we recommend viewing this manuscript in html format at https://egouldo.github.io/ManyAnalysts/

Same data, different analysts: variation in effect sizes due to analytical decisions in ecology and evolutionary biology.

Elliot Gould, School of Agriculture Food and Ecosystem Sciences, University of Melbourne, Australia

Hannah S. Fraser, School of Historical and Philosophical Studies, University of Melbourne, Australia

Timothy H. Parker, Department of Biology, Whitman College, USA. Author for Correspondence: parkerth@whitman.edu

Shinichi Nakagawa, School of Biological, Earth & Environmental Sciences, University of New South Wales, Australia

Simon C. Griffith, School of Natural Sciences, Macquarie University, Australia

Peter A. Vesk, School of Agriculture Food and Ecosystem Sciences, University of Melbourne, Australia

Fiona Fidler, School of Historical and Philosophical Studies, University of Melbourne, Australia

Daniel G. Hamilton, School of Public Health and Preventive Medicine, Monash University, Australia

Robin N Abbey-Lee, Länsstyrelsen Östergötland, Sweden

Jessica K. Abbott, Biology Department, Lund University, Sweden

Luis A. Aguirre, Department of Biology, University of Massachusetts, USA

Carles Alcaraz, Marine and Continental Waters, IRTA, Spain

Irith Aloni, Deptartment of Life Sciences, Ben Gurion University of the Negev, Israel

Drew Altschul, Department of Psychology, The University of Edinburgh, UK

Kunal Arekar, Centre for Ecological Sciences, Indian Institute of Science, India

Jeff W. Atkins, Southern Research Station, USDA Forest Service, USA

Joe Atkinson, Center for Ecological Dynamics in a Novel Biosphere (ECONOVO), Department of Biology, Aarhus University, Denmark

Christopher M. Baker, School of Mathematics and Statistics, University of Melbourne, Australia

Meghan Barrett, Biology, Indiana University Purdue University Indianapolis, USA

Kristian Bell, School of Life and Environmental Sciences, Deakin University, Australia

Suleiman Kehinde Bello, Department of Arid Land Agriculture, King Abdulaziz University, Kingdom of Saudi Arabia

Iván Beltrán, Department of Biological Sciences, Macquarie University, Australia

Bernd J. Berauer, Department of Plant Ecology, University of Hohenheim, Institute of Landscape and Plant Ecology, Germany

Michael Grant Bertram, Department of Wildlife, Fish, and Environmental Studies, Swedish University of Agricultural Sciences, Sweden

we recommend viewing this manuscript in html format at https://egouldo.github.io/ManyAnalysts/

Peter D. Billman, Department of Ecology and Evolutionary Biology, University of Connecticut, USA

Charlie K. Blake, STEM Center, Southern Illinois University Edwardsville, USA

Shannon Blake, University of Guelph, Canada

Louis Bliard, Department of Evolutionary Biology and Environmental Studies, University of Zurich, Switzerland

Andrea Bonisoli-Alquati, Department of Biological Sciences, California State Polytechnic University, Pomona, USA

Timothée Bonnet, Centre d'Études Biologiques de Chizé, UMR 7372 Université de la Rochelle - Centre National de la Recherche Scientifique, France

Camille Nina Marion Bordes, Faculty of Life Sciences, Bar Ilan University, Israel

Aneesh P. H. Bose, Department of Wildlife, Fish, and Environmental Studies, Swedish University of Agricultural Sciences, Sweden

Thomas Botterill-James, School of Natural Sciences, University of Tasmania, Australia

Melissa Anna Boyd, Whitebark Institute, USA

Sarah A. Boyle, Department of Biology, Rhodes College, USA

Tom Bradfer-Lawrence, Centre for Conservation Science, RSPB, UK

Jennifer Bradham, Environmental Studies, Wofford College, USA

Jack A. Brand, Department of Wildlife, Fish and Environmental Studies, Swedish University of Agricultural Sciences, Sweden

Martin I. Brengdahl, IFM Biology, Linköping University, Sweden

Martin Bulla, Faculty of Environmental Sciences, Czech University of Life Sciences Prague, Czech Republic

Luc Bussière, Biological and Environmental Sciences & Gothenburg Global Biodiversity Centre, University of Gothenburg, Sweden

Ettore Camerlenghi, School of Biological Sciences, Monash University, Australia

Sara E. Campbell, Ecology and Evolutionary Biology, University of Tennessee Knoxville, USA

Leonardo L. F. Campos, Departamento de Ecologia e Zoologia, Universidade Federal de Santa Catarina, Brazil

Anthony Caravaggi, School of Biological and Forensic Sciences, University of South Wales, UK

Pedro Cardoso, Centre for Ecology, Evolution and Environmental Changes (cE3c) & CHANGE - Global Change and Sustainability Institute, Faculdade de Ciências, Universidade de Lisboa, Portugal

Charles J.W. Carroll, Forest and Rangeland Stewardship, Colorado State University, USA

Therese A. Catanach, Department of Ornithology, Academy of Natural Sciences of Drexel University, USA

we recommend viewing this manuscript in html format at https://egouldo.github.io/ManyAnalysts/

Xuan Chen, Biology, Salisbury University, USA

Heung Ying Janet Chik, Groningen Institute for Evolutionary Life Sciences, University of Groningen, Netherlands

Emily Sarah Choy, Department of Biology, McMaster University, Canada

Alec Philip Christie, Department of Zoology, University of Cambridge, UK

Angela Chuang, Entomology and Nematology, University of Florida, USA

Amanda J. Chunco, Environmental Studies, Elon University, USA

Bethany L. Clark, BirdLife International, UK

Andrea Contina, School of Integrative Biological and Chemical Sciences, The University of Texas Rio Grande Valley, USA

Garth A. Covernton, Department of Ecology and Evolutionary Biology, University of Toronto, Canada

Murray P. Cox, Department of Statistics, University of Auckland, New Zealand

Kimberly A. Cressman, Catbird Stats, LLC, USA

Marco Crotti, School of Biodiversity, One Health & Veterinary Medicine, University of Glasgow, UK

Connor Davidson Crouch, School of Forestry, Northern Arizona University, USA

Pietro B. D'Amelio, Department of Behavioural Neurobiology, Max Planck Institute for Biological Intelligence, Germany

Alexandra Allison de Sousa, School of Sciences: Center for Health and Cognition, Bath Spa University, UK

Timm Fabian Döbert, Department of Biological Sciences, University of Alberta, Canada

Ralph Dobler, Applied Zoology, TU Dresden, Germany

Adam J. Dobson, School of Molecular Biosciences, College of Medical Veterinary & Life Sciences, University of Glasgow, UK

Tim S. Doherty, School of Life and Environmental Sciences, The University of Sydney, Australia

Szymon Marian Drobniak, Institute of Environmental Sciences, Jagiellonian University, Poland

Alexandra Grace Duffy, Biology Department, Brigham Young University, USA

Alison B. Duncan, Institute of Evolutionary Sciences Montpellier, University of Montpellier, CNRS, IRD., France

Robert P. Dunn, Baruch Marine Field Laboratory, University of South Carolina, USA

Jamie Dunning, Department of Life Sciences, Imperial College London, UK

Trishna Dutta, European Forest Institute, Germany

Luke Eberhart-Hertel, Department of Ornithology, Max Planck Institute for Biological Intelligence, Germany

Jared Alan Elmore, Forestry and Environmental Conservation, National Bobwhite and Grassland Initiative, Clemson University, USA

Mahmoud Medhat Elsherif, Department of Psychology and Vision Science, University of Birmingham, Baily Thomas Grant, UK

Holly M. English, School of Biology and Environmental Science, University College Dublin, Ireland

David C. Ensminger, Department of Biological Sciences, San José State University, USA

Ulrich Rainer Ernst, Apicultural State Institute, University of Hohenheim, Germany

Stephen M. Ferguson, Department of Biology, St. Norbert College, USA

Esteban Fernandez-Juricic, Department of Biological Sciences, Purdue University, USA

Thalita Ferreira-Arruda, Biodiversity, Macroecology & Biogeography, Faculty of Forest Sciences and Forest Ecology, University of Göttingen, Germany

John Fieberg, Department of Fisheries, Wildlife, and Conservation Biology, University of Minnesota, USA

Elizabeth A. Finch, CABI, UK

Evan A. Fiorenza, Department of Ecology and Evolutionary Biology, School of Biological Sciences, University of California, Irvine, USA

David N. Fisher, School of Biological Sciences, University of Aberdeen, UK

Amélie Fontaine, Department of Natural Resource Sciences, McGill University, Canada

Wolfgang Forstmeier, Department of Ornithology, Max Planck Institute for Biological Intelligence, Germany

Yoan Fourcade, Institute of Ecology and Environmental Sciences (iEES), Univ. Paris-Est Creteil, France

Graham S. Frank, Department of Forest Ecosystems and Society, Oregon State University, USA

Cathryn A. Freund, Wake Forest University, USA

Eduardo Fuentes-Lillo, Laboratorio de Invasiones Biológicas (LIB), Instituto de Ecología y Biodiversidad, Chile

Sara L. Gandy, Institute for Biodiversity, Animal Health and Comparative Medicine, University of Glasgow, UK

Dustin G. Gannon, Department of Forest Ecosystems and Society, College of Forestry, Oregon State University, USA

Ana I. García-Cervigón, Biodiversity and Conservation Area, Rey Juan Carlos University, Spain

Alexis C. Garretson, Graduate School of Biomedical Sciences, Tufts University, USA

Xuezhen Ge, Department of Integrative Biology, University of Guelph, Canada

William L. Geary, School of Life and Environmental Sciences (Burwood Campus), Deakin University, Australia

Charly Géron, CNRS, University of Rennes, France

Marc Gilles, Department of Behavioural Ecology, Bielefeld University, Germany

Antje Girndt, Fakultät für Biologie, Arbeitsgruppe Evolutionsbiologie, Universität Bielefeld, Germany

Daniel Gliksman, Chair of Meteorology, Institute for Hydrology and Meteorology, Faculty of Environmental Sciences, Technische Universität Dresden, Germany

Harrison B. Goldspiel, Department of Wildlife, Fisheries, and Conservation Biology, University of Maine, USA

Dylan G. E. Gomes, Department of Biological Sciences, Boise State University, USA

Megan Kate Good, School of Agriculture, Food and Ecosystem Sciences, The University of Melbourne, Australia

Sarah C. Goslee, Pastures Systems and Watershed Management Research Unit, USDA Agricultural Research Service, USA

J. Stephen Gosnell, Department of Natural Sciences, Baruch College, City University of New York, USA

Eliza M. Grames, Department of Biological Sciences, Binghamton University, USA

Paolo Gratton, Dipartimento di Biologia, Università di Roma "Tor Vergata", Italy

Nicholas M. Grebe, Department of Anthropology, University of Michigan, USA

Skye M. Greenler, College of Forestry, Oregon State University, USA

Maaike Griffioen, University of Antwerp, Belgium

Daniel M. Griffith, Earth & Environmental Sciences, Wesleyan University, USA

Frances J. Griffith, Yale School of Medicine, Department of Psychiatry, Yale University, USA

Jake J. Grossman, Biology Department and Environmental Studies Department, St. Olaf College, USA

Ali Güncan, Department of Plant Protection, Faculty of Agriculture, Ordu University, Turkey

Stef Haesen, Department of Earth and Environmental Sciences, KU Leuven, Belgium

James G. Hagan, Department of Marine Sciences, University of Gothenburg, Sweden

Heather A. Hager, Department of Biology, Wilfrid Laurier University, Canada

Jonathan Philo Harris, Natural Resource Ecology and Management, Iowa State University, USA

Natasha Dean Harrison, School of Biological Sciences, University of Western Australia, Australia

Sarah Syedia Hasnain, Department of Biological Sciences, Middle East Technical University, Turkey

Justin Chase Havird, Dept. of Integrative Biology, University of Texas at Austin, USA

Andrew J. Heaton, Grand Bay National Estuarine Research Reserve, USA

María Laura Herrera-Chaustre, Universidad de los Andes, Colombia

Tanner J. Howard

Bin-Yan Hsu, Department of Biology, University of Turku, Finland

Fabiola Iannarilli, Dept of Fisheries, Wildlife and Conservation Biology, University of Minnesota, USA

Esperanza C. Iranzo, Instituto de Ciencia Animal. Facultad de Ciencias Veterinarias, Universidad Austral de Chile, Chile

Erik N. K. Iverson, Department of Integrative Biology, The University of Texas at Austin, USA

Saheed Olaide Jimoh, Department of Botany, University of Wyoming, USA

Douglas H. Johnson, Department of Fisheries, Wildlife, and Conservation Biology, University of Minnesota, USA

Martin Johnsson, Department of Animal Breeding and Genetics, Swedish University of Agricultural Sciences, Sweden

Jesse Jorna, Department of Biology, Brigham Young University, Brigham Young University, USA

Tommaso Jucker, School of Biological Sciences, University of Bristol, UK

Martin Jung, International Institute for Applied Systems Analysis (IIASA), Austria

Ineta Kačergytė, Department of Ecology, Swedish University of Agricultural Sciences, Sweden

Oliver Kaltz, Université de Montpellier, France

Alison Ke, Department of Wildlife, Fish, and Conservation Biology, University of California, Davis, USA

Clint D. Kelly, Département des Sciences biologiques, Université du Québec à Montréal, Canada

Katharine Keogan, Institute of Evolutionary Biology, University of Edinburgh, UK

Friedrich Wolfgang Keppeler, Center for Limnology, Center for Limnology, University of Wisconsin - Madison, USA

Alexander K. Killion, Center for Biodiversity and Global Change, Yale University, USA

Dongmin Kim, Department of Ecology, Evolution, and Behavior, University of Minnesota, St. Paul, USA

David P. Kochan, Institute of Environment and Department of Biological Sciences, Florida International University, USA

Peter Korsten, Department of Life Sciences, Aberystwyth University, UK

Shan Kothari, Institut de recherche en biologie végétale, Université de Montréal, Canada

Jonas Kuppler, Institute of Evolutionary Ecology and Conservation Genomics, Ulm University, Germany

Jillian M. Kusch, Department of Biology, Memorial University of Newfoundland, Canada

Malgorzata Lagisz, Evolution & Ecology Research Centre and School of Biological, Earth & Environmental Sciences, University of New South Wales, Australia

Kristen Marianne Lalla, Department of Natural Resource Sciences, McGill University, Canada

Daniel J. Larkin, Department of Fisheries, Wildlife and Conservation Biology, University of Minnesota-Twin Cities, USA

Courtney L. Larson, The Nature Conservancy, USA

Katherine S. Lauck, Department of Wildlife, Fish, and Conservation Biology, University of California, Davis, USA

M. Elise Lauterbur, Ecology and Evolutionary Biology, University of Arizona, USA

Alan Law, Biological and Environmental Sciences, University of Stirling, UK

Don-Jean Léandri-Breton, Department of Natural Resource Sciences, McGill University, Canada

Jonas J. Lembrechts, Department of Biology, University of Antwerp, Belgium

Kiara L'Herpiniere, Natural sciences, Macquarie University, Australia

Eva J. P. Lievens, Aquatic Ecology and Evolution Group, Limnological Institute, University of Konstanz, Germany

Daniela Oliveira de Lima, Campus Cerro Largo, Universidade Federal da Fronteira Sul, Brazil

Shane Lindsay, School of Psychology and Social Work, University of Hull, UK

Martin Luquet, UMR 1224 ECOBIOP, Université de Pau et des Pays de l'Adour, France

Ross MacLeod, School of Biological & Environmental Sciences, Liverpool John Moores University, UK

Kirsty H. Macphie, Institute of Ecology and Evolution, University of Edinburgh, UK

Kit Magellan, Cambodia

Magdalena M. Mair, Statistical Ecotoxicology, Bayreuth Center of Ecology and Environmental Research (BayCEER), University of Bayreuth, Germany

Lisa E. Malm, Ecology and Environmental Science, Umeå University, Sweden

Stefano Mammola, Molecular Ecology Group (MEG), Water Research Institute (IRSA), National Research Council of Italy (CNR), Italy

Caitlin P. Mandeville, Department of Natural History, Norwegian University of Science and Technology, Norway

Michael Manhart, Center for Advanced Biotechnology and Medicine, Rutgers University Robert Wood Johnson Medical School, USA

Laura Milena Manrique-Garzon, Departamento de Ciencias Biológicas, Universidad de los Andes, Colombia

Elina Mäntylä, Department of Biology, University of Turku, Finland

Philippe Marchand, Institut de recherche sur les forêts, Université du Québec en Abitibi-Témiscamingue, Canada

Benjamin Michael Marshall, Biological and Environmental Sciences, University of Stirling, UK

Charles A. Martin, Université du Québec à Trois-Rivières, Canada

Dominic Andreas Martin, Institute of Plant Sciences, University of Bern, Switzerland

we recommend viewing this manuscript in html format at https://egouldo.github.io/ManyAnalysts/

Jake Mitchell Martin, Department of Wildlife, Fish, and Environmental Studies, Swedish University of Agricultural Sciences, Sweden

April Robin Martinig, School of Biological, Earth and Environmental Sciences, University of New South Wales, Australia

Erin S. McCallum, Department of Wildlife, Fish and Environmental Studies, Swedish University of Agricultural Sciences, Sweden

Mark McCauley, Whitney Laboratory for Marine Bioscience, University of Florida, USA

Sabrina M. McNew, Ecology and Evolutionary Biology, University of Arizona, USA

Scott J. Meiners, Biological Sciences, Eastern Illinois University, USA

Thomas Merkling, Centre d'Investigations Clinique Plurithématique - Institut Lorrain du Coeur et des Vaisseaux, Université de Lorraine, Inserm1433 CIC-P CHRU de Nancy, France

Marcus Michelangeli, Department of Wildlife, Fish and Environmental Studies, Swedish University of Agricultural Sciences, Sweden

Maria Moiron, Evolutionary biology department, Bielefeld University, Germany

Bruno Moreira, Department of Ecology and global change, Centro de Investigaciones sobre Desertificación, Consejo Superior de Investigaciones Cientificas (CIDE-CSIC/UV/GV), Spain

Jennifer Mortensen, Department of Biological Sciences, University of Arkansas, USA

Benjamin Mos, School of the Environment, Faculty of Science, The University of Queensland, Australia

Taofeek Olatunbosun Muraina, Department of Animal Health and Production, Oyo State College of Agriculture and Technology, Nigeria

Penelope Wrenn Murphy, Department of Forest & Wildlife Ecology, University of Wisconsin-Madison, USA

Luca Nelli, School of Biodiversity, One Health and Veterinary Medicine, University of Glasgow, UK

Petri Niemelä, Organismal and Evolutionary Biology Research Programme, Faculty of Biological and Environmental Sciences, University of Helsinki, Finland

