Year: 2024

Cultural Evolution and Social Learning: An Evolutionary Approach to Group and Individual Decision-Making Dynamics

FAESSLER Lisa

UNIL | Université de Lausanne

FACULTÉ DES HAUTES ÉTUDES COMMERCIALES

DÉPARTEMENT DE COMPORTEMENT ORGANISATIONNEL

Cultural Evolution and Social Learning:

An Evolutionary Approach to Group and Individual Decision-Making Dynamics

THÈSE DE DOCTORAT

présentée à la

Faculté des Hautes Études Commerciales
de l'Université de Lausanne

pour l'obtention du grade de
Doctorat en Management

par

Lisa FAESSLER

Directeur de thèse
Prof. Charles Efferson

Jury

Prof. Felicitas Morhart, Présidente
Prof. Christian Zehnder, expert interne
Prof. Michael Muthukrishna, expert externe

LAUSANNE
2024

UNIL | Université de Lausanne

FACULTÉ DES HAUTES ÉTUDES COMMERCIALES

DÉPARTEMENT DE COMPORTEMENT ORGANISATIONNEL

Cultural Evolution and Social Learning:

An Evolutionary Approach to Group and Individual Decision-Making Dynamics

THÈSE DE DOCTORAT

présentée à la

Faculté des Hautes Études Commerciales
de l'Université de Lausanne

pour l'obtention du grade de
Doctorat en Management

par

Lisa FAESSLER

Directeur de thèse
Prof. Charles Efferson

Jury

Prof. Felicitas Morhart, Présidente
Prof. Christian Zehnder, expert interne
Prof. Michael Muthukrishna, expert externe

LAUSANNE
2024

UNIL | Université de Lausanne
Faculté des hautes études commerciales
Internef - 261
CH-1015 Lausanne

IMPRIMATUR

La Faculté des hautes études commerciales de l'Université de Lausanne autorise l'impression de la thèse de doctorat rédigée par

Lisa FAESSLER

intitulée

Cultural Evolution and Social Learning: an Evolutionary Approach to Group and Individual Decision-Making Dynamics

sans se prononcer sur les opinions exprimées dans cette thèse.

Lausanne, le 04.07.2024

Professeure Marianne Schmid Mast, Doyenne

Thesis Committee

Prof. Charles EFFERSON

Thesis Supervisor

University of Lausanne

Prof. Christian ZEHNDER

Internal Expert

University of Lausanne

Prof. Michael MUTHUKRISHNA

External Expert

London School of Economics

University of Lausanne
Faculty of Business and Economics

Ph.D. in Management

I hereby certify that I have examined the doctoral thesis of

Lisa FAESSLER

and have found it to meet the requirements for a doctoral thesis.
All revisions that I or committee members
made during the doctoral colloquium
have been addressed to my entire satisfaction.

Signature: ______________________ Date: 19 June 2024

Prof. Charles EFFERSON
Thesis supervisor

University of Lausanne
Faculty of Business and Economics

Ph.D. in Management

I hereby certify that I have examined the doctoral thesis of

Lisa FAESSLER

and have found it to meet the requirements for a doctoral thesis.
All revisions that I or committee members
made during the doctoral colloquium
have been addressed to my entire satisfaction.

Signature: ______________________________ Date: 18.6.2024

Prof. Christian ZEHNDER
Internal expert

University of Lausanne
Faculty of Business and Economics

Ph.D. in Management

I hereby certify that I have examined the doctoral thesis of

Lisa FAESSLER

and have found it to meet the requirements for a doctoral thesis.
All revisions that I or committee members
made during the doctoral colloquium
have been addressed to my entire satisfaction.

Signature: ______________________________ Date: 18 June 2024

Prof. Michael MUTHUKRISHNA
External expert

Contents

Introduction

This thesis is about cultural evolution. Cultural evolution is a specific research field that does not tell by its name what it is about. "Culture" here does not refer to intellectual or artistic creations. Instead, culture means the information stored in people's heads, learned by others, and susceptible to modify people's behaviors (Richerson & Boyd 2005). "Evolution" does not refer to some recent development. Instead, evolution means the change of these cultural characteristics over hundreds of thousands of generations under the influence of evolutionary processes such as natural selection and genetic drift.

Evolutionary approached to social sciences

As an evolutionary approach to social sciences, cultural evolution offers to apply principles of evolution, such as natural selection, adaptation, and survival of the fittest, to understand human behavior and societal dynamics. Evolutionary approaches suggest that human behaviors, social norms, and institutions evolve over time in response to social learning dynamics, much like biological traits. They offer a comprehensive framework for analyzing how behaviors and social systems develop and change. An illustrative example of the power of evolutionary explanations is the study of cooperation (Axelrod & Hamilton 1981, Efferson et al. 2024). Evolutionary theories propose that cooperative behaviors may have evolved because they increase individuals' chances of survival and reproduction within a group. This perspective helps us design policies and interventions promoting cooperative behaviors in modern contexts, such as business environments or conservation efforts.

Cultural evolution

Cultural evolution is a subset of evolutionary approaches that focuses on how culture changes and adapts over time. Evolution operates through two primary processes: genetic and cultural evolution (Laland 2008). Genetic evolution involves changes in gene frequencies within a population over time, driven by natural selection, mutation, and genetic drift. On the other hand, cultural evolution involves changes in cultural traits driven by social learning and cultural transmission. Cultural evolution researchers study the transmission and transformation of cultural traits—such as beliefs, practices, and technologies—across generations.

Culture is defined as the set of learned behaviors, beliefs, and technologies transmitted through social learning rather than genetic inheritance or individual learning (Richerson & Boyd 2005). This includes everything we learned from others, from language and religious practices to technological innovations and social norms. The field of cultural evolution relies on two key assumptions:

- Learning from others influences behavior: Individuals adopt behaviors, beliefs, and knowledge from others within their society. This social learning is a primary mechanism through which culture is transmitted.

- Social learning is not random: Individuals do not learn randomly(McElreath et al. 2008). Instead, they are selective in whom they learn from and what they learn (Mesoudi et al. 2016, Morgan et al. 2012). This selectivity implies that cultural evolution is a systematic process.

These assumptions lead to the conclusion that culture evolves in a systematic way at the group level (Granovetter 1978, Young 2009, Efferson et al. 2008). Cultural traits that confer advantages—such as higher prestige or greater success—are more likely to be transmitted and adopted, leading to cumulative cultural evolution (Henrich & Gil-White 2001, Muthukrishna et al. 2016). The consequences of these assumptions are profound. Cultural evolution would be an adaptive process, much like biological evolution. Understanding this process would help us address societal challenges by leveraging cultural dynamics for positive outcomes.

Causal effect of culture

The first assumption of cultural evolution—that learning from others influences behavior—focuses on understanding the causal effect of culture. Understanding the causal effect of culture allows us to disentangle the influences of culture from other factors, such as genetics or environmental factors. However, identifying culture in the field is an empirical challenge as culture often covaries with many cofounds such as country, institutions, environment, and genetics. That is the challenge my coauthor and I tackled in my first thesis chapter.

To meet this challenge, we exploit the Rostigraben, a linguistic and cultural border that divides Switzerland in ways independent of institutional, environmental, and genetic variation. Using a regression discontinuity design, we estimate discontinuities at the border regarding preferences related to fertility and mortality, the two basic components of genetic fitness. We specifically select six referenda related to health and fertility and analyze differences in the proportion of yes votes across municipalities on the two sides of the border. Our results show multiple discontinuities in voting behaviors at the language border. Of those significant discontinuities, three are related to fertility, and one is related to health. These discontinuities suggest the potential for culture to create stable differences between groups in domains related to health and fertility, where cultural explanations are distinct from institutional, genetic, and environmental explanations.

Further, for each of the referenda, we speculate how these cultural differences could affect the relative fitness values of individuals in the two cultural groups. In this way, although we do not examine genetic fitness directly, we do lean in this direction by focusing on cultural variation in support for policies that should influence fertility, health, and survival. The variation in question is a group-level phenomenon based on cultural evolutionary processes, but it should have consequences for individual reproduction and, by extension, fitness. Our ultimate goal would be to demonstrate the interplay between genetic and cultural evolution processes. Genes could influence culture, and culture could influence genes, leading to a gene-culture co-evolutionary process.

Social learning complexity

Back to cultural evolution, the second assumption—social learning is systematic—emphasizes the importance of understanding how individuals learn from each other. Social learning is the primary mechanism through which cultural traits are transmitted. Despite recognizing that social learning is non-random, we still face challenges defining what "non-random" means. Further, small differences at the individual level can have far-reaching consequences on cultural evolution at the aggregate level (Granovetter 1978, Young 2009, Efferson et al. 2020). In my second chapter, I aim to answer part of that question by exploring the complexity of success-biases social learning strategies.

People tend to imitate successful individuals more than those who are not successful, and leaders leverage this tendency when "leading by success." However, social learning strategies are more complex than we often assume. People do not simply follow success; they can do much more. The added complexity and flexibility could challenge current success-based leadership techniques and cultural evolution dynamics. Focusing on success-biased social learning, we studied how individuals adapt their behaviors based on the actions of successful leaders.

Through an incentivized experiment, we examine participants' decision-making after observing the behavior of a successful leader and two additional signals: the leader's group affiliation and affiliation implications. An accompanying gene-culture co-evolutionary agent-based simulation integrates cognitive mechanisms to process the same three types of social information, along with private information. Our findings highlight three critical aspects of success-biased social learning. First, strategies are multi-dimensional. Participants and agents adjust their responses based on multiple pieces of social information. They integrate all the available pieces of information into their strategies: success, group membership of the successful leader, and relevance of that information. Second, while adjustments are symmetric in the simulation, the experiment shows participants perform better in certain conditions than others, indicating cognitive biases. Learning is easier in certain conditions than in others. Finally, the simulation and the experiment demonstrate significant heterogeneity and flexibility in the use of social learning strategies. Our results show that followers adjust to success-dependent social information in complex and heterogeneous ways, including using successful leaders as negative examples.

Collective intelligence

Culture is not a human's prerogative, but cumulative cultural evolution is, which would be the key to human adaptation success (Henrich 2016). Humans' groups and societies became collective brains (Muthukrishna & Henrich 2016). Collective brains refer to the enhanced capacity for problem-solving, innovation, and decision-making that arises when individuals work together compared to isolated individuals. The performance of the collective brain is a function of three levers: size and connectivity, transmission fidelity, and cultural trait diversity (Schimmelpfennig et al. 2022). Each of them has the potential to improve or impair innovation rate. This phenomenon is called the paradox of diversity and can lead to counterintuitive findings.

For example, larger populations tend to generate more ideas and innovations due to the diverse range of experiences and knowledge within the group. Increased connectivity facilitates the sharing and refinement of these ideas, enhancing the group's overall problem-solving. However, highly connected teams perform worse than moderately connected teams. Connectivity can lead people to over-rely on social information. Conformity homogenizes the population and stifles innovation.

My third chapter examines the issue of overexploitation in highly connected teams. I wonder whether turnover and new team members' arrival can compensate for a team's high connectivity. Turnover is costly in many ways, but its virtue may lie in the diversity that newcomers bring. In a lab experiment, participants had to solve a complex task in teams of three people. Two forms of disruption were introduced. First, some teams experienced turnover, while others remained stable. Second, among teams that had not reached the highest-performing solution, some were selectively informed about the existence of superior solutions, while others were not. The results indicate an increase in search distance following the disruptions in all treatment conditions. However, the increase was modest and temporary, insufficient to create alternative solutions and improve performance. Instead, payoffs decreased significantly after the treatment in all conditions before returning to their initial improving trend. Interestingly, newcomers did not explore more than oldtimers. Instead, they quickly conformed to the group solution and exploitation rate. Comparison with simulated performance benchmarks suggests that increased exploration would have led to a better overall performance.

