personal past: Socially shared retrieval-induced forgetting of autobiographical memories. *Journal of Experimental Psychology: General*, *142*, 1084–99. doi:10.1037/a0030739
- Stone, C. B., Coman, A., Brown, A. D., Koppel, J., & Hirst, W. (2012). Toward a science of silence: The consequences of leaving a memory unsaid. *Perspectives on Psychological Science*, *7*, 39–53. doi:10.1177/1745691611427303
- Sutton, J., Harris, C. B., Keil, P. G., & Barnier, A. J. (2010). The psychology of memory, extended cognition, and socially distributed remembering. *Phenomenology and the Cognitive Sciences*, *9*, 521–60. doi:10.1007/s11097-010-9182-y
- Sutton, J., McIlwain, D., Christensen, W., & Geeves, A. (2011). Applying intelligence to the reflexes: Embodied skills and habits between Dreyfus and Descartes. *Journal of the British Society for Phenomenology*, *42*, 78–103. doi:10.1080/00071773.2011.11006732
- Thagard, P. (1997). Collaborative knowledge. *Noûs*, *31*, 242–61. doi:10.1111/0029-4624.00044
- Theiner, G. (2013). Transactive memory systems: A mechanistic analysis of emergent group memory. *Review of Philosophy and Psychology*, *4*, 65–89. doi:10.1007/s13164-012-0128-x
- Theiner, G., Allen, C., & Goldstone, R. L. (2010). Recognizing group cognition. *Cognitive Systems Research*, *11*, 378–95. doi:10.1016/j.cogsys.2010.07.002
- Thompson, R. (2008). Collaborative and social remembering. In: G. Cohen, & M. A. Conway (Eds.), *Memory in the real world* (pp. 249–67). London, UK: Psychology Press.
- Vrij, A. (2008). *Detecting lies and deceit: Pitfalls and opportunities* (2nd ed.). West Sussex, UK: John Wiley & Sons Ltd.
- Wegner, D. M. (1987). Transactive memory: A contemporary analysis of the group mind. In B. Mullen, & G. R. Goethals (Eds.), *Theories of group behavior* (pp. 185–208). New York, NY: Springer.
- Wegner, D. M., Erber, R., & Raymond, P. (1991). Transactive memory in close relationships. *Journal of Personality and Social Psychology*, *61*, 923–9. doi:10.1037/0022-3514.61.6.923
- Weldon, M. S. (2000). Remembering as a social process. *Psychology of Learning and Motivation*, *40*, 67–120. doi:10.1016/s0079-7421(00)80018-3
- Weldon, M. S., & Bellinger, K. D. (1997). Collective memory: Collaborative and individual processes in remembering. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *23*, 1160–75. doi:10.1037/0278-7393.23.5.1160
- Wilson, R. A. (2005). Collective memory, group minds, and the extended mind thesis. *Cognitive Processing*, *6*, 227–36. doi:10.1007/s10339-005-0012-z
- Wimsatt, W. C. (1986). Forms of aggregativity. In: A. Donagan, A. N. Perovich, & M. V. Wedin (Eds.), *Human nature and natural knowledge, Boston studies in the philosophy of science* (pp. 259–291). Dordrecht, Netherlands: Springer.
- Wu, M., Birnholtz, J., Richards, B., Baecker, R., & Massimi, M. (2008). Collaborating to remember: A distributed cognition account of families coping with memory impairments. In: *Proceedings of the SIGCHI conference on human factors in computing systems*, CHI '08 (pp. 825–34). New York, NY: ACM.