Josh Nightingale, South Iceland Research Centre, University of Iceland, Iceland

Gustav Nilsonne, Department of Clinical Neuroscience, Karolinska Institutet, Sweden

Sergio Nolazco, School of Biological Sciences, Monash University, Australia

Sabine S. Nooten, Animal Ecology and Tropical Biology, University of Würzburg, Germany

Jessie Lanterman Novotny, Biology, Hiram College, USA

Agnes Birgitta Olin, Department of Aquatic Resources, Swedish University of Agricultural Sciences, Sweden

Chris L. Organ, Department of Earth Sciences, Montana State University, USA

Kate L. Ostevik, Department of Evolution, Ecology, and Organismal Biology, University of California, Riverside, USA

Facundo Xavier Palacio, Sección Ornitología, Universidad Nacional de La Plata, Argentina

Matthieu Paquet, Department of Ecology, Swedish University of Agricultural Sciences, Sweden

Darren James Parker, Bangor University, UK

David J. Pascall, MRC Biostatistics Unit, University of Cambridge, UK

Valerie J. Pasquarella, Harvard Forest, Harvard University, USA

John Harold Paterson, Biological and Environmental Sciences, University of Stirling, Scotland

Ana Payo-Payo, Departamento de Biodiversidad, Ecología y Evolución., Universidad Complutense de Madrid, Spain

Karen Marie Pedersen, Biology Department, Technische Universität Darmstadt, Germany

Grégoire Perez, UMR 1309 ASTRE, CIRAD, France

Kayla I. Perry, Department of Entomology, The Ohio State University, USA

Patrice Pottier, Evolution & Ecology Research Centre, School of Biological, Earth and Environmental Sciences, The University of New South Wales, Australia

Michael J. Proulx, Department of Psychology, University of Bath, UK

Raphaël Proulx, Chaire de recherche en intégrité écologique, Université du Québec à Trois-Rivières, Canada

Jessica L Pruett, Mississippi Based RESTORE Act Center of Excellence, University of Southern Mississippi, USA

Veronarindra Ramananjato, Department of Integrative Biology, University of California, Berkeley, USA

Finaritra Tolotra Randimbiarison, Mention Zoologie et Biodiversité Animale, Université d'Antananarivo, Madagascar

Onja H. Razafindratsima, Department of Integrative Biology, University of California, Berkeley, USA

Diana J. Rennison, Department of Ecology, Behavior and Evolution, University of California, San Diego, USA

Federico Riva, Institute for Environmental Sciences, VU Amsterdam, The Netherlands

Sepand Riyahi, Department of Evolutionary Anthropology, University of Vienna, Austria

Michael James Roast, Konrad Lorenz Institute for Ethology, University of Veterinary Medicine, Austria

Felipe Pereira Rocha, School of Biological Sciences, The University of Hong Kong, China

Dominique G. Roche, Institut de biologie, Université de Neuchâtel, Switzerland

Cristian Román-Palacios, School of Information, University of Arizona, USA

Michael S. Rosenberg, Center for Biological Data Science, Virginia Commonwealth University, USA

Jessica Ross, University of Wisconsin, USA

Freya E. Rowland, School of the Environment, Yale University, USA

Deusdedith Rugemalila, Institute of the Environment, Florida International University, USA

Avery L. Russell, Department of Biology, Missouri State University, USA

Suvi Ruuskanen, Department of Biological and Environmental Science, University of Jyväskylä, Finland

Patrick Saccone, Institute for Interdisciplinary Mountain Research, OeAW (Austrian Academy of Sciences), Austria

Asaf Sadeh, Department of Natural Resources, Newe Ya'ar Research Center, Agricultural Research Organization (Volcani Institute), Israel

Stephen M. Salazar, Department of Animal Behaviour, Bielefeld University, Germany

Kris Sales, Office for National Statistics, UK

Pablo Salmón, Institute of Avian Research "Vogelwarte Helgoland", Germany

Alfredo Sánchez-Tójar, Department of Evolutionary Biology, Bielefeld University, Germany

Leticia Pereira Santos, Ecology Department, Universidade Federal de Goiás, Brazil

Francesca Santostefano, University of Exeter, University of Exeter, UK

Hayden T. Schilling, New South Wales Department of Primary Industries Fisheries, Australia

Marcus Schmidt, Research Data Management, Leibniz Centre for Agricultural Landscape Research (ZALF), Germany

Tim Schmoll, Evolutionary Biology, Bielefeld University, Germany

Adam C. Schneider, Biology Department, University of Wisconsin-La Crosse, USA

Allie E. Schrock, Department of Evolutionary Anthropology, Duke University, USA

Julia Schroeder, Department of Life Sciences, Imperial College London, UK

Nicolas Schtickzelle, Earth and Life Institute, Ecology and Biodiversity, UCLouvain, Belgium

Nick L. Schultz, Future Regions Research Centre, Federation University Australia, Australia

Drew A. Scott, United States Department of Agriculture- Agricultural Research Service-, USA

Michael Peter Scroggie, Arthur Rylah Insitute for Environmental Research, Australia

Julie Teresa Shapiro, Epidemiology and Surveillance Support Unit, University of Lyon - French Agency for Food, Environmental and Occupational Health and Safety (ANSES), France

Nitika Sharma, UCLA Anderson Center for Impact, University of California, Los Angeles, USA

Caroline L. Shearer, Department of Evolutionary Anthropology, Duke University, USA

Diego Simón, Facultad de Ciencias, Universidad de la República, Uruguay

Michael I. Sitvarin, Independent researcher, USA

Fabrício Luiz Skupien, Programa de Pós-Graduação em Ecologia, Instituto de Biologia, Centro de Ciências da Saúde, Universidade Federal do Rio de Janeiro, Brazil

Heather Lea Slinn, Vive Crop Protection, Canada

Grania Polly Smith, University of Cambridge, UK

Jeremy A. Smith, British Trust for Ornithology, UK

Rahel Sollmann, Department of Wildlife, Fish, and Conservation Biology, University of California, Davis, USA

Kaitlin Stack Whitney, Science, Technology & Society Department, Rochester Institute of Technology, USA

Shannon Michael Still, Nomad Ecology, USA

Erica F. Stuber, Wildland Resources Department, Utah State University, USA

Guy F. Sutton, Center for Biological Control, Department of Zoology and Entomology, Rhodes University, South Africa

Ben Swallow, School of Mathematics and Statistics and Centre for Research in Ecological and Environmental Modelling, University of St Andrews, UK

Conor Claverie Taff, Department of Ecology and Evolutionary Biology, Cornell University, USA

Elina Takola, Department of Computational Landscape Ecology, Helmholtz Centre for Environmental Research – UFZ, Germany

Andrew J. Tanentzap, Ecosystems and Global Change Group, School of the Environment, Trent University, Canada

Rocío Tarjuelo, Instituto Universitario de Investigación en Gestión Forestal Sostenible (iuFOR), Universidad de Valladolid, Spain

Richard J. Telford, Department of Biological Sciences, University of Bergen, Norway

Christopher J. Thawley, Department of Biological Science, University of Rhode Island, USA

Hugo Thierry, Department of Geography, McGill University, Canada

Jacqueline Thomson, Integrative Biology, University of Guelph, Canada

Svenja Tidau, School of Biological and Marine Sciences, University of Plymouth, UK

Emily M. Tompkins, Biology Deptartment, Wake Forest University, USA

Claire Marie Tortorelli, Plant Sciences, University of California, Davis, USA

Andrew Trlica, College of Natural Resources, North Carolina State University, USA

Biz R. Turnell, Institute of Zoology, Technische Universität Dresden, Germany

Lara Urban, Helmholtz AI, Helmholtz Zentrum Muenchen, Germany

Stijn Van de Vondel, Department of Biology, University of Antwerp, Belgium

Jessica Eva Megan van der Wal, FitzPatrick Institute of African Ornithology, University of Cape Town, South Africa

Jens Van Eeckhoven, Department of Cell & Developmental Biology, Division of Biosciences, University College London, UK

Francis van Oordt, Natural Resource Sciences, McGill University, Canada

K. Michelle Vanderwel, Biology, University of Saskatchewan, Canada

Mark C. Vanderwel, Department of Biology, University of Regina, Canada

Karen J. Vanderwolf, Biology, University of Waterloo, Canada

Juliana Vélez, Department of Fisheries, Wildlife and Conservation Biology, University of Minnesota, USA

Diana Carolina Vergara-Florez, Department of Ecology & Evolutionary Biology, University of Michigan, USA

Brian C. Verrelli, Center for Biological Data Science, Virginia Commonwealth University, USA

Marcus Vinícius Vieira, Dept. Ecologia, Instituto de Biologia, Universidade Federal do Rio de Janeiro, Brazil

Nora Villamil, Lothian Analytical Services, Public Health Scotland, UK

Valerio Vitali, Institute for Evolution and Biodiversity, University of Muenster, Germany

Julien Vollering, Department of Environmental Sciences, Western Norway University of Applied Sciences, Norway

Jeffrey Walker, Department of Biological Sciences, University of Southern Maine, USA

Xanthe J. Walker, Center for Ecosystem Science and Society, Northern Arizona University, USA

Jonathan A. Walter, Center for Watershed Sciences, University of California, Davis, USA

Pawel Waryszak, School of Agriculture and Environmental Science, University of Southern Queensland, Australia

Ryan J. Weaver, Department of Ecology, Evolution, and Organismal Biology, Iowa State University, USA

Ronja E. M. Wedegärtner, Fram Project AS, Norway

Daniel L. Weller, Department of Food Science & Technology, Virginia Polytechnic Institute and State University, USA

Shannon Whelan, Department of Natural Resource Sciences, McGill University, Canada

Rachel Louise White, School of Applied Sciences, University of Brighton, UK

David William Wolfson, Department of Fisheries, Wildlife and Conservation Biology, University of Minnesota, USA

Andrew Wood, Department of Biology, University of Oxford, UK

we recommend viewing this manuscript in html format at https://egouldo.github.io/ManyAnalysts/

Scott W. Yanco, Department of Integrative Biology, University of Colorado, Denver, USA

Jian D. L. Yen, Arthur Rylah Institute for Environmental Research, Australia

Casey Youngflesh, Ecology, Evolution, and Behavior Program, Michigan State University, USA

Giacomo Zilio, ISEM, University of Montpellier, CNRS, France

Cédric Zimmer, Laboratoire d'Ethologie Expérimentale et Comparée, LEEC, UR4443, Université Sorbonne Paris Nord, USA

Gregory Mark Zimmerman, Department of Science and Environment, Lake Superior State University, USA

Rachel A. Zitomer, Department of Forest Ecosystems and Society, Oregon State University, USA

Abstract

Although variation in effect sizes and predicted values among studies of similar phenomena is inevitable, such variation far exceeds what might be produced by sampling error alone. One possible explanation for variation among results is differences among researchers in the decisions they make regarding statistical analyses. A growing array of studies has explored this analytical variability in different (mostly social science) fields, and has found substantial variability among results, despite analysts having the same data and research question. We implemented an analogous study in ecology and evolutionary biology, fields in which there have been no empirical exploration of the variation in effect sizes or model predictions generated by the analytical decisions of different researchers. We used two unpublished datasets, one from evolutionary ecology (blue tit, *Cyanistes caeruleus*, to compare sibling number and nestling growth) and one from conservation ecology (*Eucalyptus*, to compare grass cover and tree seedling recruitment), and the project leaders recruited 174 analyst teams, comprising 246 analysts, to investigate the answers to prespecified research questions. Analyses conducted by these teams yielded 141 usable effects for the blue tit dataset, and 85 usable effects for the *Eucalyptus* dataset. We found substantial heterogeneity among results for both datasets, although the patterns of variation differed between them. For the blue tit analyses, the average effect was convincingly negative, with less growth for nestlings living with more siblings, but there was near continuous variation in effect size from large negative effects to effects near zero, and even effects crossing the traditional threshold of statistical significance in the opposite direction. In contrast, the average relationship between grass cover and *Eucalyptus* seedling number was only slightly negative and not convincingly different from zero, and most effects ranged from weakly negative to weakly positive, with about a third of effects crossing the traditional threshold of significance in one direction or the other. However, there were also several striking outliers in the *Eucalyptus* dataset, with effects far from zero. For both datasets, we found substantial variation in the variable selection and random effects structures among analyses, as well as in the ratings of the analytical methods by peer reviewers, but we found no strong relationship between any of these

and deviation from the meta-analytic mean. In other words, analyses with results that were far from the mean were no more or less likely to have dissimilar variable sets, use random effects in their models, or receive poor peer reviews than those analyses that found results that were close to the mean. The existence of substantial variability among analysis outcomes raises important questions about how ecologists and evolutionary biologists should interpret published results, and how they should conduct analyses in the future.

Key Words

credibility revolution, heterogeneity, meta-analysis, metascience, Replicability, reproducibility

Introduction

One value of science derives from its production of replicable, and thus reliable, results. When we repeat a study using the original methods we should be able to expect a similar result. However, perfect replicability is not a reasonable goal. Effect sizes will vary, and even reverse in sign, by chance alone [1]. Observed patterns can differ for other reasons as well. It could be that we do not sufficiently understand the conditions that led to the original result so when we seek to replicate it, the conditions differ due to some 'hidden moderator'. This hidden moderator hypothesis is described by meta-analysts in ecology and evolutionary biology as 'true biological heterogeneity' [2]. This idea of true heterogeneity is popular in ecology and evolutionary biology, and there are good reasons to expect it in the complex systems in which we work [3]. However, despite similar expectations in psychology, recent evidence in that discipline contradicts the hypothesis that moderators are common obstacles to replicability, as variability in results in a large 'many labs' collaboration was mostly unrelated to commonly hypothesized moderators such as the conditions under which the studies were administered [4]. Another possible explanation for variation in effect sizes is that researchers often present biased samples of results, thus reducing the likelihood that later studies will produce similar effect sizes [5–9]. It also may be that although researchers did

successfully replicate the conditions, the experiment, and measured variables, analytical decisions differed sufficiently among studies to create divergent results [10, 11].

Analytical decisions vary among studies because researchers have many options. Researchers need to decide how to exclude possibly anomalous or unreliable data, how to construct variables, which variables to include in their models, and which statistical methods to use. Depending on the dataset, this short list of choices could encompass thousands or millions of possible alternative specifications [10]. However, researchers making these decisions presumably do so with the goal of doing the best possible analysis, or at least the best analysis within their current skill set. Thus it seems likely that some specification options are more probable than others, possibly because they have previously been shown (or claimed) to be better, or because they are more well known. Of course, some of these different analyses (maybe many of them) may be equally valid alternatives. Regardless, on probably any topic in ecology and evolutionary biology, we can encounter differences in choices of data analysis. The extent of these differences in analyses and the degree to which these differences influence the outcomes of analyses and therefore studies' conclusions are important empirical questions. These questions are especially important given that many papers draw conclusions after applying a single method, or even a single statistical model, to analyze a dataset.

The possibility that different analytical choices could lead to different outcomes has long been recognized [12], and various efforts to address this possibility have been pursued in the literature. For instance, one common method in ecology and evolutionary biology involves creating a set of candidate models, each consisting of a different (though often similar) set of predictor variables, and then, for the predictor variable of interest, averaging the slope across all models (i.e. model averaging) [13, 14]. This method reduces the chance that a conclusion is contingent upon a single model specification, though use and interpretation of this method is not without challenges [14]. Further, the models compared to each other typically differ only in the inclusion or exclusion of certain predictor variables and not in other important ways, such as methods of parameter

484 estimation. More explicit examination of outcomes of differences in model structure, model type,
485 data exclusion, or other analytical choices can be implemented through sensitivity
486 analyses [e.g., 15]. Sensitivity analyses, however, are typically rather narrow in scope, and are
487 designed to assess the sensitivity of analytical outcomes to a particular analytical choice rather than
488 to a large universe of choices. Recently, however, analysts in the social sciences have proposed
489 extremely thorough sensitivity analysis, including 'multiverse analysis' [16] and the 'specification
490 curve' [10], as a means of increasing the reliability of results. With these methods, researchers
491 identify relevant decision points encountered during analysis and conduct the analysis many times
492 to incorporate many plausible decisions made at each of these points. The study's conclusions are
493 then based on a broad set of the possible analyses and so allow the analyst to distinguish between
494 robust conclusions and those that are highly contingent on particular model specifications. These are
495 useful outcomes, but specifying a universe of possible modelling decisions is not a trivial
496 undertaking. Further, the analyst's knowledge and biases will influence decisions about the
497 boundaries of that universe, and so there will always be room for disagreement among analysts
498 about what to include. Including more specifications is not necessarily better. Some analytical
499 decisions are better justified than others, and including biologically implausible specifications may
500 undermine this process. Regardless, these powerful methods have yet to be adopted, and even
501 more limited forms of sensitivity analyses are not particularly widespread. Most studies publish a
502 small set of analyses and so the existing literature does not provide much insight into the degree to
503 which published results are contingent on analytical decisions.

504 Despite the potential major impacts of analytical decisions on variance in results, the outcomes of
505 different individuals' data analysis choices have received limited empirical attention. The only formal
506 exploration of this that we were aware of when we submitted our Stage 1 manuscript were (1) an
507 analysis in social science that asked whether male professional football (soccer) players with darker
508 skin tone were more likely to be issued red cards (ejection from the game for rule violation) than

509 players with lighter skin tone [11] and (2) an analysis in neuroimaging which evaluated nine separate
510 hypotheses involving the neurological responses detected with fMRI in 108 participants divided
511 between two treatments in a decision making task [17]. Several others have been published
512 since [e.g., 18, 19–21]. In the red card study, twenty-nine teams designed and implemented analyses
513 of a dataset provided by the study coordinators [11]. Analyses were peer reviewed (results blind) by
514 at least two other participating analysts; a level of scrutiny consistent with standard pre-publication
515 peer review. Among the final 29 analyses, odds-ratios varied from 0.89 to 2.93, meaning point
516 estimates varied from having players with lighter skin tones receive more red cards (odds ratio < 1)
517 to a strong effect of players with darker skin tones receiving more red cards (odds ratio > 1). Twenty
518 of the 29 teams found a statistically-significant effect in the predicted direction of players with
519 darker skin tones being issued more red cards. This degree of variation in peer-reviewed analyses
520 from identical data is striking, but the generality of this finding has only just begun to be formally
521 investigated.

522 In the neuroimaging study, 70 teams evaluated each of the nine different hypotheses with the
523 available fMRI data [17]. These 70 teams followed a divergent set of workflows that produced a wide
524 range of results. The rate of reporting of statistically significant support for the nine hypotheses
525 ranged from 21% to 84%, and for each hypothesis on average, 20% of research teams observed
526 effects that differed substantially from the majority of other teams. Some of the variability in results
527 among studies could be explained by analytical decisions such as choice of software package,
528 smoothing function, and parametric versus non-parametric corrections for multiple comparisons.
529 However, substantial variability among analyses remained unexplained, and presumably emerged
530 from the many different decisions each analyst made in their long workflows. Such variability in
531 results among analyses from this dataset and from the very different red-card dataset suggests that
532 sensitivity of analytical outcome to analytical choices may characterize many distinct fields, as
533 several more recent many-analyst studies also suggest [18–20].