Evolutionary approaches to social sciences and cultural evolution offer powerful frameworks for understanding human behavior and group dynamics, whether at the team or the societal level. By integrating these perspectives into various fields, including management, leadership, and organizational theory, we can develop more effective strategies for addressing complex organizational and social challenges.

Bibliography

Axelrod, R. & Hamilton, W. D. (1981), 'The evolution of cooperation', *science* **211**(4489), 1390–1396.

Efferson, C., Bernhard, H., Fischbacher, U. & Fehr, E. (2024), 'Super-additive cooperation', *Nature* pp. 1–8.

Efferson, C., Lalive, R. & Fehr, E. (2008), 'The coevolution of cultural groups and ingroup favoritism', *Science* **321**(5897), 1844–1849.

Efferson, C., Vogt, S. & Fehr, E. (2020), 'The promise and the peril of using social influence to reverse harmful traditions', *Nature human behaviour* **4**(1), 55–68.

Granovetter, M. (1978), 'Threshold models of collective behavior', *American journal of sociology* **83**(6), 1420–1443.

Henrich, J. (2016), *The Secret of Our Success: How culture Is Driving Human Evolution, Domesticating Our Species, and Making Us Smarter*, Princeton University Press.

Henrich, J. & Gil-White, F. J. (2001), 'The evolution of prestige: Freely conferred deference as a mechanism for enhancing the benefits of cultural transmission', *Evolution and human behavior* **22**(3), 165–196.

Laland, K. N. (2008), 'Exploring gene–culture interactions: insights from handedness, sexual selection and niche-construction case studies', *Philosophical Transactions of the Royal Society B: Biological Sciences* **363**(1509), 3577–3589.

McElreath, R., Bell, A. V., Efferson, C., Lubell, M., Richerson, P. J. & Waring, T. (2008), 'Beyond existence and aiming outside the laboratory: estimating frequency-dependent and pay-off-biased social learning strategies', *Philosophical Transactions of the Royal Society B: Biological Sciences* **363**(1509), 3515–3528.

Mesoudi, A., Chang, L., Dall, S. R. & Thornton, A. (2016), 'The evolution of individual and cultural variation in social learning', *Trends in ecology & evolution* **31**(3), 215–225.

Morgan, T., Rendell, L., Ehn, M., Hoppitt, W. & Laland, K. (2012), 'The evolutionary basis of human social learning', *Proceedings of the Royal Society B: Biological Sciences* **279**(1729), 653–662.

Muthukrishna, M. & Henrich, J. (2016), 'Innovation in the collective brain', *Philosophical Transactions of the Royal Society B: Biological Sciences* **371**(1690), 20150192.

Muthukrishna, M., Morgan, T. J. & Henrich, J. (2016), 'The when and who of social learning and conformist transmission', *Evolution and Human Behavior* **37**(1), 10–20.

Richerson, P. J. & Boyd, R. (2005), *Not by genes alone: how culture transformed human evolution*, The University of Chicago Press.

Schimmelpfennig, R., Razek, L., Schnell, E. & Muthukrishna, M. (2022), 'Paradox of diversity in the collective brain', *Philosophical Transactions of The Royal Society B* **377**(1843), 20200316.

Young, H. P. (2009), 'Innovation diffusion in heterogeneous populations: Contagion, social influence, and social learning', *American economic review* **99**(5), 1899–1924.

Chapter 1

How culture shapes choices related to fertility and mortality: Causal evidence at the Swiss language border

Lisa Faessler, Rafael Lalive, Charles Efferson

Faculty of Business and Economics, University of Lausanne, Switzerland

Abstract

Results from cultural evolutionary theory often suggest that social learning can lead cultural groups to differ markedly in the same environment. Put differently, cultural evolutionary processes can in principle stabilise behavioural differences between groups, which in turn could lead selection pressures to vary across cultural groups. Separating the effects of culture from other confounds, however, is often a daunting, sometimes intractable challenge for the working empiricist. To meet this challenge, we exploit a cultural border dividing Switzerland in ways that are independent of institutional, environmental, and genetic variation. Using a regression discontinuity design, we estimate discontinuities at the border in terms of preferences related to fertility and mortality, the two basic components of genetic fitness. We specifically select six referenda related to health and fertility and analyse differences in the proportion of yes votes across municipalities on the two sides of the border. Our results show multiple discontinuities and thus indicate a potential role of culture in shaping stable differences between groups in preferences and choices related to individual health and fertility. These findings further suggest that at least one of the two groups, in order to uphold its cultural values, has supported policies that could impose fitness costs on individuals relative to the alternative policy under consideration.

Social media summary: Discontinuities at a language border in Switzerland show that culture can shape choices related to health and fertility.

Keywords: gene-culture coevolution, social learning, cultural variation, cultural border, regression discontinuity design

1.1 Introduction

Gene-culture coevolutionary theory argues that human populations are subject to two evolutionary processes, genetic and cultural (Laland 2008). Genetic variants influence the development and spread of cultural traits, while cultural practices affect selection on genes. As a result, genes and culture coevolve as linked dynamical processes. As a kind of corollary hypothesis, an especially controversial claim is that social learning stabilises cultural differences at the group level, which in turn is a necessary but not sufficient condition for any kind of selection at the level of the cultural group (Henrich 2004, Richerson et al. 2016).

We examine a kind of proof of concept for these ideas. Specifically, we do not directly consider culture's influence on genetic fitness, but we do insist on an attempt to identify cleanly the causal influence of culture on decisions affecting health and fertility. Identifying cultural variation as a group-level phenomenon is often a difficult empirical challenge because culture typically covaries with many other variables related to institutions, the environment, and possibly even genes. To meet this challenge, we exploit a distinctive feature of Switzerland's geography, a linguistic and cultural border that separates the German-speaking part of the country from the French-speaking part. Right at the border, the environments for French speakers and German speakers are necessarily identical. Moreover, the French- and German-speaking parts of the country are genetically similar in general (Buhler et al. 2012). Finally, in some regions, the border does not match any institutional boundary. Thus, right at the border, we have the possibility of observing variation in preferences and norms that we can say is cultural in the precise sense that it cannot be institutional, environmental, or genetic. This situation represents an unusual opportunity because cultures often covary with one or more of these variables.

Consider two examples that illustrate the challenges of isolating culture in domains that could influence selection on genes. First, lactase persistence is a classic example. In most mammals, including humans, lactase production declines after weaning, but some populations have evolved the ability to produce lactase throughout adulthood, a condition known as lactase persistence. This adaptation is thought to have arisen in response to the cultural practice of dairy farming, which allowed people to consume milk and dairy products as a significant part of their diet. Nonetheless, recent evidence suggests that multiple factors, including different environmental conditions, have contributed to lactase persistence, and that dairying alone is probably insufficient to explain the spread of the trait. In particular, exposure to famine and diseases has played a crucial role in the evolution of lactase persistence (Evershed et al. 2022).

Second, the cultural practice of cooking and its influence on human gut size is another classic example. Cooking allows us to pre-digest our food over the campfire or on the stove, which improves the biological availability of the nutrients in the food. Cooking as a cultural innovation likely allowed our ancestors to evolve smaller guts because they were able to extract more energy from their food for a given metabolic cost (Wrangham & Carmody 2010). Thus, energetic resources within the body became available for other functions such as brain growth and development. This shift in energy allocation is thought to have played a key role in the evolution of larger brains and shorter digestive tracts in humans compared to our primate relatives (Navarrete et al. 2011). However, cooking is one of the few human cross-cultural universals. There is no such thing as a human group that does not engage in cooking. Therefore, establishing a causal link between the cultural practice of cooking and alterations in the human gut remains impossible, given the absence of a counterfactual. Stories of this sort are interesting and compelling, and they may very well be correct. They are not, however, causal explanations. Valid comparisons that we could rely on to represent the counterfactual state are not available to us and probably never will be.

1.1.1 Identifying Culture

Identifying the causal influence of culture on gene selection is a challenge. Comparing the average behaviours of two populations (Bell et al. 2009) often cannot provide evidence for cultural variation . If environmental conditions, institutions, and other socioeconomic variables covary with culture, isolating the extent to which group-level variation is specifically cultural can be exceedingly difficult. Lamba & Mace (2011), for example, compared groups within the same culture but living in different locations, and they found substantial variation across the groups. This kind of result suggests that large differences among groups can be environmental just as surely as they can be cultural, and indeed recent evidence suggests that ecology can explain a substantial amount of human population diversity (Wormley et al. 2022).

That said, a number of new tools have been developed to allow the identification of causal effects without randomised experiments, and these tools can potentially help us identify culture. These quasi-experimental methods include the regression discontinuity design (RDD). The basic idea of the regression discontinuity design is to compare the outcomes of individuals just above and below some threshold. Intuitively, researchers estimate two regression lines, one on each side of the threshold, and doing so identifies any discontinuities in the response variable that occur right at the threshold (Lee & Lemieux 2010, Cattaneo et al. 2019). A few studies have used a variant of this method, the spatial regression discontinuity design, to identify cultural discontinuities and the Swiss language border. We adopt the same basic approach here.

These studies are known as “Röstigraben studies”, a type of spatial regression discontinuity design that examines cultural differences in behaviour in Switzerland. The term “Röstigraben” – German for “hash brown trench” – refers to a linguistic and cultural border within Switzerland. The border separates the German-speaking part from the French-speaking part of the country, and in some regions it does not match any institutional boundary. With appropriate data, researchers could in principle check for discontinuities in any variable of interest right at the language border, and by doing so the researcher would effectively isolate cultural differences, as a group-level phenomenon, in identical institutional and ecological settings. Using this technique, Eugster et al. (2011) document a persistent difference in the demand for social insurance at the border, and Eugster et al. (2017) also found a significant discontinuity in unemployment duration. Focusing on the bilingual canton of Fribourg, Brown et al. (2018) discovered a systematic difference in the financial literacy of students across the border, and their analyses suggest that the effect is driven by cultural differences rather than unobserved heterogeneity in policies.

1.1.2 Switzerland’s Linguistic and Cultural Landscape

Switzerland is a multilingual country with four official languages: German, French, Italian, and Romansh. German is the most widely spoken language at home (62%), while French is second (22.8%). Switzerland’s linguistic diversity is a unique feature that has played a significant role in shaping its culture and society. Multilingualism is a common characteristic among Swiss people. However, the historical border between the French- and the German-speaking regions has remained clear-cut. A sharp change in the main language spoken at home persists when switching from one side of the border to the other (OFS 2022*a*). Because the language border is clear and well-defined in space, we can meaningfully isolate discontinuous differences that occur right at the border.

Beyond language, conventional wisdom posits that this linguistic border also captures differences in values, norms, and preferences. Swiss media and citizens often view it as a cultural divide that marks contrasting attitudes. During federal elections, when voting on shared issues, these differences become especially apparent (Etter et al. 2014). Furthermore, the French- and

German-speaking regions show distinct patterns of health-related behaviours *on average*. For instance, French speakers typically consume more red meat but less butter, milk, and coffee than their German-speaking counterparts (Chatelan et al. 2017, Rochat et al. 2019). These comparisons of group averages do not provide causal evidence, but they do fit with the conventional wisdom within Switzerland. When you cross the Röstigraben, it's not just the language that changes; culture more broadly changes, too. That said, we can check to see if this is the case with a spatial regression discontinuity design. The basic idea is to code variables of interest as a function of distance from the language border, and then use the method to estimate any discontinuities in the variables right at the border. Doing so is effectively like comparing what happens one meter to the east of the border to what happens one meter west of the border.