To further develop the empirical understanding of the effects of analytical decisions on study outcomes, we chose to estimate the extent to which researchers' data analysis choices drive differences in effect sizes, model predictions, and qualitative conclusions in ecology and evolutionary biology. This is an important extension of the meta-research agenda of evaluating factors influencing replicability in ecology, evolutionary biology, and beyond [22]. To examine the effects of analytical decisions, we used two different datasets and recruited researchers to analyze one or the other of these datasets to answer a question we defined. The first question was "To what extent is the growth of nestling blue tits (*Cyanistes caeruleus*) influenced by competition with siblings?" To answer this question, we provided a dataset that includes brood size manipulations from 332 broods conducted over three years at Wytham Wood, UK. The second question was "How does grass cover influence *Eucalyptus* spp. seedling recruitment?" For this question, analysts used a dataset that includes, among other variables, number of seedlings in different size classes, percentage cover of different life forms, tree canopy cover, and distance from canopy edge from 351 quadrats spread among 18 sites in Victoria, Australia.

We explored the impacts of data analysts' choices with descriptive statistics and with a series of tests to attempt to explain the variation among effect sizes and predicted values of the dependent variable produced by the different analysis teams for both datasets separately. To describe the variability, we present forest plots of the standardized effect sizes and predicted values produced by each of the analysis teams, estimate heterogeneity (both absolute, τ^2, and proportional, I^2) in effect size and predicted values among the results produced by these different teams, and calculate a similarity index that quantifies variability among the predictor variables selected for the different statistical models constructed by the different analysis teams. These descriptive statistics provide the first estimates of the extent to which explanatory statistical models and their outcomes in ecology and evolutionary biology vary based on the decisions of different data analysts. We then quantified the degree to which the variability in effect size and predicted values could be explained by (1) variation

we recommend viewing this manuscript in html format at https://egouldo.github.io/ManyAnalysts/

in the quality of analyses as rated by peer reviewers and (2) the similarity of the choices of predictor variables between individual analyses.

Methods

This project involved a series of steps (1-6) that began with identifying datasets for analyses and continued through recruiting independent groups of scientists to analyze the data, allowing the scientists to analyze the data as they saw fit, generating peer review ratings of the analyses (based on methods, not results), evaluating the variation in effects among the different analyses, and producing the final manuscript.

Step 1: Select Datasets

We used two previously unpublished datasets, one from evolutionary ecology and the other from ecology and conservation.

Evolutionary Ecology

Our evolutionary ecology dataset is relevant to a sub-discipline of life-history research which focuses on identifying costs and trade-offs associated with different phenotypic conditions.

These data were derived from a brood-size manipulation experiment imposed on wild birds nesting in boxes provided by researchers in an intensively studied population.

Understanding how the growth of nestlings is influenced by the numbers of siblings in the nest can give researchers insights into factors such as the evolution of clutch size, determination of provisioning rates by parents, and optimal levels of sibling competition (Vander Werf 1992; DeKogel 1997; Royle et al. 1999; Verhulst, Holveck, and Riebel 2006; Nicolaus et al. 2009). Data analysts were provided this dataset and instructed to answer the following question: “To what extent is the growth of nestling blue tits (*Cyanistes caeruleus*) influenced by competition with siblings?”

Researchers conducted brood size manipulations and population monitoring of blue tits at Wytham Wood, a 380ha woodland in Oxfordshire, U.K (1º 20’W, 51º 47’N). Researchers regularly checked

we recommend viewing this manuscript in html format at https://egouldo.github.io/ManyAnalysts/

584 approximately 1100 artificial nest boxes at the site and monitored the 330 to 450 blue tit pairs
585 occupying those boxes in 2001-2003 during the experiment. Nearly all birds made only one breeding
586 attempt during the April to June study period in a given year. At each blue tit nest, researchers
587 recorded the date the first egg appeared, clutch size, and hatching date. For all chicks alive at age 14
588 days, researchers measured mass and tarsus length and fitted a uniquely numbered, British Trust for
589 Ornithology (BTO) aluminium leg ring. Researchers attempted to capture all adults at their nests
590 between day 6 and day 14 of the chick-rearing period. For these captured adults, researchers
591 measured mass, tarsus length, and wing length and fitted a uniquely numbered BTO leg ring. During
592 the 2001-2003 breeding seasons, researchers manipulated brood sizes using cross fostering. They
593 matched broods for hatching date and brood size and moved chicks between these paired nests one
594 or two days after hatching. They sought to either enlarge or reduce all manipulated broods by
595 approximately one fourth. To control for effects of being moved, each reduced brood had a portion
596 of its brood replaced by chicks from the paired increased brood, and vice versa. Net manipulations
597 varied from plus or minus four chicks in broods of 12 to 16 to plus or minus one chick in broods of 4
598 or 5. Researchers left approximately one third of all broods unmanipulated. These unmanipulated
599 broods were not selected systematically to match manipulated broods in clutch size or laying date.
600 We have mass and tarsus length data from 3720 individual chicks divided among 167 experimentally
601 enlarged broods, 165 experimentally reduced broods, and 120 unmanipulated broods. The full list of
602 variables included in the dataset is publicly available (https://osf.io/hdv8m), along with the data
603 (https://osf.io/qjzby).

Additional explanation:

Shortly after beginning to recruit analysts, several analysts noted a small set of related errors in the blue tit dataset. We corrected the errors, replaced the dataset on our OSF site, and emailed the analysts on 19 April 2020 to instruct them to use the revised data. The email to analysts is available here (https://osf.io/4h53z). The errors are explained in that email.

604

we recommend viewing this manuscript in html format at https://egouldo.github.io/ManyAnalysts/

Ecology and Conservation

Our ecology and conservation dataset is relevant to a sub-discipline of conservation research which focuses on investigating how best to revegetate private land in agricultural landscapes. These data were collected on private land under the Bush Returns program, an incentive system where participants entered into a contract with the Goulburn Broken Catchment Management Authority and received annual payments if they executed predetermined restoration activities. This particular dataset is based on a passive regeneration initiative, where livestock grazing was removed from the property in the hopes that the Eucalyptus spp. overstorey would regenerate without active (and expensive) planting. Analyses of some related data have been published (Miles 2008; Vesk et al. 2016) but those analyses do not address the question analysts answered in our study. Data analysts were provided this dataset and instructed to answer the following question: "How does grass cover influence *Eucalyptus* spp. seedling recruitment?".

Researchers conducted three rounds of surveys at 18 sites across the Goulburn Broken catchment in northern Victoria, Australia in winter and spring 2006 and autumn 2007. In each survey period, a different set of 15 x 15 m quadrats were randomly allocated across each site within 60 m of existing tree canopies. The number of quadrats at each site depended on the size of the site, ranging from four at smaller sites to 11 at larger sites. The total number of quadrats surveyed across all sites and seasons was 351. The number of Eucalyptus spp. seedlings was recorded in each quadrat along with information on the GPS location, aspect, tree canopy cover, distance to tree canopy, and position in the landscape. Ground layer plant species composition was recorded in three 0.5 x 0.5 m sub-quadrats within each quadrat. Subjective cover estimates of each species as well as bare ground, litter, rock and moss/lichen/soil crusts were recorded. Subsequently, this was augmented with information about the precipitation and solar radiation at each GPS location. The full list of variables included in the dataset is publicly available (https://osf.io/r5gbn), along with the data (https://osf.io/qz5cu).

we recommend viewing this manuscript in html format at https://egouldo.github.io/ManyAnalysts/

Step 2: Recruitment and initial survey of analysts

The lead team (TP, HF, SN, EG, SG, PV, DH, FF) created a publicly available document providing a general description of the project (https://osf.io/mn5aj/). The project was advertised at conferences, via Twitter, using mailing lists for ecological societies (including Ecolog, Evoldir, and lists for the Environmental Decisions Group, and Transparency in Ecology and Evolution), and via word of mouth. The target population was active ecology, conservation, or evolutionary biology researchers with a graduate degree (or currently studying for a graduate degree) in a relevant discipline. Researchers could choose to work independently or in a small team. For the sake of simplicity, we refer to these as 'analysis teams' though some comprised one individual. We aimed for a minimum of 12 analysis teams independently evaluating each dataset (see sample size justification below). We simultaneously recruited volunteers to peer review the analyses conducted by the other volunteers through the same channels. Our goal was to recruit a similar number of peer reviewers and analysts, and to ask each peer reviewer to review a minimum of four analyses. If we were unable to recruit at least half the number of reviewers as analysis teams, we planned to ask analysts to serve also as reviewers (after they had completed their analyses), but this was unnecessary. All analysts and reviewers were offered the opportunity to share co-authorship on this manuscript and we planned to invite them to participate in the collaborative process of producing the final manuscript. All analysts signed [digitally] a consent (ethics) document (https://osf.io/xyp68/) approved by the Whitman College Institutional Review Board prior to being allowed to participate.

Preregistration Deviation:

Due to the large number of recruited analysts and reviewers and the anticipated challenges of receiving and integrating feedback from so many authors, we limited analyst and reviewer participation in the production of the final manuscript to an invitation to call attention to serious problems with the manuscript draft.

we recommend viewing this manuscript in html format at https://egouldo.github.io/ManyAnalysts/

We identified our minimum number of analysts per dataset by considering the number of effects needed in a meta-analysis to generate an estimate of heterogeneity (τ^2) with a 95% confidence interval that does not encompass zero. This minimum sample size is invariant regardless of τ^2. This is because the same t-statistic value will be obtained by the same sample size regardless of variance (τ^2). We see this by first examining the formula for the standard error, SE for variance, (τ^2) or SE(τ^2) assuming normality in an underlying distribution of effect sizes [30]:

$$SE(\tau 2) = \sqrt{\frac{t^4}{n-1}}$$

and then rearranging the above formula to show how the t-statistic is independent of τ^2, as seen below.

$$t = \frac{\tau^2}{SE\ \tau^2} = \sqrt{\frac{n-1}{2}}$$

We then find a minimum n = 12 according to this formula.

Step 3: Primary Data Analysis

Analysis teams registered and answered a demographic and expertise survey (https://osf.io/seqzy/). We then provided them with the dataset of their choice and requested that they answer a specific research question. For the evolutionary ecology dataset that question was “To what extent is the growth of nestling blue tits (*Cyanistes caeruleus*) influenced by competition with siblings?” and for the conservation ecology dataset it was “How does grass cover influence *Eucalyptus* spp. seedling recruitment?” Once their analysis was complete, they answered a structured survey (https://osf.io/neyc7/), providing analysis technique, explanations of their analytical choices, quantitative results, and a statement describing their conclusions. They also were asked to upload their analysis files (including the dataset as they formatted it for analysis and their analysis code [if applicable]) and a detailed journal-ready statistical methods section.

we recommend viewing this manuscript in html format at https://egouldo.github.io/ManyAnalysts/

Preregistration Deviation:

We originally planned to have analysts complete a single survey (https://osf.io/neyc7/), but after we evaluated the results of that survey, we realized we would need a second survey (https://osf.io/8w3v5/) to adequately collect the information we needed to evaluate heterogeneity of results (step 5). We provided a set of detailed instructions with the follow-up survey, and these instructions are publicly available and can be found within the following files (blue tit: https://osf.io/kr2g9, *Eucalyptus*: https://osf.io/dfvym).

Step 4: Peer Review of Analysis

At minimum, each analysis was evaluated by four different reviewers, and each volunteer peer reviewer was randomly assigned methods sections from at least four analyst teams (the exact number varied). Each peer reviewer registered and answered a demographic and expertise survey identical to that asked of the analysts, except we did not ask about 'team name' since reviewers did not work in teams. Reviewers evaluated the methods of each of their assigned analyses one at a time in a sequence determined by the project leaders. We systematically assigned the sequence so that, if possible, each analysis was allocated to each position in the sequence for at least one reviewer. For instance, if each reviewer were assigned four analyses to review, then each analysis would be the first analysis assigned to at least one reviewer, the second analysis assigned to another reviewer, the third analysis assigned to yet another reviewer, and the fourth analysis assigned to a fourth reviewer. Balancing the order in which reviewers saw the analyses controls for order effects, e.g. a reviewer might be less critical of the first methods section they read than the last.

The process for a single reviewer was as follows. First, the reviewer received a description of the methods of a single analysis. This included the narrative methods section, the analysis team's answers to our survey questions regarding their methods, including analysis code, and the dataset. The reviewer was then asked, in an online survey (https://osf.io/4t36u/), to rate that analysis on a

scale of 0-100 based on this prompt: “Rate the overall appropriateness of this analysis to answer the research question (one of the two research questions inserted here) with the available data. To help you calibrate your rating, please consider the following guidelines:

100: A perfect analysis with no conceivable improvements from the reviewer
75: An imperfect analysis but the needed changes are unlikely to dramatically alter outcomes
50: A flawed analysis likely to produce either an unreliable estimate of the relationship or an over-precise estimate of uncertainty
25: A flawed analysis likely to produce an unreliable estimate of the relationship and an over-precise estimate of uncertainty
0: A dangerously misleading analysis, certain to produce both an estimate that is wrong and a substantially over-precise estimate of uncertainty that places undue confidence in the incorrect estimate.
*Please note that these values are meant to calibrate your ratings. We welcome ratings of any number between 0 and 100.”

After providing this rating, the reviewer was presented with this prompt, in multiple-choice format: “Would the analytical methods presented produce an analysis that is (a) publishable as is, (b) publishable with minor revision, (c) publishable with major revision, (d) deeply flawed and unpublishable?” The reviewer was then provided with a series of text boxes and the following prompts: “Please explain your ratings of this analysis. Please evaluate the choice of statistical analysis type. Please evaluate the process of choosing variables for and structuring the statistical model. Please evaluate the suitability of the variables included in (or excluded from) the statistical model. Please evaluate the suitability of the structure of the statistical model. Please evaluate choices to exclude or not exclude subsets of the data. Please evaluate any choices to transform data (or, if there were no transformations, but you think there should have been, please discuss that choice).” After

we recommend viewing this manuscript in html format at https://egouldo.github.io/ManyAnalysts/

submitting this review, a methods section from a second analysis was then made available to the reviewer. This same sequence was followed until all analyses allocated to a given reviewer were provided and reviewed. After providing the final review, the reviewer was simultaneously provided with all four (or more) methods sections the reviewer had just completed reviewing, the option to revise their original ratings, and a text box to provide an explanation. The invitation to revise the original ratings was as follows: "If, now that you have seen all the analyses you are reviewing, you wish to revise your ratings of any of these analyses, you may do so now." The text box was prefaced with this prompt: "Please explain your choice to revise (or not to revise) your ratings."

Additional Explanation:

To determine how consistent peer reviewers were in their ratings, we assessed inter-rater reliability among reviewers for both the categorical and quantitative ratings combining blue tit and *Eucalyptus* data using Krippendorff's alpha for ordinal and continuous data respectively. This provides a value that is between -1 (total disagreement between reviewers) and 1 (total agreement between reviewers).

Step 5: Evaluate Variation

The lead team conducted the analyses outlined in this section. We described the variation in model specification in several ways. We calculated summary statistics describing variation among analyses, including mean, SD, and range of number of variables per model included as fixed effects, the number of interaction terms, the number of random effects, and the mean, SD, and range of sample sizes. We also present the number of analyses in which each variable was included. We summarized the variability in standardized effect sizes and predicted values of dependent variables among the individual analyses using standard random effects meta-analytic techniques. First, we derived standardized effect sizes from each individual analysis. We did this for all linear models or generalized linear models by converting the t value and the degree of freedom (df) associated with

regression coefficients (e.g. the effect of the number of siblings [predictor] on growth [response] or the effect of grass cover [predictor] on seedling recruitment [response]) to the correlation coefficient (r), using the following:

$$r = \sqrt{\frac{t^2}{t^2 + df}}$$

This formula can only be applied if t and df values originate from linear or generalized linear models [GLMs; [31]]. If, instead, linear mixed-effects models (LMMs) or generalized linear mixed-effects models (GLMMs) were used by a given analysis, the exact df cannot be estimated. However, adjusted df can be estimated, for example, using the Satterthwaite approximation of df, df_s, [note that SAS uses this approximation to obtain df for LMMs and GLMMs; [32]]. For analyses using either LMMs or GLMMs that do not produce df_s we planned to obtain df_s by rerunning the same (G)LMMs using the lmer() or glmer() function in the lmerTest package in R [33, 34].

Preregistration Deviation

Rather than re-run these analyses ourselves, we sent a follow-up survey (referenced above under "Primary data analyses") to analysts and asked them to follow our instructions for producing this information. The instructions are publicly available and can be found within the following files (blue tit: https://osf.io/kr2g9, *Eucalyptus*: https://osf.io/dfvym).

We then used the t values and df_s from the models to obtain r as per the formula above. All r and accompanying df (or df_s) were converted to Zr and it's sampling variance 1/(n-3) where n=df+1. Any analyses from which we could not derive a signed Zr, for instance one with a quadratic function in which the slope changed sign, were excluded from the analyses of Fisher's Zr. We expected such analyses would be rare. In fact, most submitted analyses excluded from our meta-analysis of Zr were excluded because of a lack of sufficient information provided by the analyst team rather than due to the use of effects that could not be converted to Zr. Regardless, as we describe below, we generated

754 a second set of standardized effects (predicted values) that could (in principle) be derived from any
755 explanatory model produced by these data.
756 Besides Zr, which describes the strength of a relationship based on the amount of variation in a
757 dependent variable explained by variation in an independent variable, we also examined differences
758 in the shape of the relationship between the independent and dependent variables. To accomplish
759 this, we derived a point estimate (out-of-sample predicted value) for the dependent variable of
760 interest for each of three values of our primary independent variable. We originally described these
761 three values as associated with the 25th percentile, median, and 75th percentile of the independent
762 variable and any covariates.

Preregistration Deviation

The original description of the out-of-sample specifications did not account for the facts that (a) some variables are not distributed in a way that allowed division in percentiles and that (b) variables could be either positively or negatively correlated with the dependent variable. We provide a more thorough description here: We derived three point-estimates (out-of-sample predicted values) for the dependent variable of interest; one for each of three values of our primary independent variable that we specified. We also specified values for all other variables that could have been included as independent variables in analysts' models so that we could derive the predicted values from a fully specified version of any model produced by analysts. For all potential independent variables, we selected three values or categories. Of the three we selected, one was associated with small, one with intermediate, and one with large values of one typical dependent variable (day 14 chick weight for the blue tit data and total number of seedlings for the *Eucalyptus* data; analysts could select other variables as their dependent variable, but the others typically correlated with the two identified here). For continuous variables, this means we identified the 25th percentile, median, and 75th percentile and, if the slope of the linear relationship between this variable and the typical dependent variable was positive, we left the quartiles ordered as is. If, instead, the slope was negative, we reversed the order of the independent variable quartiles so that the 'lower' quartile value was the one associated with the lower value for the dependent variable. In the case of categorical variables, we identified categories associated with the 25th percentile, median, and 75th percentile values of the typical dependent variable after averaging the values for each category. However, for some continuous and categorical predictors, we also made selections based on the principle of internal consistency between certain related variables, and we fixed a few categorical variables as identical across all three levels where doing so would simplify the modelling process (specification tables available: blue tit: https://osf.io/86akx; *Eucalyptus*: https://osf.io/jh7g5).

764 We used the 25th and 75th percentiles rather than minimum and maximum values to reduce the
765 chance of occupying unrealistic parameter space. We planned to derive these predicted values from
766 the model information provided by the individual analysts. All values (predictions) were first
767 transformed to the original scale along with their standard errors (SE); we used the delta method
768 (Ver Hoef 2012) for the transformation of SE. We used the square of the SE associated with predicted
769 values as the sampling variance in the meta-analyses described below, and we planned to analyze
770 these predicted values in exactly the same ways as we analyzed Zr in the following analyses.

Preregistration Deviation

Because analysts of blue tit data chose different dependent variables on different scales, after transforming out-of-sample values to the original scales, we standardized all values as z scores ('standard scores') to put all dependent variables on the same scale and make them comparable. This involved taking each relevant value on the original scale (whether a predicted point estimate or a SE associated with that estimate) and subtracting the value in question from the mean value of that dependent variable derived from the full dataset and then dividing this difference by the standard deviation, SD, corresponding to the mean from the full dataset. Thus, all our out-of-sample prediction values from the blue tit data are from a distribution with the mean of 0 and SD of 1. We did not add this step for the Eucalyptus data because (a) all responses were on the same scale (counts of *Eucalyptus* stems) and were thus comparable and (b) these data, with many zeros and high skew, are poorly suited for z scores.

771

772 We plotted individual effect size estimates (Zr) and predicted values of the dependent variable (y_i)
773 and their corresponding 95% confidence / credible intervals in forest plots to allow visualization of
774 the range and precision of effect size and predicted values. Further, we included these estimates in
775 random effects meta-analyses [36, 37] using the metafor package in R [34, 38]:

776 $Zr \sim 1 + 1|analysisId$

777 $y_i \sim 1 + 1|analysisId$

778 where y_i is the predicted value for the dependent variable at the 25th percentile, median, or 75th
779 percentile of the independent variables. The individual Zr effect sizes were weighted with the inverse
780 of sampling variance for Zr. The individual predicted values for dependent variable (y_i) were weighted
781 by the inverse of the associated SE^2 original registration omitted “inverse of the” in error). These
782 analyses provided an average Zr score or an average y_i with corresponding 95% confidence interval
783 and allowed us to estimate two heterogeneity indices, τ^2 and I^2. The former, τ^2, is he absolute
784 measure of heterogeneity or the between-study variance (in our case, between-effect variance)
785 whereas I^2 is a relative measure of heterogeneity. We obtained the estimate of relative heterogeneity
786 (I^2) by dividing the between-effect variance by the sum of between-effect and within-effect variance
787 (sampling error variance). I^2 is thus, in a standard meta-analysis, the proportion of variance that is
788 due to heterogeneity as opposed to sampling error. When calculating I^2, within-study variance is
789 amalgamated across studies to create a “typical” within-study variance which serves as the sampling
790 error variance [36, 37]. Our goal here was to visualize and quantify the degree of variation among
791 analyses in effect size estimates [31]. We did not test for statistical significance.

Additional Explanation

Our use of I^2 to quantify heterogeneity violates an important assumption, but this violation does not invalidate our use of I^2 as a metric of how much heterogeneity can derive from analytical decisions. In standard meta-analysis, the statistic I^2 quantifies the proportion of variance that is greater than we would expect if differences among estimates were due to sampling error alone [39]. However, it is clear that this interpretation does not apply to our value of I^2 because I^2 assumes that each estimate is based on an independent sample (although these analyses can account for non-independence via hierarchical modelling), whereas all our effects were derived from largely or entirely overlapping subsets of the same dataset. Despite this, we believe that I^2 remains a useful statistic for our purposes. This is because, in calculating I^2, we are still setting a benchmark of expected variation due to sampling error based on the variance associated with each separate effect size estimate, and we are assessing how much (if it all) the variability among our effect sizes exceeds what would be expected had our effect sizes been based on independent data. In other words, our estimates can tell us how much proportional heterogeneity is possible from analytical decisions alone when sample sizes (and therefore meta-analytic within-estimate variance) are similar to the ones in our analyses. Among other implications, our violation of the independent sample assumption means that we (dramatically) over-estimate the variance expected due to sampling error, and because I^2 s a proportional estimate, we thus underestimate the actual proportion of variance due to differences among analyses other than sampling error. However, correcting this underestimation would create a trivial value since we designed the study so that much of the variance would derive from analytic decisions as opposed to differences in sampled data. Instead, retaining the I^2 value as typically calculated provides a useful comparison to I^2 values from typical meta-analyses.

Interpretation of τ^2 also differs somewhat from traditional meta-analysis, and we discuss this further in the Results.

Finally, we assessed the extent to which deviations from the meta-analytic mean by individual effect sizes (Zr) or the predicted values of the dependent variable (y_i) were explained by the peer rating of each analysis team's method section, by a measurement of the distinctiveness of the set of predictor variables included in each analysis, and by the choice of whether or not to include random effects in the model. The deviation score, which served as the dependent variable in these analyses, is the absolute value of the difference between the meta-analytic mean Zr (or y_i) and the individual Zr (or y_i) estimate for each analysis. We used the Box-Cox transformation on the absolute values of deviation scores to achieve an approximately normal distribution [c.f. 40, 41]. We described variation in this dependent variable with both a series of univariate analyses and a multivariate analysis. All these analyses were general linear (mixed) models. These analyses were secondary to our estimation of variation in effect sizes described above. We wished to quantify relationships among variables, but we had no a priori expectation of effect size and made no dichotomous decisions about statistical significance.

When examining the extent to which reviewer ratings (on a scale from 0 to 100) explained deviation from the average effect (or predicted value), each analysis had been rated by multiple peer reviewers, so for each reviewer score to be included, we include each deviation score in the analysis multiple times. To account for the non-independence of multiple ratings of the same analysis, we planned to include analysis identity as a random effect in our general linear mixed model in the lme4 package in R [34, 42]. To account for potential differences among reviewers in their scoring of analyses, we also planned to include reviewer identity as a random effect:

$$\text{DeviationScore}_j = \text{BoxCox}(\textit{abs}(\text{DeviationFromMean}_j))$$

$$\text{DeviationScore}_{ij} \sim \text{Rating}_{ij} + \text{ReviewerID}_i + \text{AnalysisID}_j$$

$$\text{ReviewerID}_i \sim \mathcal{N}(0, \sigma^2)$$

$$\text{AnalysisID}_i \sim \mathcal{N}(0, \sigma^2)$$

Where $\text{DeviationFromMean}_j$ is the deviation from the meta-analytic mean for the jth analysis, ReviewerID_i is the random intercept assigned to each i reviewer, and AnalysisID_j is the random

intercept assigned to each j analysis, both of which are assumed to be normally distributed with a mean of 0 and a variance of σ^2 Absolute deviation scores were Box-Cox transformed using the step_box_cox() function from the timetk package in R [34, 43].

We conducted a similar analysis with the four categories of reviewer ratings ((1) deeply flawed and unpublishable, (2) publishable with major revision, (3) publishable with minor revision, (4) publishable as is) set as ordinal predictors numbered as shown here. As with the analyses above, we planned for these analyses to also include random effects of analysis identity and reviewer identity. Both of these analyses (1: 1-100 ratings as the fixed effect, 2: categorical ratings as the fixed effects) were planned to be conducted eight times for each dataset. Each of the four responses (Zr, y_{25}, y_{50}, y_{75}) were to be compared once to the initial ratings provided by the peer reviewers, and again based on the revised ratings provided by the peer reviewers.

we recommend viewing this manuscript in html format at https://egouldo.github.io/ManyAnalysts/

Preregistration Deviation

1. We planned to include random effects of both analysis identity and reviewer identity in these models comparing reviewer ratings with deviation scores. However, after we received the analyses, we discovered that a subset of analyst teams had either conducted multiple analyses and/or identified multiple effects per analysis as answering the target question. We therefore faced an even more complex potential set of random effects. We decided that including team ID, analysis ID, and effect ID along with reviewer ID as random effects in the same model would almost certainly lead to model fit problems, and so we started with simpler models including just effect ID and reviewer ID. However, even with this simpler structure, our dataset was sparse, with reviewers rating a small number of analyses, resulting in models with singular fit (Section C.2). Removing one of the random effects was necessary for the models to converge. The models that included the categorical quality rating converged when including reviewer ID, and the models that included the continuous quality rating converged when including effect ID.
2. We conducted analyses only with the final peer ratings after the opportunity for revision, not with the initial ratings. This was because when we recorded the final ratings, they over-wrote the initial ratings, and so we did not have access to those initial values.

The next set of univariate analyses sought to explain deviations from the mean effects based on a measure of the distinctiveness of the set of variables included in each analysis. As a 'distinctiveness' score, we used Sorensen's Similarity Index (an index typically used to compare species composition across sites), treating variables as species and individual analyses as sites. To generate an individual Sorensen's value for each analysis required calculating the pairwise Sorensen's value for all pairs of analyses (of the same dataset), and then taking the average across these Sorensen's values for each analysis. We calculated the Sorensen's index values using the betapart package [44] in R:

we recommend viewing this manuscript in html format at https://egouldo.github.io/ManyAnalysts/

838 $$\beta Sorensen = \frac{b + c}{2a + b + c}$$

839 Where *a* is the number of variables common to both analyses, *b* is the number of variables that
840 occur in the first analysis but not in the second and *c* is he number of variables that occur in the
841 second analysis. We then used the per-model average Sorensen's index value as an independent
842 variable to predict the deviation score in a general linear model, and included no random effect since
843 each analysis is included only once, in R [34]:

844 $$DeviationScore_j \sim \beta Sorensen$$

Additional Explanation

When we planned this analysis, we anticipated that analysts would identify a single primary effect from each model, so that each model would appear in the analysis only once. Our expecation was incorrect because some analysts identified >1 effect per analysis, but we still chose to specify our model as registered and not use a random effect. This is because most models produced only one effect and so we expected that specifying a random effect to account for the few cases where >1 effect was included for a given model would prevent model convergence.

Note that this analysis contrasts with the analyses in which we used reviewer ratings as predictors because in the analyses with reviewer ratings, each effect appeared in the analysis approximately four times due to multiple reviews of each analysis, and so it was much more important to account for that variance through a random effect.

845

846 Finally, we conducted a multivariate analysis with the five predictors described above (peer ratings 0-
847 100 and peer ratings of publishability 1-4; both original and revised and Sorensen's index, plus a
848 sixth, presence /absence of random effects) with random effects of analysis identity and reviewer
849 identity in the *lme4* package in R [34, 42]. We had stated here in the text that we would use only the
850 revised (final) peer ratings in this analysis, so the absence of the initial ratings is not a deviation from
851 our plan:

we recommend viewing this manuscript in html format at https://egouldo.github.io/ManyAnalysts/

852 $DeviationScore_j \sim RatingsContinuous_{ij} + RatingsCategorical_{ij} + \beta Sorensen_j + AnalysisID_j + ReviewerID_i$

853 $ReviewerID_i \sim \mathcal{N}(0, \sigma^2)$

854 $AnalysisID_j \sim \mathcal{N}(0, \sigma^2)$

855 We conducted all the analyses described above eight times; for each of the four responses (Zr, y_{25},

856 y_{50}, y_{75}) one time for each of the two datasets.

857 We have publicly archived all relevant data, code, and materials on the Open Science Framework

858 (https://osf.io/mn5aj/). Archived data includes the original datasets distributed to all analysts, any

859 edited versions of the data analyzed by individual groups, and the data we analyzed with our meta-

860 analyses, which include the effect sizes derived from separate analyses, the statistics describing

861 variation in model structure among analyst groups, and the anonymized answers to our surveys of

862 analysts and peer reviewers. Similarly, we have archived both the analysis code used for each

863 individual analysis (where available) and the code from our meta-analyses. We have also archived

864 copies of our survey instruments from analysts and peer reviewers.

865 Our rules for excluding data from our study were as follows. We excluded from our synthesis any

866 individual analysis submitted after we had completed peer review or those unaccompanied by

867 analysis files that allow us to understand what the analysts did. We also excluded any individual

868 analysis that did not produce an outcome that could be interpreted as an answer to our primary

869 question (as posed above) for the respective dataset. For instance, this means that in the case of the

870 data on blue tit chick growth, we excluded any analysis that did not include something that can be

871 interpreted as growth or size as a dependent (response) variable, and in the case of the *Eucalyptus*

872 establishment data, we excluded any analysis that did not include a measure of grass cover among

873 the independent (predictor) variables. Also, as described above, any analysis that could not produce

874 an effect that could be converted to a signed Zr was excluded from analyses of Zr.

we recommend viewing this manuscript in html format at https://egouldo.github.io/ManyAnalysts/

Preregistration Deviation

Some analysts had difficulty implementing our instructions to derive the out-of-sample predictions, and in some cases (especially for the *Eucalyptus* data), they submitted predictions with implausibly extreme values. We believed these values were incorrect and thus made the conservative decision to exclude out-of-sample predictions where the estimates were > 3 standard deviations from the mean value from the full dataset.

we recommend viewing this manuscript in html format at https://egouldo.github.io/ManyAnalysts/

Additional Explanation

1. Evaluating model fit.

We evaluated all fitted models using the performance() function from the *performance* package [45] and the glance() function from the *broom.mixed* package [46]. For all models, we calculated the square root of the residual variance (Sigma) and the root mean squared error (RMSE). For GLMMs performance ()calculates the marginal and conditional R^2 values as well as the contribution of random effects (ICC), based on Nakagawa et al. [47]. The conditional R^2 accounts for both the fixed and random effects, while the marginal R^2 considers only the variance of the fixed effects. The contribution of random effects is obtained by subtracting the marginal R^2 from the conditional R^2.

2. Exploring outliers and analysis quality.

After seeing the forest plots of Zr values and noticing the existence of a small number of extreme outliers, especially from the Eucalyptus analyses, we wanted to understand the degree to which our heterogeneity estimates were influenced by these outliers. To explore this question, we removed the highest two and lowest two values of Zr in each dataset and re-calculated our heterogeneity estimates.

To help understand the possible role of the quality of analyses in driving the heterogeneity we observed among estimates of Zr, we recalculated our heterogeneity estimates after removing all effects from analysis teams that had received at least one rating of “deeply flawed and unpublishable” and then again after removing all effects from analysis teams with at least one rating of either “deeply flawed and unpublishable” or “publishable with major revisions”. We also used self-identified levels of statistical expertise to examine heterogeneity when we retained analyses only from analysis teams that contained at least one member who rated themselves as “highly proficient” or “expert” (rather than “novice” or “moderately proficient”) in conducting statistical analyses in their research area in our intake survey.

we recommend viewing this manuscript in html format at https://egouldo.github.io/ManyAnalysts/

Additional Explanation

3. Exploring possible impacts of lower quality estimates of degrees of freedom.

Our meta-analyses of variation in Zr required variance estimates derived from estimates of the degrees of freedom in original analyses from which Zr estimates were derived. While processing the estimates of degrees of freedom submitted by analysts, we identified a subset of these estimates in which we had lower confidence because two or more effects from the same analysis were submitted with identical degrees of freedom. We therefore conducted a second set of (more conservative) meta-analyses that excluded these Zr estimates with identical estimates of degrees of freedom and we present these analyses in the supplement.

Step 6: Facilitated Discussion and Collaborative Write-Up of Manuscript

We planned for analysts and initiating authors to discuss the limitations, results, and implications of the study and collaborate on writing the final manuscript for review as a stage-2 Registered Report.

Preregistration Deviation

As described above, due to the large number of recruited analysts and reviewers and the anticipated challenges of receiving and integrating feedback from so many authors, we limited analyst and reviewer participation in the production of the final manuscript to an invitation to call attention to serious problems with the manuscript draft.

Results

Summary Statistics

In total, 173 analyst teams, comprising 246 analysts, contributed 182 usable analyses of the two datasets examined in this study which yielded 215 effects. Analysts produced 135 distinct effects that met our criteria for inclusion in at least one of our meta-analyses for the blue tit dataset. Analysts

887 produced 81 distinct effects meeting our criteria for inclusion for the Eucalyptus dataset. Excluded
888 analyses and effects either did not answer our specified biological questions, were submitted with
889 insufficient information for inclusion in our meta-analyses, or were incompatible with production of
890 our effect size(s). We expected this final scenario (incompatible analyses), for instance we cannot
891 extract a Zr from random forest models, which is why we analyzed two distinct types of effects, Zr
892 and out-of-sample (y_i). Effects included in only a subset of our meta-analyses provided sufficient
893 information for inclusion in only that subset (see Table A.1). For both datasets, most submitted
894 analyses incorporated mixed effects. Submitted analyses of the blue tit dataset typically specified
895 normal error and analyses of the Eucalyptus dataset typically specified a non-normal error
896 distribution (Supplementary Table A.1).
897 For both datasets, the composition of models varied substantially in regards to the number of fixed
898 and random effects, interaction terms, and the number of data points used, and these patterns
899 differed somewhat between the blue tit and Eucalyptus analyses (See Supplementary Table A.2).
900 Focusing on the models included in the Zr analyses (because this is the larger sample), blue tit
901 models included a similar number of fixed effects on average (mean 5.2 ± 2.92 SD) as *Eucalyptus*
902 models (mean 5.01 ± 3.83 SD), but the standard deviation in number of fixed effects was somewhat
903 larger in the Eucalyptus models. The average number of interaction terms was much larger for the
904 blue tit models (mean 0.44 ± 1.11 SD) than for the Eucalyptus models (mean 0.16 ± 0.65 SD), but still
905 under 0.5 for both, indicating that most models did not contain interaction terms. Blue tit models
906 also contained more random effects (mean 3.53 ± 2.08 SD) than Eucalyptus models (mean 1.41 ±
907 1.09 SD). The maximum possible sample size in the blue tit dataset (3720 nestlings) was an order of
908 magnitude larger than the maximum possible in the Eucalyptus dataset (351 plots), and the means
909 and standard deviations of the sample size used to derive the effects eligible for our study were also
910 an order of magnitude greater for the blue tit dataset (mean 2622.07 ± 939.28 SD) relative to the
911 *Eucalyptus* models (mean 298.43 ± 106.25 SD). However, the standard deviation in sample size from
912 the *Eucalyptus* models was heavily influenced by a few cases of dramatic sub-setting (described

we recommend viewing this manuscript in html format at https://egouldo.github.io/ManyAnalysts/

913 below). Approximately three quarters of *Eucalyptus* models used sample sizes within 3% of the
914 maximum. In contrast, fewer than 20% of blue tit models relied on sample sizes within 3% of the
915 maximum, and approximately 50% of blue tit models relied on sample sizes 29% or more below the
916 maximum.
917 Analysts provided qualitative descriptions of the conclusions of their analyses. Each analysis team
918 provided one conclusion per dataset. These conclusions could take into account the results of any
919 formal analyses completed by the team as well as exploratory and visual analyses of the data. Here
920 we summarize all qualitative responses, regardless of whether we had sufficient information to use
921 the corresponding model results in our quantitative analyses below. We classified these conclusions
922 into the categories summarized below (Table 1):
923
924 Mixed: some evidence supporting a positive effect, some evidence supporting a negative effect
925 Conclusive negative: negative relationship described without caveat
926 Qualified negative: negative relationship but only in certain circumstances or where analysts express
927 uncertainty in their result
928 Conclusive none: analysts interpret the results as conclusive of no effect
929 None qualified: analysts describe finding no evidence of a relationship but they describe the
930 potential for an undetected effect
931 Qualified positive: positive relationship described but only in certain circumstances or where analysts
932 express uncertainty in their result
933 Conclusive positive: positive relationship described without caveat
934
935 For the blue tit dataset, most analysts concluded that there was negative relationship between
936 measures of sibling competition and nestling growth, though half the teams expressed qualifications
937 or described effects as mixed or absent. For the *Eucalyptus* dataset, there was a broader spread of
938 conclusions with at least one analyst team providing conclusions consistent with each conclusion

category. The most common conclusion for the *Eucalyptus* dataset was that there was no relationship between grass cover and *Eucalyptus* recruitment (either conclusive or qualified description of no relationship), but more than half the teams concluded that there were effects; negative, positive, or mixed.