1.1.3 The Cultural Components of Fitness

Having explained our strategy to isolate culture's causal effect, we now turn to the second consideration. Namely, what kinds of available data connect possible cultural differences within Switzerland to fertility and mortality, the two basic components of genetic fitness? In our study, we focus on the tendency of people to vote for or against policies that should impact either the survival or reproduction of individuals. In Switzerland, the leading causes of death are predominantly disease. In 2018, cardiovascular diseases contributed to 31% of the deaths, while cancer accounted for 26%. Dementia is third at 10%. Because the majority of deaths are related to (the absence of) health, we focus on choices related to health to understand how culture could influence survival rate. Specifically, we investigate choices related to the healthcare system and the management of pandemics.

Shifting to fertility and drawing on Hrdy's work on the evolutionary basis of parenthood (1999), we focus on women's freedom of choice regarding investments in offspring. Human infants are highly resource-intensive, and raising a human child requires cooperation among multiple caregivers. Humans are cooperative breeders, and presumably women have long been subject to selection for the ability to assess the social support available for raising a child. If adequate support is lacking, women may choose not to invest in the child and prioritise potential future offspring instead. In terms of genetic fitness, women need the freedom to manage trade-offs between investing in current offspring versus conserving resources for potential future offspring. In that sense, cultural practices that limit women's autonomy could be viewed as imposing a detrimental effect on the fitness of women who have not completed reproduction and on the inclusive fitness of any genetic relatives. We investigate potential differences in support for three types of policy that should influence women's freedom of choice and degree of social support during and after pregnancy. These three types of policy pertain to abortion access, assisted reproduction, and paid parental leave.

1.1.4 Discontinuities in Voting Behaviours and Fitness Costs

We would like to explain the generic argument for why discontinuities in these voting behaviours are interesting from a gene-culture coevolutionary perspective. First, we select referenda regarding health and reproduction policies that affect fitness through their implications on mortality and fertility. Some policies might favour more children and other policies fewer children. This would mean, in turn, that policies, if enacted, would vary in terms of how they incentivise individuals to manage the trade-offs between the quantity and quality of their offspring. Analogously, some policies might augment the scope for individuals to rely on social support when raising offspring, while other policies might do the opposite. In this way, if enacted, policies would vary in terms of how they incentivise individuals to manage the trade-offs between current and future offspring. Lastly, policies related to pandemics should affect the risk of infectious disease and by extension the risk of mortality. Policies related to healthcare more broadly should affect the extent to which individuals invest in their health and in turn survival. For example, one of the referenda below concerned how to organise health insurance. Even if we imagine that the alternatives would have no consequences in terms of the quality of healthcare supplied, we can easily imagine that different insurance schemes would affect behaviour on the demand side. Some schemes might incentivise healthy lifestyles and preventative treatments, while other schemes might tip the balance in favour of treating people after they get sick.

Second, we estimate potential differences in voting behaviour at the language border. Right at the border, we assume that, among the policies under consideration, one policy is better than the other on average in terms of expected fitness. By this, we do not mean that one policy is best or optimal in absolute terms. Rather, we mean that, between the policies for which citizens vote, one is better than the other. We call this the "better" policy, and our working assumption is that this better policy is the same on both sides of the border. This is a fundamental assumption for our approach. The assumption might be wrong, of course, but focusing on discontinuities right at the border maximises the chances that it is correct. In any case, our task is to examine both the implications and the limitations of this assumption.

This assumption does not mean that the same policy is better than the other throughout all of French- and German-speaking Switzerland. It simply means that the better policy is the same immediately to the west and immediately to the east of the border. In addition, this assumption does not mean that one policy is better for the expected fitness of every individual. Instead, we assume that one is better than the other on average in terms of the expected fitness of individuals in the group. We do not deny individual heterogeneity. Instead, by using the regression discontinuity design and the linguistic border, we aggregate individual differences and focus on the average outcome at the group level.

We do not know which policy is better, nor does the answer to this question matter for present purposes. We simply assume that at the border one is better than the other in terms of average expected fitness. If this is true, then a discontinuity implies that at least one of the two groups does not favour the best of the two policies for cultural reasons, where cultural reasons, by this account, must be separate from institutions, genes, and environment. If enacted, the inferior policy would bring an expected fitness cost, however small, on some individuals of the group relative to the other policy under consideration.

Nonetheless, one can challenge the assumption that right at the border one policy is better than the other in fitness terms. We would like to highlight two possibilities. First, in high-dimensional choice spaces with a complex fitness topography, multiple optima can easily exist. Two groups can thus favour two different policies, both of which are roughly equivalent local optima. The two optima in question may or may not be globally optimal. Regardless, the point is that the two policies differ in the details, but they are extremely similar in terms of ultimate outcomes. Second, by examining discontinuities at the border, our approach controls for institutional, geographic, and genetic variation as potential confounds. It does not, however, control for all sources of social variation. People living on the two sides of the border may have different social networks, which could lead the value of a given policy to vary as we move from one side of the border to the other. Such scenarios could challenge our assumption that right at the border, on both sides of the border, one single policy is better than the other in terms of average fitness. We cannot definitively rule out such scenarios, but we should reduce the probability they play a large role precisely because we limit attention to discontinuities.

In sum, our study aims to investigate the causal influence of culture on health- and fertility-related choices and to discuss how any differences might relate to genetic fitness. To meet this goal, we use a quasi-experimental design based on distance from the Röstigraben, a linguistic and cultural border in Switzerland. We are looking for discontinuities in choices at the border. Any discontinuities at the border would suggest a cleanly identified cultural difference that shapes preferences and behaviour. We will then discuss, somewhat speculatively, how these cultural differences could affect the relative fitness of individuals in the two cultural groups.

1.2 Methods

1.2.1 Referanda Data

To explore potential cultural differences in decision-making domains related to fertility and mortality, we use data from referenda in Switzerland. The use of a regression discontinuity design necessitates a substantial amount of geographically precise data, which the referenda data provide. We focus on referenda held at the Swiss level and thus common to all cantons. Importantly, referenda occur multiple times a year and encompass a wide range of topics, including health, the healthcare system, and fertility. However, our sample represents only the voting population and excludes non-voters' opinions on both sides of the border. Nonetheless, the laws are based on the decisions of voters. As such, even though our data are not fully representative of the entire Swiss population, they can help identify cultural differences in the voting population. Further, our data are aggregated at the municipal level, rather than at the individual level, presenting a notable limitation in assessing the impact of cultural differences on individual fitness within the two groups. A more direct assessment would involve individual-level data. However, the requirements of our study for large datasets with precise geographical accuracy, combined with the sensitive nature of voting, health, and fertility data, ensure that access to individual-level data is strictly limited. Therefore, we have employed municipal-level data as a feasible and effective solution.

We use the percentage of "yes" votes by municipality in referenda as our response variables, and we estimate discontinuities in referenda results across municipalities on both sides of the border. Our analysis focuses on a preregistered list of referenda related to health or fertility in the past decade (Faessler et al. 2022). The data are provided by the Federal Statistical Office and include referenda results across municipalities, with our unit of analysis being the municipality. We selected municipalities within 100 km of the language border, totalling 1,409 municipalities.

1.2.2 Regression Discontinuity Design

A regression discontinuity design has three essential elements: a threshold, a running variable, and a treatment. In our case, the threshold is the cultural border, the continuous variable is the distance from this border, and the treatment is the culture. Starting from these elements, we estimate two regression lines on each side of the border to examine whether voting results are discontinuous at the border. As such, we study the effect of moving from one side of the language and cultural boundary to the other on referenda outcomes and the distribution of policy preferences these outcomes represent.

The generic regression model for these regression discontinuity designs can be represented as follows.

$$y_m = \beta_0 + \beta_1 German_m + \beta_2 f_0(Distance_m) + \beta_3 German_m * f_1(Distance_m) + controls + \epsilon_m \quad (1.1)$$

In detail, y_m denotes the outcome of interest for municipality m, which is the proportion of "yes" votes for a referendum. $German_m$ is a dummy variable that takes the value of 1 if the municipality is on the German side of the border and 0 otherwise. In that sense, β_1 captures the discontinuity of interest, the cultural discontinuity at the border. A significant β_1 value indicates a causal effect of culture on voting decisions at the border. $Distance_m$ is the running variable that measures the distance from the border. $f_0()$ and $f_1()$ are functions of distance to the border that will be estimated. Both $Distance_m$ and its interaction with $German_m$ take care of controlling for effects that happen away from the border and that could be driven by environmental differences. Throughout the study, we will estimate different versions of this generic regression discontinuity model, each of which will focus on a distinct referendum.

In this analysis, municipality language $German_m$ and distance from the language border $Distance_m$ are our main independent variables. Distance from the border, in particular, plays a crucial role, and we explain in detail how the measure is constructed. First, using the same distance data as Eugster et al. (2011), each municipality is assigned a language according to the language spoken by most of its population. Second, the distance to the language border is calculated by determining the shortest road distance between the focal municipality and the nearest municipality where the other language is spoken. Further, the distance is set as negative for French-speaking municipalities and positive for German-speaking municipalities.

Our statistical model controlled for municipality type because rural and urban areas could exhibit different voting patterns. We control for this possibility by including a dummy for municipality type, i.e. whether the municipality is located in an urban or rural area. We also include canton fixed effects. In Switzerland, a federal system divides power between the state and the cantons. Cantons are administrative subdivisions of the country and have authority over education, health care, policing, and taxation. In particular, institutions related to health and fertility may vary across cantons. By incorporating a canton fixed effect, we address disparities among cantons and restrict our analysis to variations within each canton. Nevertheless, the language border crosses some cantons and does not correspond to an institutional boundary for this part.

A fundamental assumption of regression discontinuity design is that at the threshold, the treated and control groups differ only by treatment. Because our unit of analysis is the municipality, we necessarily move from one municipality to another at the threshold. However, while municipalities have a certain degree of autonomy, their powers are limited by cantonal and federal laws. Municipalities are mainly responsible for local governance, including waste management, water supply, social welfare and public transport. Thus, even though institutions change from one municipality to another, the institutional changes are limited and not directly related to health and fertility.

	Mean all	French l.	German l.	Difference	At the border
Population size	3163.38	2609.39	3576.18	966.79	-796.118
Population variation (%)	8.83	13.36	5.45	-7.91^{***}	-6.784^{***}
Density	323.73	203.00	413.69	210.69^{***}	126.498
Immigrants (%)	14.19	16.14	12.74	-3.40^{***}	-5.219^{***}
Average household size	2.31	2.37	2.27	-0.10^{***}	-0.016
0-19 years (%)	20.45	21.97	19.32	-2.65^{***}	-2.184^{***}
20-64 years (%)	59.74	59.75	59.73	-0.02	0.303
+65 years (%)	19.81	18.28	20.95	2.67^{***}	1.882^{***}
Young dependency ratio	34.52	36.81	32.81	-4.00^{***}	-3.571^{***}
Mean taxable revenue	69,762	66,617	72,275	$5{,}658^{***}$	$6{,}079^{*}$
Tax rate for families	5.30	4.94	5.57	0.63^{***}	-0.182^{***}
Tax rate for singles	15.40	15.33	15.45	0.12	-0.373^{***}
Social assistance (%)	2.69	2.80	2.61	-0.19	-1.087^{***}

Table 1.1: Municipalities and population characteristics around the border. Notes: “Mean all” refers to the mean of municipalities within 50 km of the language border. “French language” includes only the municipalities where most of the population speaks French, within a 50km range. “German language”, the municipalities where most of the population speaks German, within a 50km range. “Difference” shows the mean difference between French-language municipalities and German-language municipalities. “At the border” shows the difference estimated at the language border using regression discontinuity design and controlling canton, and whether the municipality is urban or rural. $^{\dagger}p<0.1$; $^{*}p<0.05$; $^{**}p<0.01$; $^{***}p<0.001$. Source: Swiss Federal Statistical Office (SFO). Distances from search.ch.

Aside from institutions, population characteristics may also vary at the border. Table 1.1

provide the statistics for a selection of population and municipality variables likely to influence choices related to health and fertility. The variables include population size and characteristics, age structure within the population, and a series of wealth indicators. Most of the variables are not perfectly balanced at the border, but regions are more balanced at the border than overall (column "Difference" has larger differences than "At the border"). In particular, the age structure and wealth seem to differ on the border's two sides. The municipalities on the German-speaking side count more older individuals and fewer younger individuals while having higher revenues and smaller tax ratios. These differences could suggest that the population on the German-speaking side of the border is more preoccupied with health, but it also benefits from higher revenues to prevent disease or provide medical care.

To ensure that our results are not influenced by these disparities across French and German-speaking municipalities, we incorporate additional municipality-level controls in a robustness analysis (presented as model (4) in the results regression tables). Specifically, we account for age structure differences by introducing the following variables: the proportion of individuals below 19 years old, those exceeding 64 years old, the youth dependency ratio, the birth rate, and the average household size. Additionally, we address wealth disparities by integrating the average taxable revenue and the tax rates for families and singles. Because we do not know if some of these variables are influenced by culture, we do not treat the analysis that includes these variables as our baseline analysis. Instead, we treat the analysis with these additional variables as a robustness validation, even though some of the variables could be colliders.

Importantly, our setting cannot exclude that some individuals decide to move to the other side of the border. If so, people would self-select their treatments, which would undermine to some extent our identification strategy. While people could decide to live in the region that best matches their values, the language border is sharp. Our data indicate that the mean proportion of French speakers shifts from 74% to 12% within a distance of only 6 km. Similarly, the proportion of German speakers shifts from 24% to 86%. Moving to another linguistic region would require the individual to learn the other language with a fluency level comparable to speaking a language at home, which necessarily constitutes a barrier. Further, the average moving distance for Switzerland is 13 km, and most of the moves (58%) happen within a distance of 5km (OFS 2022*b*). Although we cannot exclude that some individuals self-select in treatments, we suspect this mechanism has limited effects.

1.3 Results

Our results show multiple discontinuities in voting behaviours at the language border. Of those significant discontinuities, three are related to fertility, and one is related to health. These discontinuities suggest the potential for culture to create stable differences between groups in domains related to health and fertility, where cultural explanations are distinct from institutional, genetic, and environmental explanations. For each of the referenda, we further speculate how these cultural differences could affect the relative fitness values of individuals in the two cultural groups.

1.3.1 Health-related Referanda

28 September 2014, the referendum for a single public health insurance company. First, we analysed the results of the referendum on creating a single public health insurance company, which took place on 28 September 2014. Under the proposed single-payer system, a public insurance company would have replaced the current private insurance companies, and all residents would have been required to enrol in the public plan. Supporters argued that the single-payer system would reduce administrative costs and improve access to healthcare. At the same time, opponents claimed that it would lead to longer waiting times and lower quality of care.

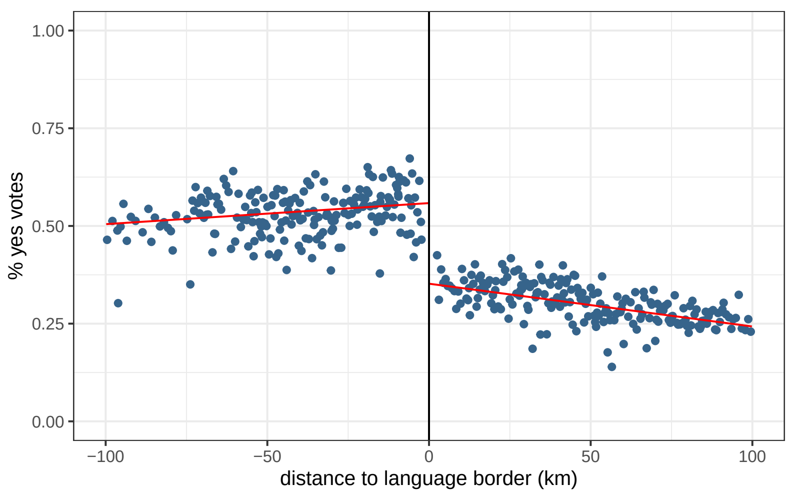

Figure 1.1: Average proportion of “yes” votes to the referendum for a single public health insurance, by distance to the language border. Notes: The left-hand side of the graph displays French-speaking municipalities; the right-hand side, German-speaking municipalities. The red lines are the linear regression lines. Source: Federal Statistical Office. Distances from search.ch.

Figure 1.1 shows a strong discontinuity at the border in the pattern of "yes" votes proportions across municipalities. The left-hand side of the graph displays French-speaking municipalities, while the right-hand side shows German-speaking municipalities. In almost all municipalities on the French-speaking side of the border, the proportion of "yes" votes is higher than in municipalities on the German-speaking side. The red lines represent linear regression lines. Linear regression results in Table 1.2 confirm the presence of a discontinuity in voting results at the border (estimate $= -0.221$, $p < 0.001$). Further, the German language estimate is not sensitive to controlling for additional municipality-level controls.

	(1)	(2)	(3)	(4)
			Baseline	
German Language	-0.199^{***}	-0.219^{***}	-0.221^{***}	-0.209^{***}
	(0.009)	(0.009)	(0.009)	(0.009)
German*Distance	-0.002^{***}	-0.002^{***}	-0.002^{***}	-0.002^{***}
	(0.0002)	(0.0002)	(0.0002)	(0.0002)
Distance	0.001^{***}	0.001^{***}	0.001^{***}	0.001^{***}
	(0.0001)	(0.0001)	(0.0002)	(0.0001)
Urban			0.022^{***}	0.021^{***}
			(0.004)	(0.004)
Constant	0.558^{***}	0.574^{***}	0.554^{***}	0.766^{***}
	(0.007)	(0.016)	(0.016)	(0.056)
Cantons FE	No	Yes	Yes	Yes
Municipality Controls	No	No	No	Yes
Observations	1,409	1,409	1,353	1,353
Adjusted R^2	0.662	0.804	0.807	0.819

Table 1.2: Referenda for a single public health insurance company: regression analysis at the language border. Notes: The regression analysis shows the impact of switching from the French-speaking side of the border to the German-speaking side on voting results, that is, the proportion of "yes" votes in a municipality. "German language" indicates that the primary language of a municipality is German and is our variable of interest. "Distance" is the road distance to the language border. "Distance" and its interaction with "German language" control for effects that happen away from the border and environmental differences. We restrict our analysis to municipalities within 100km of the language border. Models (2) and (3) include controls for the canton. Model (3) includes a control variable for municipality characteristics, whether the municipality is located in a rural or urban area. Model (4) includes additional controls at the municipality level. Controls include population age structure, average household size, birth rates, average revenue, and tax rates. Robust standard errors are in parenthesis. $^{\dagger}p<0.1$; $^{*}p<0.05$; $^{**}p<0.01$; $^{***}p<0.001$. Source: Federal Statistical Office. Distances from search.ch.

The referendum on a single health insurance company highlights an interesting example of the potential influence of culture on fitness. Swiss citizens were asked if they would like a single public health insurance system or multiple private health insurance companies. To illustrate the significance of this choice, imagine two extremes. At one extreme, a single insurance company would pool risk over the entire Swiss population. At the other extreme, each individual would self-insure and be responsible for her own healthcare and associated costs. Whatever the details, the best system in terms of an individual's health, survival, and fitness for the people at the border must lie between these two extremes. We observed that the two groups supported different policies at the border. If, however, the better policy on average is the same right at the border, the discontinuity in preferences at the border means that at least one of the two groups supported, for cultural reasons, a worse policy in terms of expected fitness compared to the other policy.

22 September 2013, revision of the law on epidemics and 13 June 2021, Covid law. The second example comes from two referenda related to the management of epidemics. The two referenda are 8 years apart. On 22 September 2013, Switzerland held a first referendum on revising the law on epidemics, and the proposed changes aimed to enhance the country's response to any future pandemics. The revised law would have expanded the government's powers to contain outbreaks, require vaccinations, and collect health data for public health reasons. However, groups such as anti-vaxxers and privacy advocates were concerned about the increased surveillance and data collection that could follow. Eight years later, on 13 June 2021, Swiss citizens voted on a related question, namely the Covid law. The proposal was to give the government extraordinary powers to manage the Covid-19 pandemic, such as imposing restrictions on public life and providing financial aid to those affected. However, the law faced opposition from groups who believed it gave the government too much power and infringed on individual freedoms. A majority vote of around 60% approved both laws.

(a) The revision of the epidemics law (b) The Covid law

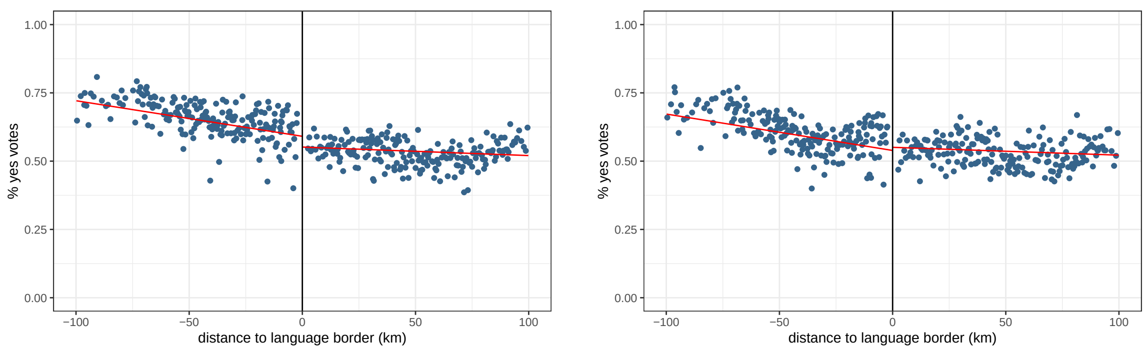

Figure 1.2: Average proportion of "yes" votes to the two referenda on epidemics management across municipalities, by distance to the language border. Notes: The left-hand side of the graph displays French-speaking municipalities; the right-hand side, German-speaking municipalities. The red lines are the linear regression lines. Source: Federal Statistical Office. Distances from search.ch.

Figure 1.2 plots the average proportion of "yes" votes for these two referenda across municipalities on the two sides of the language border. The two figures present similar patterns, namely a negative slope on both sides and a steeper slope on the French side. However, these two graphs by themselves do not allow us to confirm or disconfirm the presence of discontinuities at the border.

	(1)	(2)	(3) Baseline	(4)
German Language	-0.050^{***}	-0.014	$-0.016^{\dagger}$	-0.026^{**}
	(0.009)	(0.010)	(0.010)	(0.009)
German*Distance	0.001^{***}	-0.0003	$-0.0004^{\dagger}$	-0.0005^{*}
	(0.0002)	(0.0002)	(0.0002)	(0.0002)
Distance	-0.001^{***}	-0.001^{***}	-0.001^{***}	-0.0004^{*}
	(0.0001)	(0.0002)	(0.0002)	(0.0002)
Urban			0.062^{***}	0.048^{***}
			(0.004)	(0.004)
Constant	0.592^{***}	0.695^{***}	0.639^{***}	0.826^{***}
	(0.007)	(0.018)	(0.018)	(0.060)
Cantons FE	No	Yes	Yes	Yes
Municipality Controls	No	No	No	Yes
Observations	1,409	1,409	1,353	1,353
Adjusted R^2	0.382	0.525	0.594	0.635

Table 1.3: Revision of the epidemics law: regression analysis at the language border. Notes: The regression analysis shows the impact of switching from the French-speaking side of the border to the German-speaking side on voting results, that is, the proportion of "yes" votes in a municipality. "German language" indicates that the primary language of a municipality is German and is our variable of interest. "Distance" is the road distance to the language border. "Distance" and its interaction with "German language" control for effects that happen away from the border and environmental differences. We restrict our analysis to municipalities within 100km of the language border. Models (2) and (3) include controls for the canton. Model (3) includes a control variable for municipality characteristics, whether the municipality is located in a rural or urban area. Model (4) includes additional controls at the municipality level. Controls include population age structure, average household size, birth rates, average revenue, and tax rates. Robust standard errors are in parenthesis. $^{\dagger}p < 0.1$; $^{*}p < 0.05$; $^{**}p<0.01$; $^{***}p<0.001$. Source: Federal Statistical Office. Distances from search.ch.