Table 1: Tallies of analysts' qualitative answers to the research questions addressed by their analyses.

Dataset	Mixed	Negative Conclusive	Negative Qualified	None Conclusive	None Qualified	Positive Qualified	Positive Conclusive
blue tit	5	37	27	4	1	0	0
Eucalytpus	8	6	12	19	12	4	2

Distribution of Effects

Effect Size Zr

Although the majority (111 of 132) of the usable Zr effects from the blue tit dataset found nestling growth decreased with sibling competition, and the meta-analytic mean Zr (Fisher's transformation of the correlation coefficient) was convincingly negative (-0.35 ± 0.06 95% CI), there was substantial variability in the strength and the direction of this effect. Zr ranged approximately continuously from -0.93 to 0.19, (Figure 1a and Table 4) and of the 111 effects with negative slopes, 92 had confidence intervals excluding 0. Of the 20 with positive slopes indicating increased nestling growth in the presence of more siblings, 3 had confidence intervals excluding zero (Figure 1a).

Meta-analysis of the *Eucalyptus* dataset also showed substantial variability in the strength of effects as measured by Zr, and unlike with the blue tits, a notable lack of consistency in the direction of effects (Figure 1b, Table 4). Zr ranged from -4.47 (Supplementary Figure A.2), indicating a strong tendency for reduced *Eucalyptus* seedling success as grass cover increased, to 0.39, indicating the opposite. Although the range of reported effects skewed strongly negative, this was due to a small number of substantial outliers. Most values of Zr were relatively small with values < 0.2 and the

961 meta-analytic mean effect size was close to zero (-0.09 ± 0.12 95% CI). Of the 79 effects, fifty-three
962 had confidence intervals overlapping zero, approximately a quarter (fifteen) crossed the traditional
963 threshold of statistical significance indicating a negative relationship between grass cover and
964 seedling success, and eleven crossed the significance threshold indicating a positive relationship
965 between grass cover and seedling success (Figure 1b).
966

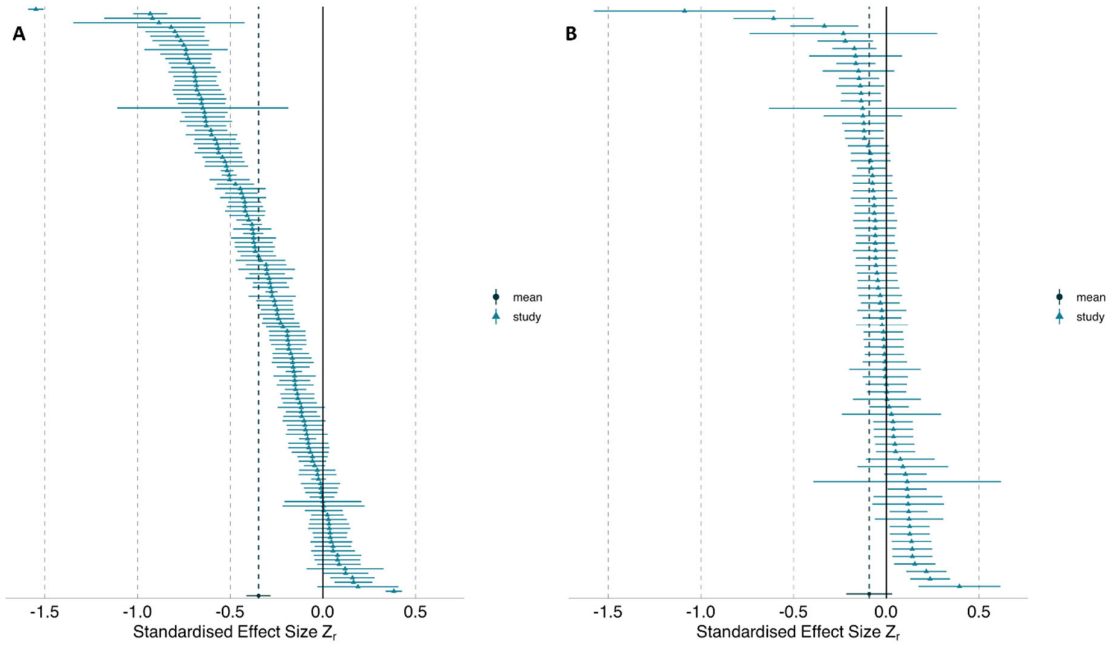

967

968 Figure 1: Forest plots of meta-analytic estimated standardized effect sizes (Zr) and their 95%
969 confidence intervals for each effect size included in the meta-analysis model for a) blue tit and b)
970 *Eucalytpus*. The meta-analytic mean effect size is noted in black and as a dashed vertical line, with
971 error bars also representing the 95% confidence interval. The solid black vertical line demarcates
972 effect size of 0, indicating no relationship between the test variable and the response variable. Note
973 that the Eucalyptus plot omits one extreme outlier with the value of -4.47 (Figure A.2) in order to
974 standardize the x-axes on these two panels.
975

we recommend viewing this manuscript in html format at https://egouldo.github.io/ManyAnalysts/

Out-of-sample predictions (y_i)

As with the effect size Zr, we observed substantial variability in the size of out-of-sample predictions derived from the analysts' models. Blue tit predictions (Figure 2a), which were z-score-standardised to accommodate the use of different response variables, always ranged far in excess of one standard deviation. In the y_{25} scenario, model predictions ranged from -1.85 to 0.42 (a range of 2.68 standard deviations), in the y_{50} scenario, they ranged from -0.53 to 1.11 (a range of 1.63 standard deviations), and in the y_{75} scenario they ranged from -0.03 to 1.58 (a range of 1.9 standard deviations). As should be expected given the existence of both negative and positive Zr values, all three out-of-sample scenarios produced both negative and positive predictions, although as with the Zr values, there is a clear trend for scenarios with more siblings to be associated with smaller nestlings. This is supported by the meta-analytic means of these three sets of predictions which were -0.66 (95% CI -0.82,–0.5) for the y_{25}, 0.34 (95% CI 0.2-0.48) for the y_{50}, and 0.67 (95% CI 0.57-0.77) for the y_{75}.

Eucalyptus out-of-sample predictions also varied substantially (Figure 2b), but because they were not z-score-standardised and are instead on the original count scale, the types of interpretations we can make differ. The predicted Eucalyptus seedling counts per 15 x 15 m plot for the y_{25} scenario ranged from 0.04 to 33.66, for the y_{50} scenario ranged from 0.03 to 13.02, and for the y_{75} scenario they ranged from 0.05 to 21.93. The meta-analytic mean predictions for these three scenarios were similar; 0.58 (95% CI, 0.21,-1.37) for the y_{25}, 0.92 (95% CI 0.36-1.65) for the y_{50}, and 1.67 (95% CI 0.8-2.83) for the y_{75} scenarios respectively.

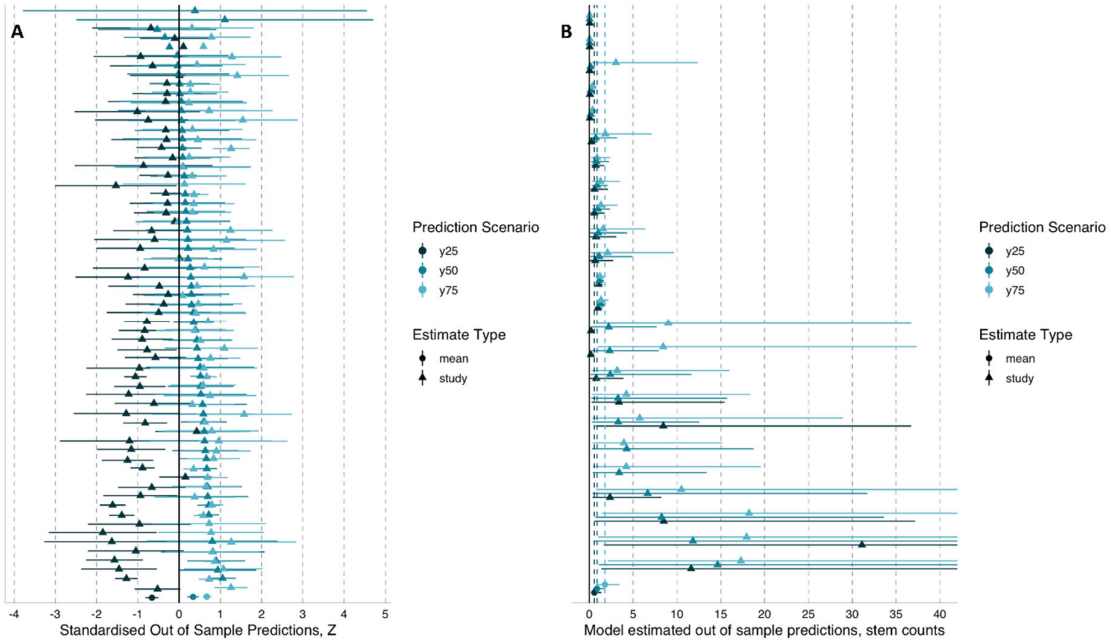

Figure 2: Forest plot of meta-analytic estimated standardized (z-score) blue tit out-of-sample predictions, y_i. for a) blue tit and b) Eucalyptus. Triangles represent individual estimates, circles represent the meta-analytic mean for each prediction scenario. Error bars are 95% confidence intervals.

Quantifying Heterogeneity

Effect Size (Zr)

We quantified both absolute (τ^2) and relative (I^2) heterogeneity resulting from analytical variation. Both measures suggest that substantial variability among effect sizes was attributable to the analytical decisions of analysts.

The total absolute level of variance beyond what would typically be expected due to sampling error, τ^2 (Table 2), among all usable blue tit effects was 0.088 and for Eucalyptus effects was 0.267. This is similar to or exceeding the median value (0.105) of τ^2 found across 31 recent meta-analyses (calculated from the data in 48]). The similarity of our observed values to values from meta-analyses of different studies based on different data suggest the potential for a large portion of heterogeneity to arise from analytical decisions. For further discussion of interpretation of τ^2 in our study, please consult discussion of post hoc analyses below.

we recommend viewing this manuscript in html format at https://egouldo.github.io/ManyAnalysts/

Table 2: Heterogeneity in the estimated effects Zr for meta-analyses of the full dataset, as well as from post hoc analyses including the dataset with outliers removed, the dataset excluding effects from analysis teams with at least one “unpublishable” rating, the dataset excluding effects from analysis teams with at least one “major revisions” rating or worse, or the dataset including only analyses from teams in which at least one analyst rated themselves as "highly proficient" or "expert" in statistical analysis. $\tau Team^2$ is the absolute heterogeneity for the random effect Team, $\tau EffectID^2$ is the absolute heterogeneity for the random effect *EffectID*, nested under Team, and $\tau total^2$ is the total absolute heterogeneity. I^2Total is the proportional heterogeneity; the proportion of the variance among effects not attributable to sampling error, I^2Team is the subset of the proportional heterogeneity due to differences among Teams and I^2Team, *EffectID* is subset of the proportional heterogeneity attributable to among-*EffectID* differences.

Dataset	τ^2Total	τ^2Team	τ^2EffectID	I^2Total	I^2Team	I^2Team,EffectID	N. Obs
All Analyses							
blue tit	0.09	0.04	0.05	97.732%	40.11%	57.63%	131
Eucalyptus	0.27	0.02	0.25	98.589%	6.88%	91.71%	79
All analyses, outliers Removed							
blue tit	0.07	0.05	0.02	97.030%	66.90%	30.13%	127
Eucalyptus	0.01	0.00	0.01	66.193%	19.27%	46.93%	75
Analyses receiving at least one 'Unpublishable' rating removed							
blue tit	0.08	0.03	0.05	97.601%	38.10%	59.50%	109
Eucalyptus	0.01	0.01	0.01	79.741%	28.32%	51.42%	55
Analyses receiving at least one 'Unpublishable' and or 'Major Revisions' rating removed							
blue tit	0.14	0.01	0.13	98.718%	5.17%	93.55%	32
Eucalyptus	0.03	0.03	0.00	88.915%	88.91%	0.00%	13
Analyses from teams that include highly proficient or expert data analysts							
blue tit	0.10	0.04	0.06	98.058%	36.27%	61.78%	89
Eucalyptus	0.58	0.02	0.56	99.412%	3.49%	95.93%	34

In our analyses, I^2 is a plausible index of how much more variability among effect sizes we have observed, as a proportion, than we would have observed if sampling error were driving variability. We discuss our interpretation of I^2 further in the methods, but in short, it is a useful metric for comparison to values from published meta-analyses and provides a plausible value for how much heterogeneity could arise in a normal meta-analysis with similar sample sizes due to analytical

variability alone. In our study, total I^2 for the blue tit Zr estimates was extremely large, at 97.73%, as was the Eucalyptus estimate (98.59% Table 2).

Although the overall I^2 values were similar for both Eucalyptus and blue tit analyses, the relative composition of that heterogeneity differed. For both datasets, the majority of heterogeneity in Zr was driven by differences among effects as opposed to differences among teams, though this was more prominent for the Eucalyptus dataset, where nearly all of the total heterogeneity was driven by differences among effects (91.71%) as opposed to differences among teams (6.88%) (Table 2).

Out-of-sample predictions (y_i)

We observed substantial heterogeneity among out-of-sample estimates, but the pattern differed somewhat from the Zr values (Table 3). Among the blue tit predictions, I^2 ranged from medium-high for the y_{25} scenario (68.36) to low (27.02) for the y_{75} scenario. Among the Eucalyptus predictions, I^2 values were uniformly high (>82%). For both datasets, most of the existing heterogeneity among predicted values was attributable to among-team differences, with the exception of the y_{50} analysis of the Eucalyptus dataset. We are limited in our interpretation of τ^2 for these estimates because, unlike for the Zr estimates, we have no benchmark for comparison with other meta-analyses.

Table 3: Heterogeneity among the out-of-sample predictions y_i for both blue tit and *Eucalyptus* datasets. $\tau Team^2$ is the absolute heterogeneity for the random effect *Team*, $\tau EffectID^2$ is the absolute heterogeneity for the random effect *EffectID*, nested under *Team*, and $\tau total^2$ is the total absolute heterogeneity. I^2Total is the proportional heterogeneity; the proportion of the variance among effects not attributable to sampling error, I^2Team is the subset of the proportional heterogeneity due to differences among Teams and I^2Team, *EffectID* is subset of the proportional heterogeneity attributable to among-*EffectID* differences.

Dataset	Scenario	N. Obs	T^2Total	τ^2Team	τ^2EffectID	I^2Total	I^2Team	I^2Team,EffectID
blue tit	y_{25}	62	0.14	0.11	0.03	68.36%	51.82%	16.54%
	y_{50}	59	0.07	0.06	0.01	50.37%	45.66%	4.71%
	y_{75}	62	0.02	0.02	0.00	27.02%	25.57%	1.45%

Eucalyptus	y_{25}	22	3.05	1.95	1.10	88.76%	56.76%	32.00%
	y_{50}	24	1.61	0.53	1.08	83.26%	27.52%	55.73%
	y_{75}	24	1.69	1.41	0.28	79.76%	66.52%	13.25%

Post-hoc Analysis: Exploring outlier characteristics and the effect of outlier removal on heterogeneity

Effect Sizes (Zr)

The outlier *Eucalyptus* Zr values were striking and merited special examination. The three negative outliers had very low sample sizes were based on either small subsets of the dataset or, in one case, extreme aggregation of data. The outliers associated with small subsets had sample sizes (n= 117, 90) that were less than half of the total possible sample size of 351. The case of extreme aggregation involved averaging all values within each of the 18 sites in the dataset.

Surprisingly, both the largest and smallest effect sizes in the blue tit analyses (Figure 1a) come from the same analyst (anonymous ID: Adelong), with identical models in terms of the explanatory variable structure, but with different response variables. However, the radical change in effect was primarily due to collinearity with covariates. The primary predictor variable (brood count after manipulation) was accompanied by several collinear variables, including the highly collinear (correlation of approximately 0.9 (Supplementary Figure D.2)) covariate (brood size at day 14) in both analyses. In the analysis of nestling weight, brood count after manipulation showed a strong positive partial correlation with weight after controlling for brood count at day 14 and treatment category (increased, decreased, unmanipulated). In that same analysis, the most collinear covariate (the day 14 count) had a negative partial correlation with weight. In the analysis with tarsus length as the response variable, these partial correlations were almost identical in absolute magnitude, but reversed in sign and so brood count after manipulation was now the collinear predictor with the negative relationship. The two models were therefore very similar, but the two collinear predictors simply switched roles, presumably because a subtle difference in the distribution of weight and tarsus length data.