Tables 1.3 and 1.4 present the results of the regression analyses. Both results show quantitatively small estimates whose significance varies across models. In 2013, the municipalities on the German-speaking side of the border were less likely to accept the law, but the difference is not significant in the baseline model (estimate $= -0.016$, $p < 0.1$). In 2021, the effect goes in the opposite direction. Municipalities on the German-speaking side of the border are more likely to vote "yes" (estimate $= 0.027$, $p < 0.05$). However, the significance disappears in model (4) with the addition of controls at the municipality level. These results should be interpreted with caution, and we treat them as neither significant nor robust. Further analyses are needed to understand whether culture can influence pandemic-related behaviours.

	(1)	(2)	(3)	(4)
			Baseline	
German Language	-0.0005	0.029^{*}	0.027^{*}	0.013
	(0.010)	(0.012)	(0.011)	(0.011)
German*Distance	0.001^{***}	-0.0003	$-0.0004^{\dagger}$	-0.001^{*}
	(0.0002)	(0.0003)	(0.0002)	(0.0002)
Distance	-0.002^{***}	-0.0002	-0.0002	0.0001
	(0.0002)	(0.0002)	(0.0002)	(0.0002)
Urban			0.094^{***}	0.079^{***}
			(0.005)	(0.005)
Constant	0.541^{***}	0.617^{***}	0.528^{***}	0.746^{***}
	(0.008)	(0.022)	(0.020)	(0.069)
Cantons FE	No	Yes	Yes	Yes
Municipality Controls	No	No	No	Yes
Observations	1,409	1,409	1,353	1,353
Adjusted R^2	0.186	0.272	0.440	0.489

Table 1.4: Revision of the Covid law: regression analysis at the language border. Notes: The regression analysis shows the impact of switching from the French-speaking side of the border to the German-speaking side on voting results, that is, the proportion of "yes" votes in a municipality. "German language" indicates that the primary language of a municipality is German and is our variable of interest. "Distance" is the road distance to the language border. "Distance" and its interaction with "German language" control for effects that happen away from the border and environmental differences. We restrict our analysis to municipalities within 100km of the language border. Models (2) and (3) include controls for the canton. Model (3) includes a control variable for municipality characteristics, whether the municipality is located in a rural or urban area. Model (4) includes additional controls at the municipality level. Controls include population age structure, average household size, birth rates, average revenue, and tax rates. Robust standard errors are in parenthesis. $^{\dagger}p < 0.1$; $^{*}p < 0.05$; $^{**}p<0.01$; $^{***}p<0.001$. Source: Federal Statistical Office. Distances from search.ch.

1.3.2 Fertility-related Referanda

9 February 2014, referendum prohibiting the reimbursement of abortion. We now provide three examples related to fertility. We start with the referendum on the reimbursement of abortion. On 9 February 2014, Swiss citizens voted on the prohibition of the reimbursement of abortion by health insurance companies. Proponents of the proposal argued that taxpayers should not be forced to pay for a procedure they consider morally objectionable. Conversely, opponents argued that women should have access to safe and affordable abortion services, regardless of their financial situation.

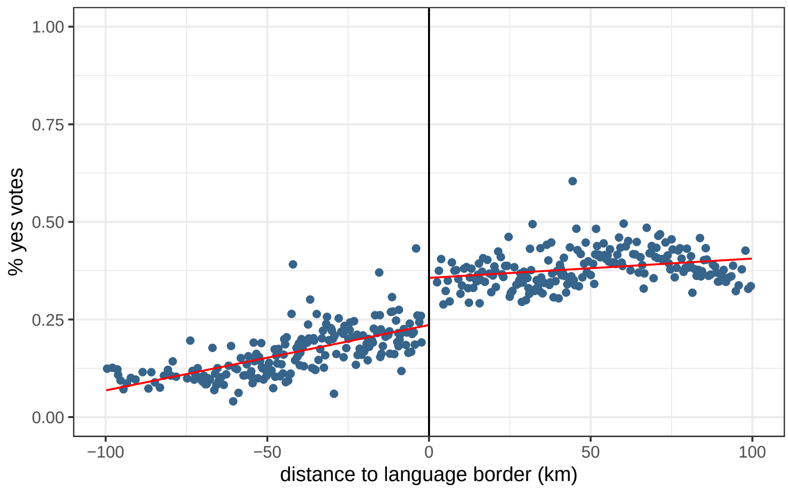

Figure 1.3: Average proportion of "yes" votes to referendum prohibiting the reimbursement of abortion across municipalities, by distance to the language border. Notes: The left-hand side of the graph displays French-speaking municipalities; the right-hand side, German-speaking municipalities. The red lines are the linear regression lines. Source: Federal Statistical Office. Distances from search.ch.

Figure 1.3 presents the percentage of votes in favour of the initiative across municipalities at different distances of the language border. The data show an evident discontinuity at the border. Municipalities on the French-speaking side of the border were less likely to vote in favour of modifying the law than municipalities on the German-speaking side. Regression analysis results in Table 1.5 confirm these descriptive results. The German language estimate is significant in the four models, and adding controls does not change this in any way (estimate = 0.127, $p < 0.001$).

Restricting women's access to abortion could have considerable genetic fitness implications, particularly for women. As cooperative breeders, mothers, and by extension fathers, require social support to raise their children. They must balance investment in their current offspring with investment in potential future offspring (Hrdy 1999). In that sense, any restrictions on access to abortion would limit women's ability to manage this trade-off and impose a fitness

	(1)	(2)	(3) Baseline	(4)
German Language	0.122***	0.125***	0.127***	0.132***
	(0.008)	(0.009)	(0.009)	(0.009)
German*Distance	−0.001***	0.0003	0.0003†	0.0005*
	(0.0002)	(0.0002)	(0.0002)	(0.0002)
Distance	0.002***	0.0005**	0.0005**	0.0003†
	(0.0001)	(0.0002)	(0.0002)	(0.0002)
Urban			−0.041***	−0.029***
			(0.004)	(0.004)
Constant	0.236***	0.133***	0.174***	−0.041
	(0.006)	(0.017)	(0.017)	(0.058)
Cantons FE	No	Yes	Yes	Yes
Municipality Controls	No	No	No	Yes
Observations	1,409	1,409	1,353	1,353
Adjusted R^2	0.676	0.756	0.770	0.786

Table 1.5: Referendum prohibiting the reimbursement of abortion: regression analysis at the language border. Notes: The regression analysis shows the impact of switching from the French-speaking side of the border to the German-speaking side on voting results, that is, the proportion of “yes” votes in a municipality. “German language” indicates that the primary language of a municipality is German and is our variable of interest. “Distance” is the road distance to the language border. “Distance” and its interaction with “German language” control for effects that happen away from the border and environmental differences. We restrict our analysis to municipalities within 100km of the language border. Models (2) and (3) include controls for the canton. Model (3) includes a control variable for municipality characteristics, whether the municipality is located in a rural or urban area. Model (4) includes additional controls at the municipality level. Controls include population age structure, average household size, birth rates, average revenue, and tax rates. Robust standard errors are in parenthesis. $^{\dagger}p<0.1$; $^{*}p<0.05$; $^{**}p<0.01$; $^{***}p<0.001$. Source: Federal Statistical Office. Distances from search.ch.

cost on women. Acknowledging men's commitment to their offspring, these constraints and costs might not pertain to women only, but in some cases may extend to the couple. The 2014 referendum prohibiting the reimbursement of abortion in Switzerland could have resulted in such a cost, given the potential restrictions on access that the initiative could have imposed. Assuming that at the border one policy is better than the other in terms of average fitness, the discontinuity in the voting results suggests that one group was more willing to support a policy that would presumably impose an additional fitness cost on some individuals in the population relative to the other policy under consideration.

5 June 2016, referendum on assisted reproduction. On 5 June 2016, Swiss citizens voted to modify the medically assisted reproduction law. The proposed amendment aimed to legalize, under certain conditions, the genetic diagnosis of embryos derived from in vitro fertilization before implanting the embryos. The amended law would have allowed pre-implementation diagnosis only for carriers of alleles associated with severe hereditary disease or those who cannot have a child naturally. Supporters argued that the law was necessary to provide couples with the same reproductive options already available in neighbouring countries. On the other hand, opponents feared that the revision would have led to an ethically unacceptable expansion of genetic testing on human embryos and undermined the traditional family structure.

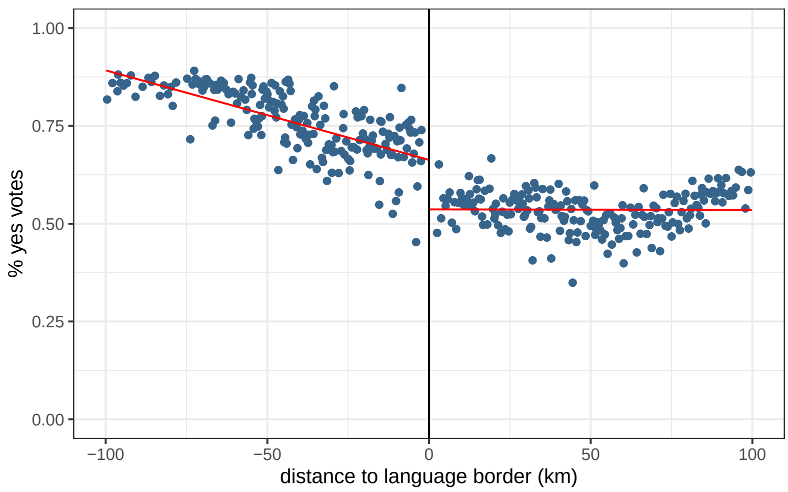

Figure 1.4: Average proportion of "yes" votes to referendum allowing genetic diagnosis of embryos, across municipalities, by distance to the language border. Notes: The left-hand side of the graph displays French-speaking municipalities; the right-hand side, German-speaking municipalities. The red lines are the linear regression lines. Source: Federal Statistical Office. Distances from search.ch.

Figure 1.4 shows the average proportion of "yes" votes across municipalities at various distances from the language border. Data present a clear discontinuity at the border. Further, most data points on the French-speaking side of the border are above the data points on the German-speaking side. At the border, the French-speaking group is more likely to favour amending the law than the German-speaking group. These results are confirmed by the regression analysis results presented in Table 1.6. The German language estimate is significant and robust to additional municipality-level controls (estimate = −0.097, $p < 0.001$).