When we dropped the *Eucalyptus* outliers, I^2 decreased from high (98.59%), using Higgins' [36] suggested benchmark, to between moderate and high (66.19%, Table 2). However, more notably, τ^2 dropped from 0.27 to 0.01, indicating that, once outliers were excluded, the observed variation in effects was similar to what we would expect if sampling error were driving the differences among effects (since τ^2 is the variance in addition to that driven by sampling error). The interpretation of this value of τ^2 in the context of our many-analyst study is somewhat different than a typical meta-analysis, however, since in our study (especially for *Eucalyptus*, where most analyses used almost exactly the same data points), there is almost no role for sampling error in driving the observed differences among the estimates. Thus, rather than concluding that the variability we observed among estimates (after removing outliers) was due only to sampling error (because τ^2 became small: 10% of the median from 48), we instead conclude that the observed variability, which must be due to the divergent choices of analysts rather than sampling error, is approximately of the same magnitude as what we would have expected if, instead, sampling error, and not analytical heterogeneity, were at work. Presumably, if sampling error had actually also been at work, it would have acted as an additional source of variability and would have led total variability among estimates to be higher. With total variability higher and thus greater than expected due to sampling error alone, τ^2 would have been noticeably larger. Conversely, dropping outliers from the set of blue tit effects did not meaningfully reduce I^2, and only modestly reduced τ^2 (Table 2). Thus, effects at the extremes of the distribution were much stronger contributors to total heterogeneity for effects from analyses of the Eucalyptus than for the blue tit dataset.

Table 4: Estimated mean value of the standardised correlation coefficient, Zr, along with its standard error and 95% confidence intervals. We re-computed the meta-analysis for different post-hoc subsets of the data: All eligible effects, removal of effects from analysis teams that received at least one peer rating of 'deeply flawed and unpublishable', removal of any effects from analysis teams that received at least one peer rating of either 'deeply flawed and unpublishable' or 'publishable with major

we recommend viewing this manuscript in html format at https://egouldo.github.io/ManyAnalysts/

revisions' , inclusion of only effects from analysis teams that included at least one member who rated themselves as "highly proficient" or "expert" at conducting statistical analyses in their research area..

Dataset	$\widehat{\mu}$	**SE**[$\widehat{\mu}$]	**95% CI**	**statistic**	**p-value**
All Analyses					
blue tit	−0.35	0.03	[−0.41,−0.28]	−10.49	<0.001
Eucalyptus	−0.09	0.06	[−0.22,0.03]	−1.47	0.14
Analyses receiving at least one 'Unpublishable' rating removed					
blue tit	−0.36	0.03	[−0.43,−0.29]	−10.49	<0.001
Eucalyptus	−0.02	0.02	[−0.07,0.02]	−1.15	0.3
Analyses receiving at least one 'Unpublishable' and or 'Major Revisions' rating removed					
blue tit	−0.37	0.07	[−0.51,−0.23]	−5.34	<0.001
Eucalyptus	−0.04	0.05	[−0.15,0.07]	−0.77	0.4
All analyses - outliers removed					
blue tit	−0.35	0.03	[−0.42,−0.29]	−10.95	<0.001
Eucalyptus	−0.03	0.01	[−0.06,0.00]	−2.23	0.026
Analyses from teams with highly proficient or expert data analysts					
blue tit	−0.35	0.04	[−0.44,−0.27]	−8.31	<0.001
Eucalyptus	−0.17	0.13	[−0.43,0.10]	−1.24	0.2

Out-of-sample predictions (y_i)

We did not conduct these post hoc analyses on the out-of-sample predictions as the number of eligible effects was smaller and the pattern of outliers differed.

Post-hoc analysis: Exploring the effect of removing analyses with poor peer ratings on heterogeneity

Effect Size (Zr)

Removing poorly rated analyses had limited impact on the meta-analytic means (Supplementary Figure B.3). For the Eucalyptus dataset, the meta-analytic mean shifted from -0.09 to -0.02 when effects from analyses rated as unpublishable were removed, and to -0.04 when effects from analyses rated, at least once, as unpublishable or requiring major revisions were removed. Further, the confidence intervals for all of these means overlapped each of the other means (Table 4). We saw similar patterns for the blue tit dataset, with only small shifts in the meta-analytic mean, and confidence intervals of all three means overlapping each other mean (Table 4). Refitting the meta-

analysis with a fixed effect for categorical ratings also showed no indication of differences in group meta-analytic means due to peer ratings (Supplementary Figure B.1).

For the blue tit dataset, removing poorly-rated analyses led to only negligible changes in I^2Total and relatively minor impacts on τ^2. However, for the Eucalyptus dataset, removing poorly-rated analyses led to notable reductions in I^2Total and substantial reductions in τ^2. When including all analyses, the *Eucalyptus* I^2Total was 98.59% and τ^2 was 0.27, but eliminating analyses with ratings of "unpublishable" reduced I^2Total to 79.74% and τ^2 to 0.01, and removing also those analyses "needing major revisions" left I^2Total at 88.91% and τ^2 at 0.03 (Table 2). Additionally, the allocations of I^2 to the team versus individual effect were altered for both blue tit and *Eucalyptus* meta-analyses by removing poorly rated analyses, but in different ways. For blue tit meta-analysis, between a third and two-thirds of the total I^2 was attributable to among-team variance in most analyses until both analyses rated "unpublishable" and analyses rated in need of "major revision" were eliminated, in which case almost all remaining heterogeneity was attributable to among-effect differences. In contrast, for Eucalyptus meta-analysis, the among-team component of I^2 was less than third until both analyses rated "unpublishable" and analyses rated in need of "major revision" were eliminated, in which case almost 90% of heterogeneity was attributable to differences among teams.

Out-of-sample predictions (y_i)

We did not conduct these post hoc analyses on the out-of-sample predictions as the number of eligible effects was smaller and our ability to interpret heterogeneity values for these analyses was limited.

we recommend viewing this manuscript in html format at https://egouldo.github.io/ManyAnalysts/

Post-hoc analysis: Exploring the effect of including only analyses conducted by analysis teams with at least one member self-rated as "highly proficient" or "expert" in conducting statistical analyses in their research area

Effect sizes (Zr)

Including only analyses conducted by teams that contained at least one member who rated themselves as "highly proficient" or "expert" in conducting the relevant statistical methods had negligible impacts on the meta-analytic means (Table 4), the distribution of Zr effects (Supplementary Figure B.4), or heterogeneity estimates (Table 2), which remained extremely high.

Out-of- sample predictions (y_i)

We did not conduct these post hoc analyses on the out-of-sample predictions as the number of eligible effects was smaller.

Post-hoc analysis: Exploring the effect of excluding estimates of Zr in which we had reduced confidence

As described in our addendum to the methods, we identified a subset of estimates of Zr in which we had less confidence because of features of the submitted degrees of freedom. Excluding these effects in which we had lower confidence had minimal impact on the meta-analytic mean and the estimates of total I^2 and τ^2 for both blue tit and Eucalyptus meta-analyses, regardless of whether outliers were also excluded (Supplementary Table B.1).

Explaining Variation in Deviation Scores

None of the pre-registered predictors explained substantial variation in deviation among submitted statistical effects from the meta-analytic mean (Table 5, Table 6). Note that the extremely high $R^2_{Conditional}$ values from the analyses of continuous peer ratings as predictors of deviation scores are a function of the random effects, not the fixed effect of interest. These high values of $R^2_{Conditional}$ result from the fact that each effect size was included in the analysis multiple times, to allow comparison

we recommend viewing this manuscript in html format at https://egouldo.github.io/ManyAnalysts/

with ratings from the multiple peer reviewers who reviewed each analysis, and therefore when we included Effect ID as a random effect, the observations within each random effect category were identical.

Table 5: Summary metrics for registered models seeking to explain deviation (Box-Cox transformed absolute deviation scores) from the mean Zr as a function of Sorensen's Index, categorical peer ratings, and continuous peer ratings for blue tit and Eucalyptus analyses, and as a function of the presence or absence of random effects (in the analyst's models) for Eucalyptus analyses. We report coefficient of determination, R^2, for our models including only fixed effects as predictors of deviation, and we report $R^2_{Conditional}$, $R^2_{Marginal}$ and the intra-class correlation (ICC) from our models that included both fixed and random effects. For all our models, we calculated the residual standard deviation σ and root mean squared error (RMSE).

Dataset	**R^2**	**$R^2_{Conditional}$**	**$R^2_{Marginal}$**	**ICC**	**σ**	**RMSE**	**N. Obs.**
Deviation explained by categorical ratings							
blue tit		0.0903	0.0067	0.0842	6.52e-01	6.32e-01	473
Eucalyptus		0.1319	0.0124	0.1209	1.06e+00	1.02e+00	346
Deviation explained by continuous ratings							
blue tit		1.0000	2.00e-26	1.0000	1.63e-05	1.56e-12	473
Eucalyptus		0.9998	6.57e-30	0.9998	7.93e-03	7.09e-14	346
Deviation explained by Sorensen's index							
blue tit	0.0011				0.681	0.676	124
Eucalyptus	0.0005				1.14	1.120	72
Deviation explained by inclusion of random effects							
blue tit	0.0268				0.658	0.653	131
Eucalyptus	8.67e-08				1.12	1.100	79

Table 6: Parameter estimates from models of Box-Cox transformed deviation scores as a function of continuous and categorical peer ratings, Sorensen scores, and the inclusion of random effects. Standard Errors (SE), 95% confidence intervals (95%CI) are reported for all estimates, while t values, degrees of freedom and p-values are presented for fixed-effects. Note that positive parameter estimates mean that as the predictor variable increases, so does the absolute value of the deviation from the meta-analytic mean.

Dataset	**Parameter**	**Effect**	**Coeff.**	**SE**	**95% CI**	**t**	**df**	**p-value**
Deviation explained by inclusion of random effects								

Eucalyptus	Intercept		2.53	0.27	-3.06,-1.99	-9.31	77	<0.001
	Random effects		0.00	0.31	-0.60, 0.60	0.00	77	>0.9
Deviation explained by mean Sorensen's index								
Eucalyptus	Intercept		-2.75	1.07	-4.85,-0.65	-2.57	70	0.010
	Sorensen Index		0.29	1.54	-2.74, 3.32	0.19	70	0.9
blue tit	Intercept		-1.56	0.38	-2.30,-0.82	-4.12	122	<0.001
	Mean Sorensen Index		0.23	0.63	-1.00, 1.46	0.37	122	0.7
Deviation explained by continuous ratings								
Eucalyptus	Intercept	Fixed	-2.52	0.06	-2.63,-2.40	-42.58	342	<0.001
	Continuous Rating	Fixed	6e-17	2e-10	-4e-10, 4e-10	-3e-07	342	>0.9
	SD (Intercept)	Random (EffectID)	0.53	0.04	0.45, 0.62			
	SD (Observations)	Random (Residual)	0.01	3e-04	0.01,0.01			
blue tit	Intercept	Fixed	-1.41	0.03	-1.47,-1.35	-46.54	469	<0.001
	Continuous Rating	Fixed	-3e-15	1e-09	-2e-09,2e-09	-2e-06	469	>0.9
	SD (Intercept)	Random (EffectID)	0.34	0.02	0.30, 0.39			
	SD (Observations)	Random (Residual)	2e-05	6e-07	2e-05,2e-05			
Deviation explained by categorical ratings								
Eucalyptus	Intercept	Fixed	-2.66	0.27	-3.18,-2.13	-9.97	340	<0.001
	Publishable with major revisions	Fixed	0.29	0.29	-0.27, 0.85	1.02	340	0.3
	Publishable with minor revisions	Fixed	0.01	0.28	-0.54, 0.56	0.04	340	>0.9
	Publishable as is	Fixed	0.05	0.31	-0.55, 0.66	0.17	340	0.9
	SD (Intercept)	Random (ReviewerID)	0.39	0.09	0.25, 0.61			
	SD (Observations	Random (Residual)	1.06	0.04	0.98,1.15			
blue tit	Intercept	Fixed	-1.21	0.15	-1.50,-0.93	-8.29	467	<0.001
	Publishable with major revisions	Fixed	-0.23	0.15	-0.53, 0.07	-1.50	467	0.13
	Publishable with minor revisions	Fixed	-0.23	0.15	-0.53, 0.07	-1.52	467	0.13
	Publishable as is	Fixed	-0.15	0.17	-0.48, 0.18	-0.89	467	0.4
	SD (Intercept)	Random (ReviewerID)	0.20	0.05	0.13, 0.31			
	SD (Observations	Random (Residual)	0.65	0.02	0.61,0.7			

Deviation Scores as explained by reviewer ratings

Effect Sizes (Zr)

We obtained reviews from 128 reviewers who reviewed analyses for a mean of 3.27 (range 1 - 11) analysis teams. Analyses of the blue tit dataset received a total of 240 reviews, each was reviewed by

a mean of 3.87 (SD 0.71, range 3-5) reviewers. Analyses of the Eucalyptus dataset received a total of 178 reviews, each was reviewed by a mean of 4.24 (SD 0.79, range 3-6) reviewers. We tested for inter-rater reliability to examine how similarly reviewers reviewed each analysis and found approximately no agreement among reviewers. When considering continuous ratings, IRR was 0.01, and for categorical ratings, IRR was -0.14.

Many of the models of deviance as a function of peer ratings faced issues of failure to converge or singularity due to sparse design matrices with our pre-registered random effects (*EffectID* and *ReviewerID*) (see Supplementary Table C.1). These issues persisted after increasing the tolerance and changing the optimizer. For both *Eucalyptus* and blue tit datasets, models with continuous ratings as a predictor were singular when both pre-registered random effects were included.

When using only categorical ratings as predictors, models converged only when specifying reviewer ID as a random effect. That model had a R^2_C of 0.09 and a R^2_M of 0.01. The model using the continuous ratings converged for both random effects (in isolation), but not both. We present results for the model using study ID as a random effect because we expected it would be a more important driver of variation in deviation scores. That model had a R^2_C of 1 and a R^2_M of 0.01 for the blue tit dataset and a R^2_C of 1 and a R^2_M of 0.01 for the *Eucalyptus* dataset. Neither continuous or categorical reviewer ratings of the analyses meaningfully predicted deviance from the meta-analytic mean (Table 6, Figure 3). We re-ran the multi-level meta-analysis with a fixed-effect for the categorical publishability ratings and found no difference in mean standardised effect sizes among publishability ratings (Supplementary Figure B.1).

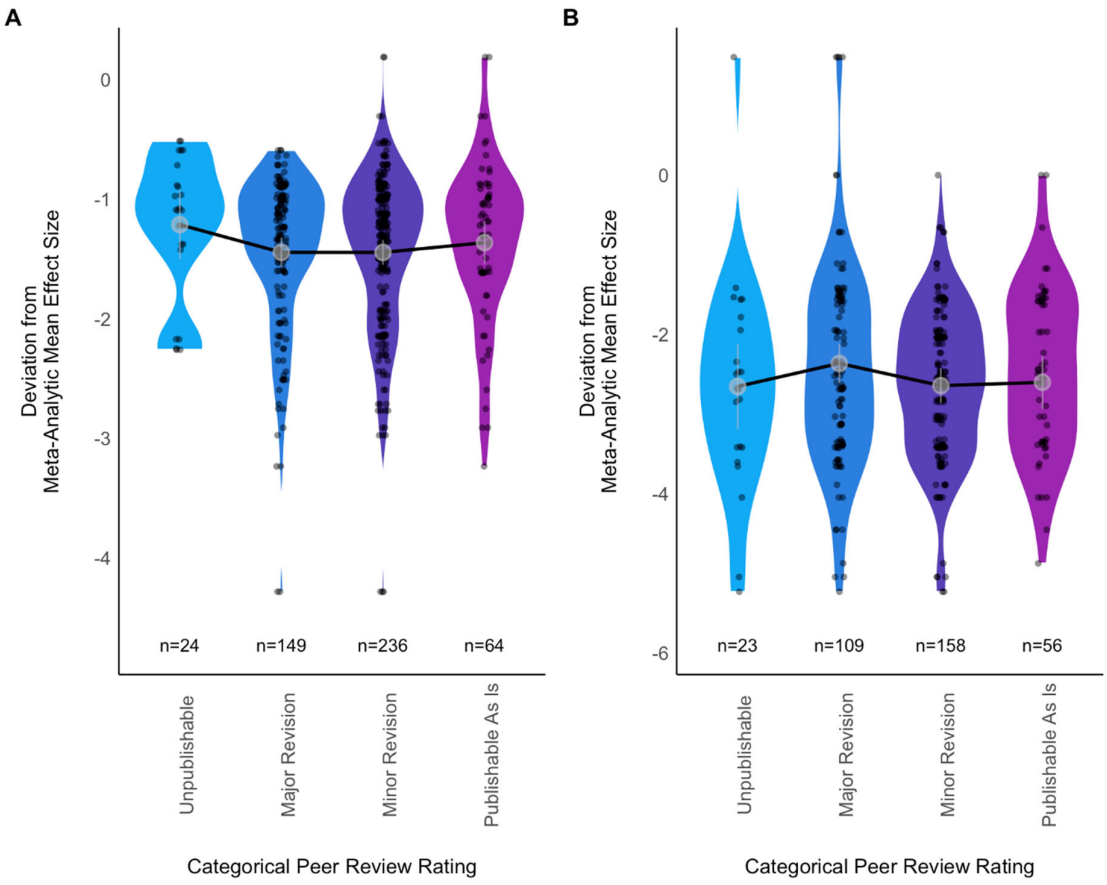

Figure 3: Violin plot of Box-Cox transformed deviation from meta-analytic mean as a function of categorical peer rating for a) blue tit and b) *Eucalyptus*. Grey points for each rating group denote model-estimated marginal mean deviation, and error bars denote 95% CI of the estimate.

Out-of-sample predictions (y_i)

Some models of the influence of reviewer ratings on out-of-sample predictions (y_i) had issues with convergence and singularity of fit (see Supplementary Table C.2) and those models that converged and were not singular showed no strong relationship (Supplementary Figures C.2, Figure C.3), as with the Zr analyses.

we recommend viewing this manuscript in html format at https://egouldo.github.io/ManyAnalysts/

Deviation scores as explained by the distinctiveness of variables in each analysis

Effect Size (Zr)

We employed Sorensen's index to calculate the distinctiveness of the set of predictor variables used in each model (Figure 5). The mean Sorensen's score for blue tit analyses was 0.69 (range 0.55-0.98), and for Eucalyptus analyses was 0.59 (range 0.43-0.86).

We found no meaningful relationship between distinctiveness of variables selected and deviation from the meta-analytic mean (Table 6, Figure 5) for either blue tit (mean 0.23, 95% CI -1,1.46) or Eucalyptus effects (mean 0.29, 95% CI -2.74,3.32).