	(1)	(2)	(3) Baseline	(4)
German Language	-0.131^{***}	-0.099^{***}	-0.097^{***}	-0.109^{***}
	(0.010)	(0.010)	(0.010)	(0.009)
German*Distance	0.002^{***}	-0.001^{*}	-0.001^{**}	-0.001^{***}
	(0.0002)	(0.0002)	(0.0002)	(0.0002)
Distance	-0.002^{***}	-0.0004^{*}	-0.0003^{*}	-0.0001
	(0.0002)	(0.0002)	(0.0002)	(0.0002)
Urban			0.050^{***}	0.033^{***}
			(0.004)	(0.004)
Constant	0.664^{***}	0.809^{***}	0.758^{***}	0.998^{***}
	(0.007)	(0.018)	(0.018)	(0.059)
Cantons FE	No	Yes	Yes	Yes
Municipality Controls	No	No	No	Yes
Observations	1,409	1,409	1,353	1,353
Adjusted R^2	0.638	0.772	0.792	0.820

Table 1.6: Referendum on assisted reproduction: regression analysis at the language border.
Notes: The regression analysis shows the impact of switching from the French-speaking side of the border to the German-speaking side on voting results, that is, the proportion of "yes" votes in a municipality. "German language" indicates that the primary language of a municipality is German and is our variable of interest. "Distance" is the road distance to the language border. "Distance" and its interaction with "German language" control for effects that happen away from the border and environmental differences. We restrict our analysis to municipalities within 100km of the language border. Models (2) and (3) include controls for the canton. Model (3) includes a control variable for municipality characteristics, whether the municipality is located in a rural or urban area. Model (4) includes additional controls at the municipality level. Controls include population age structure, average household size, birth rates, average revenue, and tax rates. Robust standard errors are in parenthesis. $^{\dagger}p<0.1$; $^{*}p<0.05$; $^{**}p<0.01$; $^{***}p<0.001$. Source: Federal Statistical Office. Distances from search.ch.

The outcome of the 5 June 2016 referendum on pre-implantation genetic diagnosis could have had fitness consequences at the individual level. By allowing couples with serious hereditary diseases to implant healthy embryos selectively, the legalisations of pre-implantation diagnosis could have increased their offspring's chances of survival and reproduction, ultimately positively impacting individual fitness. However, genetic screening implies an opportunity cost. Using genetic screening for non-medical reasons, such as selecting specific traits such as eye colour or height, could result in a waste of resources. Unnecessary screening might divert limited resources away from other procedures that could matter more in terms of health. We do not know what screening level maximised individual fitness in that particular environment. Nonetheless, we observed that at the border the two groups had different preferences and associated voting behaviours. Assuming that at the border one policy is better than the other in terms of average fitness, supporting one policy would presumably impose a fitness cost on individuals relative to the other policy.

27 September 2020, referendum on paternity leave. Our last example focuses on paternity leave. On 27 September 2020, Swiss citizens had to decide whether fathers should be granted two weeks of paid paternity leave. The proposed amendment to the Swiss Federal Constitution aimed to give fathers the right to take two weeks off work after the birth of a child. This leave would have been financed by the government. Proponents of the amendment argued that paternity leave would have provided fathers with the opportunity to bond with their newborns and help reduce gender inequality in the workplace and society. On the other hand, opponents claimed that the proposed paternity leave policy would have increased costs for employers and should not be legislated at the federal level.

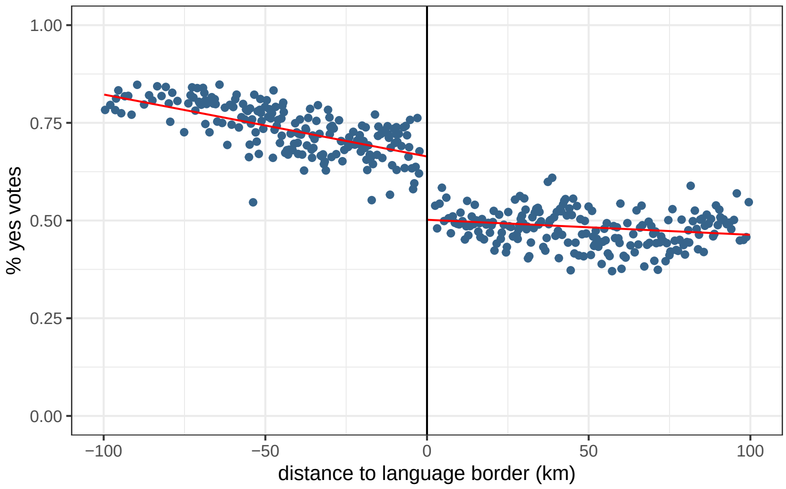

Figure 1.5: Average proportion of "yes" votes to referendum on paternity leave across municipalities, by distance to language border. Notes: The left-hand side of the graph displays French-speaking municipalities; the right-hand side, German-speaking municipalities. The red lines are the linear regression lines. Source: Federal Statistical Office. Distances from search.ch.

Figure 1.5 presents the average proportion of "yes" votes for the referendum on paid paternity leave across municipalities at different distances from the language border. We observe an apparent discontinuity at the language border. Municipalities on the French-speaking side of the border were more likely to approve a paid paternity leave than those on the German-speaking side. Table 1.7 presents the regression analysis. The results confirm the descriptive evidence from the graph. The German language estimate is significant and not sensitive to additional controls (estimate $= -0.160$, $p < 0.001$).

	(1)	(2)	(3) Baseline	(4)
German Language	-0.171^{***}	-0.155^{***}	-0.160^{***}	-0.159^{***}
	(0.009)	(0.010)	(0.009)	(0.010)
German*Distance	0.001^{***}	-0.0003	-0.0003	-0.001^{*}
	(0.0002)	(0.0002)	(0.0002)	(0.0002)
Distance	-0.002^{***}	-0.001^{***}	-0.001^{**}	$-0.0003^{\dagger}$
	(0.0001)	(0.0002)	(0.0002)	(0.0002)
Urban			0.067^{***}	0.058^{***}
			(0.004)	(0.004)
Constant	0.665^{***}	0.769^{***}	0.704^{***}	0.843^{***}
	(0.007)	(0.019)	(0.018)	(0.061)
Cantons FE	No	Yes	Yes	Yes
Municipality Controls	No	No	No	Yes
Observations	1,409	1,409	1,353	1,353
Adjusted R^2	0.715	0.773	0.811	0.822

Table 1.7: Referendum on paternity leave: regression analysis at the language border. Notes: The regression analysis shows the impact of switching from the French-speaking side of the border to the German-speaking side on voting results, that is, the proportion of “yes” votes in a municipality. “German language” indicates that the primary language of a municipality is German and is our variable of interest. “Distance” is the road distance to the language border. “Distance” and its interaction with “German language” control for effects that happen away from the border and environmental differences. We restrict our analysis to municipalities within 100km of the language border. Models (2) and (3) include controls for the canton. Model (3) includes a control variable for municipality characteristics, whether the municipality is located in a rural or urban area. Model (4) includes additional controls at the municipality level. Controls include population age structure, average household size, birth rates, average revenue, and tax rates. Robust standard errors are in parenthesis. $^{\dagger}p<0.1$; $^{*}p<0.05$; $^{**}p<0.01$; $^{***}p<0.001$. Source: Federal Statistical Office. Distances from search.ch.

Paternity leave may have had positive fitness consequences. Paternity leave allows fathers to spend more time with their newborn children. The more the father invests, the better the outcomes should tend to be for the current offspring. However, we could also imagine a countervailing effect for men. By investing time and resources in current offspring, fathers are potentially hindering their careers, which could make them less attractive in the future. Consequently, fathers are potentially hindering their ability to identify opportunities to mate with other women. In this sense, paternity leave could partially harm fathers’ fitness. We observe that the two groups adopted different voting behaviours at the border. Assuming the acceptance of paternity leave has fitness consequences and that at the border the optimal policy was the same, then one group showed stronger support for a policy that would have imposed fitness costs on some individuals compared to the other policy.

1.4 Discussion

We have investigated the causal influence of culture on health- and fertility-related choices using a spatial regression discontinuity design and Swiss referenda data. Our results show multiple discontinuities at the language border, especially with regard to fertility. Such discontinuities isolate cultural variation in preferences for policies that, if enacted, would have presumably affected health and fertility choices at the individual level. We have also speculated about connections between possible referenda outcomes and downstream effects on genetic fitness. Although the details of these speculations differ, the generic logic is always the same. For a given referendum, assume that one policy was better than the other policy in the sense that it would have promoted choices and created incentives that would have been better – in terms of individual expected fitness. We do not know which policy was better in this sense, but we assume that one was better, and the other was worse. If, in addition, the better policy right at the border was the same on both sides of the border, then any discontinuity in voting at the border implies that one of the two groups showed relative support for the worse policy for cultural reasons. More to the point, one of the two groups supported a policy that would have negatively affected health, survival, and fertility relative to the other policy. By extension, the individuals in this group were ready to pay an opportunity cost in terms of fitness, and they were willing to impose this fitness cost on their Swiss fellows who would have been subject to the policy if enacted. We can view this opportunity cost in two ways. First, it would have represented an opportunity cost relative to the other policy under consideration. Second, it would have represented an opportunity cost in the form of reduced fitness relative to other societies, for example other countries in continental Europe.

While our findings emphasise cultural differences in health- and fertility-related voting decisions at the language border, our study comes with several limitations. First, our study employs municipal-level data and not individual-level data. This approach is well-suited for the central part of our analysis. We effectively demonstrate the capacity of culture to create stable differences between groups in domains related to health and fertility. However, this approach presents a significant limitation in exploring the potential impact of these cultural differences on variation in individual fitness values. Future research would benefit greatly by using individual-level data to assess more accurately how cultural differences affect the individuals who make up the cultural groups under study.

Second, individuals could have, in principle, self-selected into treatments. People born on the French-speaking side of the border could have moved to the German-speaking region in search of a cultural environment more aligned with their personal values and vice versa. Although we suspect associated effects are trivial, we cannot definitively dismiss the potential impact of endogenous sorting into location at the border. Future research, equipped with more extensive data regarding the place of birth in lieu of the place of residence, would be better poised to control for any possible selection bias of this sort. Third, our sample consists solely of voters and is thus unrepresentative of the Swiss population. That said, laws and policies are enacted precisely on the basis of the preferences and decisions of voters, and in this sense our sample represents the politically engaged part of the population. As such, our data demonstrate how culture can shape voting decisions and policy outcomes.

Fourth, we do not know how cultural variation in voting translates into cultural variation in behaviour. For instance, we found clear distinctions in voting about paternity leave. Yet, we do not know how these kinds of differences might relate to the time fathers spend with their children, and we do not know how people on both sides of the border might react to one policy versus another. In general, we can imagine that the two groups might often support different policies, but they might also react differently to the policy that prevails after all the votes are tallied. Future research could examine these kinds of questions by exploring cultural differences in behavioural responses to political outcomes.

Finally, the data only pertain to referenda results and do not distinguish between the different reasons people vote one way or another. Our task was to isolate, as much as possible, the effects of culture from the effects of environments, institutions, and even genes. Our approach separates the influence of culture on voting in this way, but it cannot identify which components of culture drive results. Similarly, we cannot control for variation in the social environment. Observed variation at the border could be driven by differences in cultural domains related to religion, political affiliation, media consumption, or secular values. Future studies could unpack the discontinuities by investigating these kinds of underlying mechanisms.

Within the boundaries of these limitations, we have attempted to add a crucial element to the discussion of gene-culture processes by pushing for the clean identification of culture as a distinct cause of health- and fertility-related choices. In particular, genetic evolutionary processes do not favour stable differences between groups. Minimal gene flow between groups is enough to render groups nearly identical genetically (Frankham et al. 2002, Bell et al. 2009), and this seems to be the state of affairs at the Röstigraben in Switzerland (Buhler et al. 2012). This is crucial because, if groups are genetically similar, selection at the group level is irrelevant. If groups are different, in contrast, selection at the group level could easily matter. In this latter case, group selection can shape evolutionary dynamics in addition to selection at the individual level, and the result can be entirely new evolutionary regimes that would not otherwise be possible. Although the workaday evolutionary ecologist generally ignores such possibilities in strictly genetic systems, cultural evolutionary processes may be completely different (Mesoudi & Danielson 2008, Richerson et al. 2016). Our results show that cultural evolution can stabilise differences between groups, even amid ongoing contact, and it can do so in decision-making domains that should have a relatively close link to genetic fitness.

In particular, under the assumption that fitness effects are equivalent right at the border on both sides of the border, our results suggest that voters on one side or another routinely support a policy that was worse in terms of expected fitness than the other policy under consideration. The policy is worse in the sense that it should impose a cost in terms of expected fitness on individuals subject to the policy, but support for the policy is to some extent a group-level cultural phenomenon. This suggests the potential for cultures to maintain preferences detrimental to fitness when compared to some relevant benchmark. However, there are exceptions where this assumption does not hold; cases with complex fitness landscapes where multiple equivalent optima exist and variations in social environments that affect policy fitness consequences. These scenarios are crucial for a comprehensive understanding, yet they complement rather than contradict our primary observation. Cultural influences have the potential to shape preferences in ways that may not always align with optimal fitness outcomes.

These results are especially surprising because they hold in contemporary Switzerland. Switzerland is one of the easiest places in the world to get from one place to another. The distances are short, and the trains are clean, pleasant, frequent, extremely long, and exceedingly reliable. Moreover, this has been the state of affairs for a long time. The flow of cultural information across the border on a daily basis must be extreme, and thus one might naively expect the Röstigraben to be a cute vestige of former times. Our results, however, show that the reality is quite the opposite.

Altogether, given the limitations of our approach, our contribution is twofold. First, we highlight the value of using a quasi-experimental design to isolate the causal influence of culture on decision making. Strangely, many of us are probably comfortable with the notion that somehow cultural differences exist. However, from a strictly empirical perspective, cultures routinely covary with other confounds, and separating the effects of culture from these confounds can often be difficult or impossible. Our approach does so by essentially identifying systematic group-level differences that cannot be genetic, environmental, or institutional. Second, we specifically isolate cultural effects of this sort in decision-making domains related to health and fertility. In this way, although we do not examine genetic fitness directly, we do lean in this direction by focusing on cultural variation in support of policies that should influence fertility, health, and survival. The variation in question is a group-level phenomenon based on cultural evolutionary processes, but it should have consequences for individual reproduction and by extension fitness.

Bibliography

Bell, A. M., Hankison, S. J. & Laskowski, K. L. (2009), 'The repeatability of behaviour: a meta-analysis', *Animal behaviour* **77**(4), 771–783.

Brown, M., Henchoz, C. & Spycher, T. (2018), 'Culture and financial literacy: Evidence from a within-country language border', *Journal of Economic Behavior & Organization* **150**, 62–85.

Buhler, S., Nunes, J. M., Nicoloso, G., Tiercy, J.-M. & Sanchez-Mazas, A. (2012), 'The heterogeneous HLA genetic makeup of the Swiss population', *PloS ONE* **7**(7), e41400.

Cattaneo, M. D., Idrobo, N. & Titiunik, R. (2019), *A practical introduction to regression discontinuity designs: Foundations*, Cambridge University Press.

Chatelan, A., Beer-Borst, S., Randriamiharisoa, A., Pasquier, J., Blanco, J. M., Siegenthaler, S., Paccaud, F., Slimani, N., Nicolas, G., Camenzind-Frey, E. et al. (2017), 'Major differences in diet across three linguistic regions of switzerland: results from the first national nutrition survey menuch', *Nutrients* **9**(11), 1163.

Etter, V., Herzen, J., Grossglauser, M. & Thiran, P. (2014), Mining democracy, *in* 'Proceedings of the second ACM conference on Online social networks', pp. 1–12.

Eugster, B., Lalive, R., Steinhauer, A. & Zweimüller, J. (2011), 'The demand for social insurance: does culture matter?', *The Economic Journal* **121**(556), F413–F448.

Eugster, B., Lalive, R., Steinhauer, A. & Zweimüller, J. (2017), 'Culture, work attitudes, and job search: Evidence from the swiss language border', *Journal of the European Economic Association* **15**(5), 1056–1100.

Evershed, R. P., Davey Smith, G., Roffet-Salque, M., Timpson, A., Diekmann, Y., Lyon, M. S., Cramp, L. J., Casanova, E., Smyth, J., Whelton, H. L. et al. (2022), 'Dairying, diseases and the evolution of lactase persistence in europe', *Nature* **608**(7922), 336–345.

Faessler, L., Lalive, R. & Efferson, C. (2022), 'How culture affects our fitness'. Preregistration.
URL: *https://doi.org/10.17605/OSF.IO/YBCAS*

Frankham, R., Ballou, S. E. J. D., Briscoe, D. A. & Ballou, J. D. (2002), *Introduction to conservation genetics*, Cambridge university press.

Henrich, J. (2004), 'Cultural group selection, coevolutionary processes and large-scale cooperation', *Journal of Economic Behavior & Organization* **53**(1), 3–35.

Hrdy, S. B. (1999), *Mother nature*, Chatto & Windus London.

Laland, K. N. (2008), 'Exploring gene–culture interactions: insights from handedness, sexual selection and niche-construction case studies', *Philosophical Transactions of the Royal Society B: Biological Sciences* **363**(1509), 3577–3589.

Lamba, S. & Mace, R. (2011), 'Demography and ecology drive variation in cooperation across human populations', *Proceedings of the National Academy of Sciences* **108**(35), 14426–14430.

Lee, D. S. & Lemieux, T. (2010), 'Regression discontinuity designs in economics', *Journal of economic literature* **48**(2), 281–355.

Mesoudi, A. & Danielson, P. (2008), 'Ethics, evolution and culture', *Theory in Biosciences* **127**(3), 229–240.

Navarrete, A., Van Schaik, C. P. & Isler, K. (2011), 'Energetics and the evolution of human brain size', *Nature* **480**(7375), 91–93.

OFS (2022*a*), The linguistic landscape in switzerland, Technical Report 2203-2200, Swiss Federal Statistical Office.

OFS (2022*b*), Movers by nationality and distance moved, Technical Report je-f-09.03.06.02.10, Swiss Federal Statistical Office.

Richerson, P., Baldini, R., Bell, A. V., Demps, K., Frost, K., Hillis, V., Mathew, S., Newton, E. K., Naar, N., Newson, L. et al. (2016), 'Cultural group selection plays an essential role in explaining human cooperation: A sketch of the evidence', *Behavioral and Brain Sciences* **39**, e30.

Rochat, C., Eap, C. B., Bochud, M. & Chatelan, A. (2019), 'Caffeine consumption in switzerland: Results from the first national nutrition survey menuch', *Nutrients* **12**(1), 28.

Wormley, A., Kwon, J. Y., Barlev, M. & Varnum, M. E. (2022), 'Ecology explains a substantial amount of human cultural variation around the globe.'.

Wrangham, R. W. & Carmody, R. N. (2010), 'Human adaptation to the control of fire', *Evolutionary Anthropology* .

Chapter 2

Success-biased social learning:

Using successful leaders as examples of how to behave and how not to behave

Lisa Faessler, Charles Efferson

Faculty of Business and Economics, University of Lausanne, Switzerland

Abstract

People tend to imitate successful individuals more than those who are not successful, and leaders leverage this tendency when "leading by success." However, recent theoretical and empirical evidence suggests that social learning strategies, including success-biased social learning, are more complex than previously thought. The added complexity and flexibility could challenge current success-based leadership techniques. Focusing on success-biased social learning, we studied how individuals adapt their behaviors based on the actions of successful leaders. Through an incentivized experiment, we examine participants' decision-making after observing the behavior of a successful leader and two additional signals: the leader's group affiliation and affiliation implications. An accompanying gene-culture coevolutionary agent-based simulation integrates cognitive mechanisms to process the same three types of social information, along with private information. Our findings highlight three critical aspects of success-biased social learning. First, strategies are multi-dimensional. Participants and agents adjust their responses based on multiple pieces of social information. Second, while adjustments are symmetric in the simulation, the experiment shows participants perform better in certain conditions than others, indicating cognitive biases. Finally, the simulation and the experiment demonstrates significant heterogeneity and flexibility in the use of social learning strategies. Our results show that followers adjust to success-dependent social information in complex and heterogeneous ways, including using successful leaders as negative examples. More effective success-based leadership strategies would integrate the complexity and flexibility of followers' cognition.

Keywords: organizational behavior, decision-making, social learning strategies, success bias, group affiliation effects

2.1 Introduction

People tend to imitate successful individuals more than those who are not successful. This tendency is called success-biased social learning and is well-documented in social sciences and evolutionary science (Offerman & Sonnemans 1998, Henrich & Gil-White 2001). Leaders can and do leverage this tendency to follow success in several ways. They showcase their own success to encourage others to emulate their behaviors. They highlight successful team members, making successful behaviors visible and hoping that followers will get inspired by these behaviors. Or they pair successful mentors with individuals in mentorship programs, hoping they will naturally follow successful examples. However, recent theoretical and empirical evidence suggests that social learning strategies, including success-biased social learning (Ehret et al. 2021), are more complex than previously thought. This added complexity would challenge current "leading by success" practices.

Followership theory emphasizes the active role followers play in shaping leadership outcomes (Uhl-Bien et al. 2014). In particular, followership research explores how followers' cognition, traits, and values can moderate the impact of leader behaviors (Matthews et al. 2021, Oc et al. 2023). Central to understanding followership is social learning, also called vicarious learning, where individuals observe and learn from the behaviors of others (Bandura & Walters 1977, Manz & Sims Jr 1981). Humans rely heavily on others to learn and adapt to their environment (Cavalli-Sforza & Feldman 1981, Boyd & Richerson 1985). By imitating what others are doing, they spread behaviors, and because humans do not imitate random fellow humans, some behaviors are more likely to spread, and others disappear (Laland 2004, McElreath et al. 2008).

A social learning strategy is a relatively simple strategy that summarizes how an agent responds to a given type of social information. We can imagine social learning strategies of all levels of complexity. However, many social scientists and cultural evolutionists focus on quite simple strategies, for example, strategies that are one-dimensional in the sense that they represent a response to one and only one observed variable (Efferson et al. 2008). Individuals focus on a single variable and vary their behavior only according to this variable. For example, they respond to the behavior of a successful demonstrator (Offerman & Sonnemans 1998). Influential leaders serve as role models, but many factors can moderate their influence on followers, including the perceived relevance of the leader to the follower's environment (Offerman & Sonnemans 1998).

However, recent empirical evidence suggests that social learning strategies are more complex than we thought. Individuals can consider multiple variables in social learning and express social learning strategies of higher orders (Mesoudi et al. 2016, Efferson et al. 2016). Strategies could involve flexibility and adjustments to multiple pieces of information such as individual characteristics, group affiliations, and the perceived relevance of specific behaviors (Myers 2018, Bellamy et al. 2022). In the context of success-biased social learning, individuals do not simply mimic successful leaders universally. Instead, they could adjust their strategies to group membership, the context, or the perceived relevance of the leader. This complexity implies that followers might even do the opposite of what successful leaders do under certain conditions.

Aside from complexity, theoretical work in gene-culture coevolution tends to assume that social learning strategies are homogeneous as it vastly simplifies, for example, developing and analyzing simulations (Boyd & Richerson 1985). This assumption leads to two erroneous ideas about social learning. First, social learning strategies tend to be the same across agents; many or even all agents rely on the same strategy. Second, social learning strategies are fixed at the individual level. Intuitively, a strategy is homogeneous in the sense that the agent relies on the same strategy across decision-making settings. With the expansion of experimental research on social learning and cultural evolution, we now know that both ideas are probably wrong (Mesoudi et al. 2016, Kendal et al. 2018). Social learning strategies are radically heterogeneous.