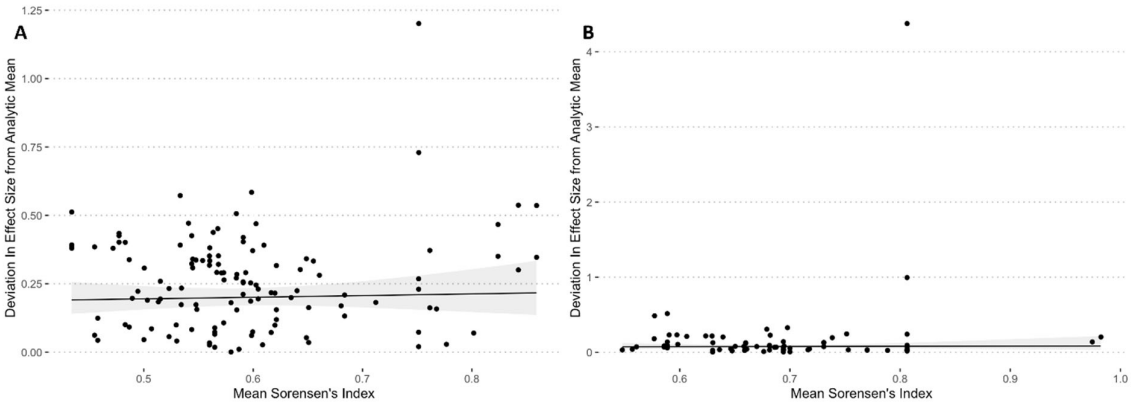

Figure 4: Fitted model of the Box-Cox-transformed deviation score (deviation in effect size from meta-analytic mean) as a function of the mean Sorensen's index showing distinctiveness of the set of predictor variables for a) blue tit, and b) *Eucalyptus*. Grey ribbons on predicted values are 95% CI's.

Out-of-sample predictions

As with the Zr estimates, we did not observe any convincing relationships between deviation scores of out-of-sample predictions and Sorensen's index values. Please see Supplementary Material C.4.2.

we recommend viewing this manuscript in html format at https://egouldo.github.io/ManyAnalysts/

Deviation scores as explained by the inclusion of random effects

Effect Size (Zr)

There were only three blue tit analyses that did not include random effects, which is below the pre-registered threshold for fitting a model of the Box-Cox transformed deviation from the meta-analytic mean as a function of whether the analysis included random-effects. However, 17 Eucalyptus analyses included only fixed effects, which crossed our pre-registered threshold. Consequently, we performed this analysis for the Eucalyptus dataset only. There was no relationship between random-effect inclusion and deviation from meta-analytic mean among the Eucalyptus analyses (Table 6, Figure 5).

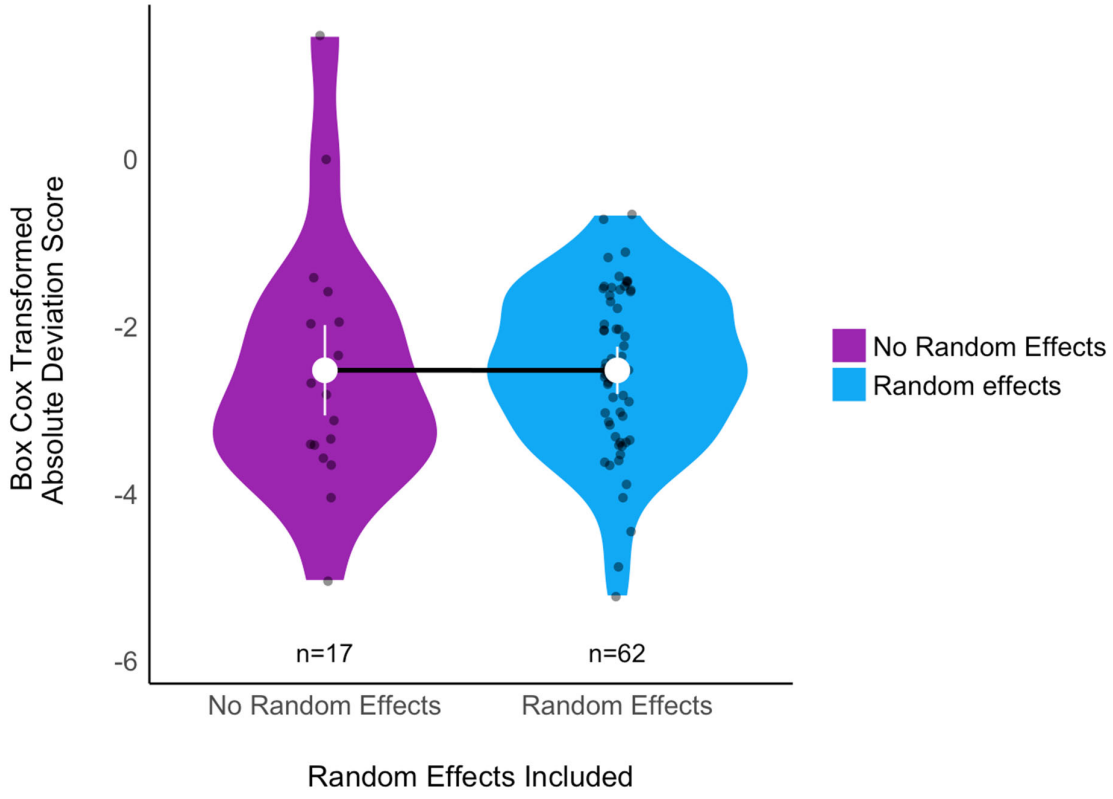

Figure 5: Violin plot of mean Box-Cox transformed deviation from meta-analytic mean as a function of random-effects inclusion in *Eucalyptus* analyses. '1' indicates random-effects were included in analyst's model, while 0 indicates no random-effects were included. White points for each group of

analyses denote model-estimated marginal mean deviation, and error bars denote 95% CI of the estimate.

Out-of-sample predictions

As with the Zr estimates, we did not examine the possibility of a relationship between the inclusion of random effects and the deviation scores of the blue tit out-of-sample predictions. When we examined the possibility of this relationship for the Eucalyptus effects, we found consistent evidence of somewhat higher Box-Cox-transformed deviation values for models including a random effect, meaning the models including random effects averaged slightly higher deviation from the meta-analytic means (Supplementary Figure C.5).

Multivariate Analysis Effect size (Zr) and out-of-sample predictions (y_i)

Like the univariate models, the multivariate models did a poor job of explaining deviations from the meta-analytic mean. Because we pre-registered a multivariate model that contained collinear predictors that produce results which are not readily interpretable, we present these models in the supplement. We also had difficulty with convergence and singularity for multivariate models of out-of-sample (y_i) result, and had to adjust which random effects we included (Supplementary Table C.7). However, no multivariate analyses of Eucalyptus out-of-sample results avoided problems of convergence or singularity, no matter which random effects we included (Supplementary Table C.7). We therefore present no multivariate Eucalyptus y_i models. We present parameter estimates from multivariate Zr models for both datasets (Supplementary Tables C.5, C.6) and from y_i models from the blue tit dataset (Supplementary Tables C.8, C.9). We include interpretation of the results from these models in the supplement, but the results do not change the interpretations we present above based on the univariate analyses.

Discussion

When a large pool of ecologists and evolutionary biologists analyzed the same two datasets to answer the corresponding two research questions, they produced substantially heterogeneous sets

of answers. Although the variability in analytical outcomes was high for both datasets, the patterns of this variability differed distinctly between them. For the blue tit dataset, there was nearly continuous variability across a wide range of Zr values. In contrast, for the Eucalyptus dataset, there was less variability across most of the range, but more striking outliers at the tails. Among out-of-sample predictions, there was again almost continuous variation across a wide range (2 SD) among blue tit estimates. For Eucalyptus, out-of-sample predictions were also notably variable, with about half the predicted stem count values at <2 but the other half being much larger, and ranging to nearly 40 stems per 15 m x 15 m plot. We investigated several hypotheses for drivers of this variability within datasets, but found little support for any of these. Most notably, even when we excluded analyses that had received one or more poor peer reviews, the heterogeneity in results largely persisted. Regardless of what drives the variability, the existence of such dramatically heterogeneous results when ecologists and evolutionary biologists seek to answer the same questions with the same data should trigger conversations about how ecologists and evolutionary biologists analyze data and interpret the results of their own analyses and those of others in the literature [e.g., 11, 20, 49, 50].

Our observation of substantial heterogeneity due to analytical decisions is consistent with a growing body of work, much of it from the quantitative social sciences [e.g., 11, 17–21]. In all of these studies, when volunteers from the discipline analyzed the same data, they produced a worryingly diverse set of answers to a pre-set question. This diversity always included a wide range of effect sizes, and in most cases, even involved effects in opposite directions. Thus, our result should not be viewed as an anomalous outcome from two particular datasets, but instead as evidence from additional disciplines regarding the heterogeneity that can emerge from analyses of complex datasets to answer questions in probabilistic science. Not only is our major observation consistent with other studies, it is, itself, robust because it derived primarily from simple forest plots that we produced based on a small set of decisions that were mostly registered before data gathering and which conform to widely accepted meta-analytic practices.

Unlike the strong pattern we observed in the forest plots, our other analyses, both registered and post hoc, produced either inconsistent patterns, weak patterns, or the absence of patterns. Our registered analyses found that deviations from the meta-analytic mean by individual effect sizes (Zr) or the predicted values of the dependent variable (y_i) were poorly explained by our hypothesized predictors: peer rating of each analysis team's method section, a measurement of the distinctiveness of the set of predictor variables included in each analysis, or whether the model included random effects. However, in our post hoc analyses, we found that dropping analyses identified as unpublishable or in need of major revision by at least one reviewer modestly reduced the observed heterogeneity among the Zr outcomes, but only for Eucalyptus analyses, apparently because this led to the dropping of the major outlier. This limited role for peer review in explaining the variability in our results should be interpreted cautiously because the inter-rater reliability among peer reviewers was extremely low, and at least some analyses that appeared flawed to us were not marked as flawed by reviewers. However, the hypothesis that poor quality analyses drove the heterogeneity we observed was also contradicted by our observation that analysts' self-declared statistical expertise appeared unrelated to heterogeneity. When we retained only analyses from teams including at least one member with high self-declared levels of expertise, heterogeneity among effect sizes remained high. Thus, our results suggest lack of statistical expertise is not the primary factor responsible for the heterogeneity we observed, although further work is merited before rejecting a role for statistical expertise. Not surprisingly, simply dropping outlier values of Zr for *Eucalyptus* analyses, which had more extreme outliers, led to less observable heterogeneity in the forest plots, and also reductions in our quantitative measures of heterogeneity. We did not observe a similar effect in the blue tit dataset because that dataset had outliers that were much less extreme and instead had more variability across the core of the distribution.

Our major observations raise two broad questions; why was the variability among results so high, and why did the pattern of variability differ between our two datasets. One important and plausible answer to the first question is that much of the heterogeneity derives from the lack of a precise

relationship between the two biological research questions we posed and the data we provided. This lack of a precise relationship between data and question creates many opportunities for different model specifications, and so may inevitably lead to varied analytical outcomes [50]. However, we believe that the research questions we posed are consistent with the kinds of research question that ecologists and evolutionary biologists typically work from. When designing the two biological research questions, we deliberately sought to represent the level of specificity we typically see in these disciplines. This level of specificity is evident when we look at the research questions posed by some recent meta-analyses in these fields:

"how [does] urbanisation impact mean phenotypic values and phenotypic variation … [in] paired urban and non-urban comparisons of avian life-history traits" [51]

"[what are] the effects of ocean acidification on the crustacean exoskeleton, assessing both exoskeletal ion content (calcium and magnesium) and functional properties (biomechanical resistance and cuticle thickness)" [52]

"[what is] the extent to which restoration affects both the mean and variability of biodiversity outcomes … [in] terrestrial restoration" [53]

"[does] drought stress [have] a negative, positive, or null effect on aphid fitness" [54]

"[what is] the influence of nitrogen-fixing trees on soil nitrous oxide emissions" [55]

There is not a single precise answer to any of these questions, nor to the questions we posed to analysts in our study. And this lack of single clear answers will obviously continue to cause uncertainty since ecologists and evolutionary biologists conceive of the different answers from the different statistical models as all being answers to the same general question. A possible response would be a call to avoid these general questions in favor of much more precise alternatives [50]. However, the research community rewards researchers who pose broad questions [56], and so researchers are unlikely to narrow their scope without a change in incentives. Further, we suspect that even if individual studies specified narrow research questions, other scientists would group

1348 these more narrow questions into broader categories, for instance in meta-analyses, because it is
1349 these broader and more general questions that often interest the research community.
1350 Although variability in statistical outcomes among analysts may be inevitable, our results raise
1351 questions about why this variability differed between our two datasets. We are particularly
1352 interested in the differences in the distribution of Zr since the distributions of out-of-sample
1353 predictions were on different scales for the two datasets, thus limiting the value of comparisons. The
1354 forest plots of Zr from our two datasets showed distinct patterns, and these differences are
1355 consistent with several alternative hypotheses. The results submitted by analysts of the Eucalyptus
1356 dataset showed a small average (close to zero) with most estimates also close to zero (± 0.2), though
1357 about a third far enough above or below zero to cross the traditional threshold of statistical
1358 significance. There were a small number of striking outliers that were very far from zero. In contrast,
1359 the results submitted by analysts of the blue tit dataset showed an average much further from zero (-
1360 0.35) and a much greater spread in the core distribution of estimates across the range of Zr values (±
1361 0.5 from the mean), with few modest outliers. So, why was there more spread in effect sizes (across
1362 the estimates that are not outliers) in the blue tit analyses relative to the Eucalyptus analyses?
1363 One possible explanation for the lower heterogeneity among most Eucalyptus Zr effects is that weak
1364 relationships may limit the opportunities for heterogeneity in analytical outcome. Some evidence for
1365 this idea comes from two sets of “many labs” studies in psychology [4, 57]. In these studies, many
1366 independent lab groups each replicated a large set of studies, including, for each study, the
1367 experiment, data collection, and statistical analyses. These studies showed that, when the meta-
1368 analytic mean across the replications from different labs was small, there was much less
1369 heterogeneity among the outcomes than when the mean effect sizes were large [4, 57]. Of course, a
1370 weak average effect size would not prevent divergent effects in all circumstances. As we saw with the
1371 Eucalyptus analyses, taking a radically smaller subset of the data can lead to dramatically divergent
1372 effect sizes even when the mean with the full dataset is close to zero.

Our observation that dramatic sub-setting in the Eucalyptus dataset was associated with correspondingly dramatic divergence in effect sizes leads us towards another hypothesis to explain the differences in heterogeneity between the Eucalyptus and blue tit analysis sets. It may be that when analysts often divide a dataset into subsets, the result will be greater heterogeneity in analytical outcome for that dataset. Although we saw sub-setting associated with dramatic outliers in the Eucalyptus dataset, nearly all other analyses of Eucalyptus data used very close to the same set of 351 samples, and as we saw, these effects did not vary substantially. However, analysts often analyzed only a subset of the blue tit data, and as we observed, sample sizes were much more variable among blue tit effects, and the effects themselves were also much more variable. Important to note here is that subsets of data may differ from each other for biological reasons, but they may also differ due to sampling error. Sampling error is a function of sample size, and sub-samples are, by definition, smaller samples, and so more subject to variability in effects due to sampling error [58]. Other features of datasets are also plausible candidates for driving heterogeneity in analytical outcomes, including features of covariates. In particular, relationships between covariates and the response variable as well as relationships between covariates and the primary independent variable (collinearity) can strongly influence the modeled relationship between the independent variable of interest and the dependent variable [59, 60]. Therefore, inclusion or exclusion of these covariates can drive heterogeneity in effect sizes (Zr). Also, as we saw with the two most extreme Zr values from the blue tit analyses, in multivariate models with collinear predictors, extreme effects can emerge when estimating partial correlation coefficients due to high collinearity, and conclusions can differ dramatically depending on which relationship receives the researcher's attention. Therefore, differences between datasets in the presence of strong and/or collinear covariates could influence the differences in heterogeneity in results among those datasets.

Although it is too early in the many-analyst research program to conclude which analytical decisions or which features of datasets are the most important drivers of heterogeneity in analytical outcomes, we must still grapple with the possibility that analytical outcomes may vary substantially based on

the choices we make as analysts. If we assume that, at least sometimes, different analysts will produce dramatically different statistical outcomes, what should we do as ecologists and evolutionary biologists? We review some ideas below.

The easiest path forward after learning about this analytical heterogeneity would be simply to continue with "business as usual", where researchers report results from a small number of statistical models. A case could be made for this path based on our results. For instance, among the blue tit analyses, the precise values of the estimated Zr effects varied substantially, but the average effect was convincingly different from zero, and a majority of individual effects (84%) were in the same direction. Arguably, many ecologists and evolutionary biologists appear primarily interested in the direction of a given effect and the corresponding p-value[61], and so the variability we observed when analyzing the blue tit dataset may not worry these researchers. Similarly, most effects from the Eucalyptus analyses were relatively close to zero, and about two-thirds of these effects did not cross the traditional threshold of statistical significance. Therefore, a large proportion of people analyzing these data would conclude that there was no effect, and this is consistent with what we might conclude from the meta-analysis.

However, we find the counter arguments to "business as usual" to be compelling. For blue tits, there were a substantial minority of calculated effects that would be interpreted by many biologists as indicating the absence of an effect (28%), and there were three traditionally 'significant' effects in the opposite direction to the average. The qualitative conclusions of analysts also reflected substantial variability, with fully half of teams drawing a conclusion distinct from the one we draw from the distribution as a whole. These teams with different conclusion were either uncertain about the negative relationship between competition and nestling growth, or they concluded that effects were mixed or absent. For the Eucalyptus analyses, this issue is more concerning. Around two-thirds of effects had confidence intervals overlapping zero, and of the third of analyses with confidence intervals excluding zero, almost half were positive, and the rest were negative. Accordingly, the

qualitative conclusions of the Eucalyptus teams were spread across the full range of possibilities. But even these problems are optimistic.

A potentially larger argument against "business as usual" is that it provides the raw material for biasing the literature. When different model specifications readily lead to different results, analysts may be tempted to report the result that appears most interesting, or that is most consistent with expectation [7, 12]. !There is growing evidence that researchers in ecology and evolutionary biology often report a biased subset of the results they produce [62, 63], and that this bias exaggerates the average size of effects in the published literature between 30 and 150% [9, 48].The bias then accumulates in meta-analyses, apparently more than doubling the rate of conclusions of "statistical significance" in published meta-analyses above what would have been found in the absence of bias [48]. Thus, "business as usual" does not just create noisy results, it helps create systematically misleading results.

Conclusions

Overall, our results suggest to us that, where there is a diverse set of plausible analysis options, no single analysis should be considered a complete or reliable answer to a research question. We contend that ecologists and evolutionary biologists typically do multiple analyses (as many of our analyst teams did) however, some of these analyses dont make it into the published manuscript. Further, because of the evidence that ecologists and evolutionary biologists often present a biased subset of the analyses they conduct [48, 62, 63], we do not expect that even a collection of different effect sizes from different studies will accurately represent the true distribution of effects [48]. Therefore, we believe that an increased level of skepticism of the outcomes of single analyses, or even single meta-analyses, is warranted going forward. We recognize that some researchers have long maintained a healthy level of skepticism of individual studies as part of sound and practical scientific practice, and it is possible that those researchers will be neither surprised nor concerned by

our results. However, we doubt that many researchers are sufficiently aware of the potential problems of analytical flexibility to be appropriately skeptical.

If we are skeptical of single analyses, the path forward may be multiple analyses per dataset. One possibility is the traditional robustness or sensitivity check [e.g., 64, 65], in which the researcher presents several alternative versions of an analysis to demonstrate that the result is 'robust' [66]. Unfortunately, robustness checks are at risk of the same potential biases of reporting found in other studies [11], especially given the relatively few models typically presented. However, these risks could be minimized by running more models and doing so with pre-registration or registered report.

Another option is model averaging. Averages across models often perform well [e.g., 67], and in some forms this may be a relatively simple solution. As most often practiced in ecology and evolutionary biology, model averaging involves first identifying a small suite of candidate models [see 13], then using Akaike weights, based on Akaike's Information Criterion (AIC), to calculate weighted averages for parameter estimates from those models. Again, the small number of models limits the exploration of specification space, but we can examine a larger number of models.

However, there are more concerning limitations. The largest of these limitations is that averaging regression coefficients is problematic when models differ in interaction terms or collinear variables [68]. Additionally, weighting by AIC may often be inconsistent with our modelling goals. AIC balances the trade-off between model complexity and predictive ability, but penalizing models for complexity may not be suited for testing hypotheses about causation. So, AIC may often not offer the weight we want to use for an average, and we may also not wish to just generate an average. Instead, if we hope to understand an extensive universe of possible modelling outcomes, we could conduct a multiverse analysis, possibly with a specification [10, 49]. This could mean running hundreds or thousands of models (or more!) to examine the distribution of possible effects, and to see how different specification choices map onto these effects. However, there is a trade-off between efficiently exploring large areas of specification space and limiting the analyses to biologically plausible specifications. Instead of simply identifying modelling decisions and creating all possible

combinations for the multiverse, a researcher could attempt to prevent implausible combinations, though the more variables in the dataset, the more difficult this becomes. To make this easier, one could recruit many analysts to each designate one or a few plausible specifications, as with our 'many analyst' study [11]. An alternative that may be more labor intensive for the primary analyst, but which may lead to a more plausible set of models, could involve hypothesizing about causal pathways with DAGs [directed acyclic graphs; [69]] to constrain the model set. Devoting this effort to thoughtful multiverse specifications, possibly combined with pre-registration to hinder undisclosed data dredging, seems worthy of consideration.

Although we have reviewed a variety of potential responses to the existence of variability in analytical outcomes, we certainly do not wish to imply that this is a comprehensive set of possible responses. Nor do we wish to imply that the opinions we have expressed about these options are correct. Determining how the disciplines of ecology and evolutionary biology should respond to knowledge of the variability in analytical outcome will benefit from the contribution and discussion of ideas from across these disciplines. We look forward to learning from these discussions and to seeing how these disciplines ultimately respond.

Declarations

Ethics approval and consent to participate

We obtained permission to conduct this research from the Whitman College Institutional Review Board (IRB). As part of this permission, the IRB approved the consent form (https://osf.io/xyp68/) that all participants completed prior to joining the study.

Consent for publication

Not applicable

Availability of data and materials

All data cleaning and preparation for our analyses was conducted in R (R Core Team 2022) and is publicly archived at (https://zenodo.org/doi/10.5281/zenodo.10046152). Please see session info for

we recommend viewing this manuscript in html format at https://egouldo.github.io/ManyAnalysts/

the full list of packages and their citations used in our analysis pipeline. We built an R package, *ManyEcoEvo* to conduct the analyses described in this chapter. This same package can be used to reproduce our analyses or replicate the analyses described here using alternate datasets.

Competing interests

The authors declare that they have no competing interests

Funding

EG's contributions were supported by an Australian Government Research Training Program Scholarship, AIMOS top-up scholarship (2022) and Melbourne Centre of Data Science Doctoral Academy Fellowship (2021). FF's contributions were supported by ARC Future Fellowship FT150100297.

Author's contributions

HF, THP and FF conceptualized the project. PV provided raw data for Eucalyptus analyses and SG and THP provided raw data for blue tit analyses. DGH, HF and THP prepared surveys for collecting participating analysts and reviewer's data. EG, HF, THP, PV, SN and FF planned the analyses of the data provided by our analysts and reviewers, EG, HF, and THP curated the data, EG and HF wrote the software code to implement the analyses and prepare data visualisations. EG ensured that analyses were documented and reproducible. THP and HF administered the project, including coordinating with analysts and reviewers. FF provided funding for the project. THP, HF, and EG wrote the manuscript. Authors listed alphabetically contributed analyses of the primary datasets or reviews of analyses. All authors read and approved the final manuscript.

Acknowledgements

Not applicable

References

1. Arif S, MacNeil MA. Applying the structural causal model framework for observational causal inference in ecology. Ecological Monographs. 2023;93:e1554.

2. Atkinson J, Brudvig LA, Mallen-Cooper M, Nakagawa S, Moles AT, Bonser SP. Terrestrial ecosystem restoration increases biodiversity and reduces its variability, but not to reference levels: A global meta-analysis. Ecology Letters. 2022;25:1725–37.

3. Auspurg K, Brüderl J. Has the credibility of the social sciences been credibly destroyed? Reanalyzing the “many analysts, one data set” project. Socius. 2021;7:23780231211024421.

4. Schloerke B, Cook D, Larmarange J, Briatte F, Marbach M, Thoen E, et al. GGally: Extension to ’ggplot2’. 2022.

5. Baselga A, Orme D, Villeger S, De Bortoli J, Leprieur F, Logez M, et al. Package “betapart”. 2023.

6. Bates D, Mächler M, Bolker B, Walker S. Fitting linear mixed-effects models using lme4. 2015. 2015;67:48.

7. Bolker B, Robinson D, Menne D, Gabry J, Buerkner P, Hau C, et al. Package “broom.mixed”. 2022.

8. Borenstein M, Higgins JPT, Hedges L, Rothstein H. Basics of meta-analysis: I2 is not an absolute measure of heterogeneity. Research Synthesis Methods. 2017;8:5–18.

9. Botvinik-Nezer R, Holzmeister F, Camerer CF, Dreber A, Huber J, Johannesson M, et al. Variability in the analysis of a single neuroimaging dataset by many teams. Nature. 2020;582:84–8.

10. Breznau N, Rinke EM, Wuttke A, Nguyen HHV, Adem M, Adriaans J, et al. Observing many researchers using the same data and hypothesis reveals a hidden universe of uncertainty. Proceedings of the National Academy of Sciences. 2022;119:e2203150119.

11. Briga M, Verhulst S. Mosaic metabolic ageing: Basal and standard metabolic rates age in opposite directions and independent of environmental quality, sex and life span in a passerine. Functional Ecology. 2021;35:1055–68.

12. Burnham KP, Anderson DR. Model selection and multimodel inference: A practical information-theoretical approach. 2nd edition. Book. New York: Springer-Verlag; 2002.

13. Cade BS. Model averaging and muddled multimodel inferences. Ecology. 2015;96:2370–82.

14. Capilla-Lasheras P, Thompson MJ, Sánchez-Tójar A, Haddou Y, Branston CJ, Réale D, et al. A global meta-analysis reveals higher variation in breeding phenology in urban birds than in their non-urban neighbours. Ecology Letters. 2022;25:2552–70.

15. Coretta S, Casillas JV, Roessig S, Franke M, Ahn B, Al-Hoorie AH, et al. Multidimensional signals and analytic flexibility: Estimating degrees of freedom in human-speech analyses. Advances in Methods and Practices in Psychological Science. 2023;6:25152459231162567.

16. DeKogel CH. Long-term effects of brood size manipulation on morphological development and sex-specific mortality of offspring. Journal of Animal Ecology. 1997;66:167–78.

17. Deressa T, Stern D, Vangronsveld J, Minx J, Lizin S, Malina R, et al. More than half of statistically significant research findings in the environmental sciences are actually not. EcoEvoRxiv. 2023. https://doi.org/https://doi.org/10.32942/X24G6Z.

18. Dormann CF, Elith J, Bacher S, Buchmann C, Carl G, Carré G, et al. Collinearity: A review of methods to deal with it and a simulation study evaluating their performance. Ecography. 2013;36:27–46.

19. Fanelli D, Costas R, Ioannidis JPA. Meta-assessment of bias in science. Proceedings of the National Academy of Sciences. 2017;114:3714–9.

20. Fanelli D, Ioannidis JPA. US studies may overestimate effect sizes in softer research. Proceedings of the National Academy of Sciences. 2013;110:15031–6.

21. Fidler F, Burgman MA, Cumming G, Buttrose R, Thomason N. Impact of criticism of null-hypothesis significance testing on statistical reporting practices in conservation biology. Conservation Biology. 2006;20:1539–44.

22. Fidler F, Chee YE, Wintle BC, Burgman MA, McCarthy MA, Gordon A. Metaresearch for evaluating reproducibility in ecology and evolution. BioScience. 2017;67:282–9.

23. Forstmeier W, Wagenmakers E-J, Parker TH. Detecting and avoiding likely false-positive findings – a practical guide. Biological Reviews. 2017;92:1941–68.

24. Fraser H, Parker T, Nakagawa S, Barnett A, Fidler F. Questionable research practices in ecology and evolution. PLOS ONE. 2018;13:e0200303.

25. Gelman A, Weakliem D. Of beauty, sex, and power. American Scientist. 2009;97:310–6.

26. Gelman A, Loken E. The garden of forking paths: Why multiple comparisons can be a problem, even when there is no “fishing expedition” or “p-hacking” and the research hypothesis was posited ahead of time. Department of Statistics, Columbia University. 2013.

27. Grueber CE, Nakagawa S, Laws RJ, Jamieson IG. Multimodel inference in ecology and evolution: Challenges and solutions. Journal of Evolutionary Biology. 2011;24:699–711.

28. Higgins JPT, Thompson SG, Deeks JJ, Altman DG. Measuring inconsistency in meta-analyses. BMJ. 2003;327:557–60.

29. Huntington-Klein N, Arenas A, Beam E, Bertoni M, Bloem JR, Burli P, et al. The influence of hidden researcher decisions in applied microeconomics. Economic Inquiry. 2021;59:944–60.

30. Jennions MD, Lortie CJ, Rosenberg MS, Rothstein HR. Publication and related biases. In: Koricheva J, Gurevitch J, Mengersen K, editors. Handbook of meta-analysis in ecology and evolution. Princeton, USA: Princeton University Press; 2013. p. 207–36.

31. Kimmel K, Avolio ML, Ferraro PJ. Empirical evidence of widespread exaggeration bias and selective reporting in ecology. Nature Ecology & Evolution. 2023. https://doi.org/10.1038/s41559-023-02144-3.

32. Klein RA, Ratliff KA, Vianello M, Jr. RBA, Bahník Š, Bernstein MJ, et al. Investigating variation in replicability: A "many labs" replication project. Social Psychology. 2014;45:142–52.

33. Klein RA, Vianello M, Hasselman F, Adams BG, Adams RB, Alper S, et al. Many labs 2: Investigating variation in replicability across samples and settings. Advances in Methods and Practices in Psychological Science. 2018;1:443–90.

34. Knight K. Mathematical statistics. Book. New York: Chapman; Hall; 2000.

35. Kou-Giesbrecht S, Menge DNL. Nitrogen-fixing trees increase soil nitrous oxide emissions: A meta-analysis. Ecology. 2021;102:e03415.

1599 36. Kuznetsova A, Brockhoff PB, Christensen RHB. lmerTest package: Tests in linear mixed effects
1600 models. Journal of Statistical Software. 2017;82:1–26.

1601 37. Leybourne DJ, Preedy KF, Valentine TA, Bos JIB, Karley AJ. Drought has negative consequences on
1602 aphid fitness and plant vigor: Insights from a meta-analysis. Ecology and Evolution. 2021;11:11915–
1603 29.

1604 38. Lu X, White H. Robustness checks and robustness tests in applied economics. Journal of
1605 Econometrics. 2014;178:194–206.

1606 39. Lüdecke D, Ben-Shachar MS, Patil I, Waggoner P, Makowski D. Performance: An r package for
1607 assessment, comparison and testing of statistical models. Journal of Open Source Software.
1608 2021;6:3139.

1609 40. Luke SG. Evaluating significance in linear mixed-effects models in r. Behavior Research Methods.
1610 2017;49:1494–502.

1611 41. Miles C. Testing market-based instruments for conservation in northern victoria. In: Norton T,
1612 Lefroy T, Bailey K, Unwin G, editors. Biodiversity: Integrating conservation and production: Case
1613 studies from australian farms, forests and fisheries. Melbourne, Australia: CSIRO Publishing; 2008. p.
1614 133–46.

1615 42. Morrissey MB, Ruxton GD. Multiple regression is not multiple regressions: The meaning of
1616 multiple regression and the non-problem of collinearity. Philosophy, Theory, and Practice in Biology.
1617 2018;10.

1618 43. Nakagawa S, Cuthill IC. Effect size, confidence interval and statistical significance: A practical
1619 guide for biologists. Biological Reviews. 2007;82:591–605.

1620 44. Nakagawa S, Noble DW, Senior AM, Lagisz M. Meta-evaluation of meta-analysis: Ten appraisal
1621 questions for biologists. BMC Biology. 2017;15:18.

1622 45. Nicolaus M, Michler SPM, Ubels R, Velde M van der, Komdeur J, Both C, et al. Sex-specific effects
1623 of altered competition on nestling growth and survival: An experimental manipulation of brood size
1624 and sex ratio. Journal of Animal Ecology. 2009;78:414–26.

1625 46. Noble DWA, Lagisz M, O'Dea RE, Nakagawa S. Nonindependence and sensitivity analyses in
1626 ecological and evolutionary meta-analyses. Molecular Ecology. 2017;26:2410–25.

1627 47. Open Science Collaboration. Estimating the reproducibility of psychological science. Science.
1628 2015;349:aac4716.

1629 48. Parker TH, Forstmeier W, Koricheva J, Fidler F, Hadfield JD, Chee YE, et al. Transparency in ecology
1630 and evolution: Real problems, real solutions. Trends in Ecology & Evolution. 2016;31:711–9.

1631 49. Parker TH, Yang Y. Exaggerated effects in ecology. Nature Ecology & Evolution. 2023.
1632 https://doi.org/10.1038/s41559-023-02156-z.

1633 50. Pei Y, Forstmeier W, Wang D, Martin K, Rutkowska J, Kempenaers B. Proximate causes of infertility
1634 and embryo mortality in captive zebra finches. The American Naturalist. 2020;196:577–96.

1635 51. R Core Team. R: A language and environment for statistical computing. Vienna, Austria: R
1636 Foundation for Statistical Computing; 2022.

52. Rosenberg MS. Moment and least-squares based approaches to metaanalytic inference. In: Koricheva J, Gurevitch J, Mengersen K, editors. Handbook of meta-analysis in ecology and evolution. Princeton, USA: Princeton University Press; 2013. p. 108–24.

53. Royle NJ, Hartley IR, Owens IPF, Parker GA. Sibling competition and the evolution of growth rates in birds. Proceedings of the Royal Society B-Biological Sciences. 1999;266:923–32.

54. Schweinsberg M, Feldman M, Staub N, Akker OR van den, Aert RCM van, Assen M van, et al. Same data, different conclusions: Radical dispersion in empirical results when independent analysts operationalize and test the same hypothesis. Organizational Behavior and Human Decision Processes. 2021;165:228–49.

55. Senior AM, Grueber CE, Kamiya T, Lagisz M, O'Dwyer K, Santos ESA, et al. Heterogeneity in ecological and evolutionary meta-analyses: Its magnitude and implications. Ecology. 2016;97:3293–9.

56. Shavit A, Ellison AM. Stepping in the same river twice: Replication in biological research. Edited Book. New Haven, Connecticut, USA: Yale University Press; 2017.

57. Siegel KR, Kaur M, Grigal AC, Metzler RA, Dickinson GH. Meta-analysis suggests negative, but pCO2-specific, effects of ocean acidification on the structural and functional properties of crustacean biomaterials. Ecology and Evolution. 2022;12:e8922.

58. Silberzahn R, Uhlmann EL, Martin DP, Anselmi P, Aust F, Awtrey E, et al. Many analysts, one data set: Making transparent how variations in analytic choices affect results. Advances in Methods and Practices in Psychological Science. 2018;1:337–56.

59. Simons DJ, Shoda Y, Lindsay DS. Constraints on generality (COG): A proposed addition to all empirical papers. Perspectives on Psychological Science. 2017. https://doi.org/10.1177/174569161770863.

60. Simonsohn U, Simmons JP, Nelson LD. Specification curve: descriptive and inferential statistics on all reasonable specifications. SSRN Electronic Journal. 2015. https://doi.org/10.2139/ssrn.2694998 .

61. Simonsohn U, Simmons JP, Nelson LD. Specification curve analysis. Nature Human Behaviour. 2020;4:1208–14.

62. Steegen S, Tuerlinckx F, Gelman A, Vanpaemel W. Increasing transparency through a multiverse analysis. Perspectives on Psychological Science. 2016;11:702–12.

63. Taylor JW, Taylor KS. Combining probabilistic forecasts of COVID-19 mortality in the united states. European Journal of Operational Research. 2023;304:25–41.

64. Dancho M, Vaughan D. Timetk: A tool kit for working with time series. 2023.

65. Vander Werf E. Lack's clutch size hypothesis: An examination of the evidence using meta-analysis. Ecology. 1992;73:1699–705.

66. Ver Hoef JM. Who invented the delta method? The American Statistician. 2012;66:124–7.

67. Verhulst S, Holveck MJ, Riebel K. Long-term effects of manipulated natal brood size on metabolic rate in zebra finches. Biology Letters. 2006;2:478–80.

68. Vesk PA, Morris WK, McCallum W, Apted R, Miles C. Processes of woodland eucalypt regeneration: Lessons from the bush returns trial. Proceedings of the Royal Society of Victoria. 2016;128:54–63.

69. Viechtbauer W. Conducting meta-analyses in r with the metafor package. 2010. 2010;36:48.

70. Yang Y, Sánchez-Tójar A, O'Dea RE, Noble DWA, Koricheva J, Jennions MD, et al. Publication bias impacts on effect size, statistical power, and magnitude (type m) and sign (type s) errors in ecology and evolutionary biology. BMC Biology. 2023;21:71.