Complexity and heterogeneity in social learning have significant implications for leaders who rely on success to guide their teams. A leader might demonstrate the success of a new tool, expecting widespread adoption, or promote in-office work by showcasing their productivity, assuming it will inspire the same behavior in employees. If followers do not perceive these leaders and their actions as relevant to their own circumstances, they might not follow or, worse, might use these examples as counter-examples. For instance, a mentorship program designed to have all followers emulate their mentors may fail if followers view the mentors as outgroup members and adjust their strategies accordingly. Understanding these dynamics is crucial for predicting and enhancing the effectiveness of leadership strategies. This study contributes to followership theory by exploring the cognitive processes of followers in response to successful leaders when additional pieces of information are available. By expanding our understanding of social learning strategies and their complexities, we aim to improve leadership practices, enabling leaders to better predict and influence the behaviors of their followers.

Further, complexity and heterogeneity in social learning strategies at the individual level can significantly impact the dynamics at the aggregate level, from teams to the entire organization

(Granovetter 1978, Young 2009, Efferson et al. 2020). How individuals interpret and react to their peers' actions shapes decision-making, team dynamics, and organizational culture dynamics (Levitt & March 1988, Argote et al. 2000, Argote & Levine 2020). Allowing multiple variables to feed social learning strategies can change the cost of social learning compared to individual learning, affecting cultural evolution dynamics (Efferson et al. 2016).

This was the point of departure for the present study. Focusing on success-biased social learning, we studied social learning complexity and heterogeneity at the individual level through a simulation and a lab experiment. In both the simulation and the experiment, followers were exposed to three pieces of social information. In the first part, we investigate the complexity of social learning strategies. Both the simulation and experiment show that followers expressed complex social learning strategies, and that includes the use of a successful demonstrator as a negative example. Followers adjust how they respond to the successful demonstrator based on the three pieces of social information: the demonstrator's allocation, group affiliation, and the meaning of group affiliations. However, these adjustments are symmetric in the simulation but not in the experiment. Holding the value of social information constant, experimental social learners perform best when observing successful ingroup demonstrators when shared group affiliation predicts similar decision-making environments. The second part focuses on the two types of heterogeneity discussed above: how learners vary their responses to social information from one situation to another and how the same situation can trigger different responses from one learner to another. Results demonstrate tremendous heterogeneity of social learning strategies at the individual level. Strategies vary between individuals in the same situation and across situations for a single individual. Our results show that followers adjust to success-dependent social information in complex and heterogeneous ways that include the use of successful leaders as negative examples. The findings from our simulation and experiment have significant implications for real-world organizations, where leaders can use these insights to foster more effective success-based leading strategies by adjusting to followers and context characteristics.

2.2 Complex, biased and heterogenous social cognition

2.2.1 Social learning strategies complexity

We do not know how complex social learning strategies are and how they vary across individuals and situations. Here, we investigate social learning complexity and flexibility when learners make decisions with information about successful others' choices. We first focus on complexity defined as the number of variables that feed into social learning strategies. In that sense, simple strategies are based on a single variable. In the case of payoff-dependent social learning, for example, learners would respond to success but would not consider the other information pieces. Complex strategies, on the contrary, would integrate multiple variables. Learners would respond to success but could also adjust to other variables, such as the leader's characteristics, context, or relevance.

To illustrate, here are different strategies of various complexity, all derived from the same generic strategy, following success. In the simplest version, an individual relies on a single variable and, for example, always follows the behavior of a successful demonstrator. From an evolutionary perspective, this makes sense; imitating successful behaviors is likely to make you successful as well. This is the first layer of complexity. In a slightly more complex version, we can imagine that two variables feed into the strategy. For example, an individual selects a successful person but from a subsample based on observable markers. That is, restricting the sample to people who share some readily observable characteristics with her, such as dress, language, or ethnicity. Thus, the learner pays attention to two variables, the successful behavior and the demonstrator's group membership, based on observable traits. Until recently, we were more likely to share our environment with people similar to us; thus, imitating people like us tended to be a successful strategy. Note that this could imply not imitating (even doing the opposite) a successful individual if she does not share some observable markers. In a third version, the learner would pay attention to three variables, for example, the successful behavior, the demonstrator's group membership, and the meaning of group membership information. The social learner observes a successful person who is more or less similar to the social learner based on some readily observable trait, and the social learner has additional information about the relevance of the observed similarity. These three strategies could initially seem very similar, yet they would produce very different cultural evolutionary dynamics at the aggregate level (Efferson et al. 2020).

To distinguish between complexity levels, we developed a gene-cultural coevolutionary simulation and conducted an incentivized experiment. Participants were given three pieces of social information: a successful demonstrator's choice, the demonstrator's group affiliation, and additional context about that affiliation. Treatments varied based on these information pieces, and we assessed how they influenced social learning strategies. Specifically, we examined the number of variables affecting these strategies. If participants respond to multiple variables, that would confirm that social learning strategies are multidimensional.

2.2.2 Biased cognition

Allowing for more complex social learning strategies does not necessarily imply that individuals do not evolve a biased cognition. The term "biased" can have two different meanings in this setting, and we want to clarify what those two concepts are and which terms we will use in the rest of the paper to call each of them. As described by Boyd & Richerson (1985), the first meaning refers to cultural evolution bias. Specific social learning strategies generate endogenous cultural evolutionary dynamics. That is what we refer to when we say social learning is not random or biased. People follow specific social learning strategies that cause behaviors and practices to spread and evolve within a culture. For example, people imitating successful individuals, leading to widespread adoption of those behaviors. In the rest of the paper, we only refer to this type of bias by saying success-biased social learning.

The second meaning pertains to cognitive bias in the context of error management theory (McKay & Efferson 2010). Here, a cognitive bias can be trivial or interesting. A trivial cognitive bias means a tendency to hold beliefs that are not uniformly distributed but may be justified by evidence. For example, believing certain outcomes are more likely based on past experiences. An interesting cognitive bias involves systematic deviations from Bayesian updating, where individuals consistently process information in a way that leads to sub-optimal decisions. A biased cognition in that sense would lead the learner to use social information better in some situations than others, making better choices in some situations than others. The intuition is that some learning settings in the ancestral past were more frequent than others, and this persistent asymmetric exposure to learning settings in the past has shaped the evolution of human cognition (Barrett 2014). Evolution has retained a higher level of learning complexity only for frequent environments in our ancestral past. Consequently, if the contemporary settings are similar to ancestral ones, learners will demonstrate a higher level of complexity in their social learning strategy. They will better integrate information and perform better, regardless of explicit contemporary incentives (Cosmides & Tooby 2013). Namely, learners would not adjust symmetrically in two informationally equivalent settings but rather perform better in a situation consistent with the ancestral setting. For example, learners would perform better when learning from an ingroup member than an outgroup member, even though the incentive structure is rigorously the same in the two situations. In the rest of the paper, we call that type of bias "cognitive bias."

In our research, both types of bias are involved. First, we study specifically success-biased social learning, where individuals tend to follow successful demonstrators, aligning with Boyd and Richerson's concept (1985). However, simultaneously, the experiment design allows us to identify cognitive biases of the second sort. We constructed four informationally equivalent treatments. All four convey the same amount of information. The only difference is the framing of information. Participants can observe either an ingroup member or an outgroup member, and they can be similar (dissimilar) with ingroup members. By extension, in principle, participants can earn equivalent amounts of money on average in these four treatments. However, if they form beliefs under a biased cognition, they will perform better in some treatments than others. In particular, they will perform better in the most "natural" scenarios. We anticipate two potential biases. First, participants will perform better when observing a demonstrator from the same group, an ingroup member, compared to an outgroup member. Second, participants will perform better when they share their environment with ingroup members rather than outgroup members.

2.2.3 Social learning strategies heterogeneity

Aside from complexity, social learning strategies might be more heterogeneous than we currently assume. Even though homogeneity has been a useful hypothesis in cultural evolution theory (Boyd & Richerson 1985), recent empirical evidence suggests that social learning is radically heterogeneous (Mesoudi et al. 2016). Social learning varies across individuals and situations for the same individual. All agents do not rely on the same strategies (Muthukrishna et al. 2016), and the same agent could rely on different strategies as she develops and moves from one situation to another (Morgan et al. 2012). Many simulations lead social learners to misapply a strategy sometimes, applying rigid social learning strategy. It might not always be correct, especially since allowing heterogeneity in social learning strategies drastically reduces situations where agents misapply some strategies. For example, an agent could adopt a different strategy from one period to the next, expressing temporal heterogeneity (Perreault et al. 2012, Efferson et al. 2016). She could also vary strategies across environments or demonstrators. In a way, social learning heterogeneity suggests that agents face different incentives and learn to play the game in different ways as a result (Efferson et al. 2008). Our experimental design allows us to describe and analyze heterogeneity in social learning strategies across participants and situations. Multiple participants will face the same scenario, and the same participant will encounter multiple scenarios.

2.3 Gene-culture coevolutionary simulation

2.3.1 Simulation setup

Simulation description

To formally test our intuitions and obtain predictions for the experiment, we built an agent-based gene-culture coevolutionary simulation. The basic process of a gene-culture coevolutionary simulation is as follows. Each generation, individuals come to the world with an inherited genetically encoded strategy. Individuals receive some information, both social and private, and make a decision. Their decision depends both on their inherited strategies and the information received. Further, these decisions result in payoffs. Based on the relative payoff values, individuals potentially reproduce and transmit their genetically encoded strategies to the next generation, subject to mutations. As generations pass, the relative value of social learning strategies and behaviors vary. Culture and the genetically encoded strategies that generate cultural evolution evolve, and the system potentially stabilizes on a set of culturally evolved behaviors and genetically evolved strategies. We then compare the equilibrium strategies with those used by the experiment participants.

Task In each generation, individuals have to allocate all of their endowment between two projects, which we label here 0 and 1. One project is optimal in the sense that the return per unit invested is greater than one. The other project is suboptimal, and the return is less than one. Project 1 is optimal in Environment 1, and Project 0 is optimal in Environment 0. Individuals receive some social and private information before making their decision but do not know with certainty which environment they face. However, they enjoy a high payoff if they allocate most or all of their endowment to the optimal project.

Cultural model Agents belong to a triangle or circle group, each with distinct optimal projects. In each generation, the triangle group's optimal project is always the opposite of the circle group's optimal project. Imagine, for example, a division of labor setting where each group specializes in opposing tasks to enhance overall efficiency and adaptability. In the simulation, if the optimal project for triangle individuals in generation t is 0, then the optimal project for circle individuals in the same generation (t) is Project 1.

Table 2.1: Simulation parameters and functions

Parameters	Function	Data condition
ϕ	**Environment change probability**. Controls the probability that the environmental state changes between generations.	$\phi \in 0.05, 0.5, 0.95$
σ	**Signal reliability**. Indicates the probability that the private signal correctly indicates the winning project.	$\sigma \in 0.5, 0.6, 0.9$
π_h	Payoff per unit invested in the optimal project.	$\pi_h = 10$
π_l	Payoff per unit invested in the suboptimal project.	$\pi_l = 0.1$
μ	**Mutation rate.** Controls the probability of mutations occurring in the genotype.	$\mu = 0.02$
n	Number of agents in the population.	$n = 1,000$
t_max	Maximum number of generations in the simulation.	$t_{max} = 3,000$
r_max	Maximum number of independent simulation runs.	$r_{max} = 50$

Individuals receive information about which project is optimal but do not know it with certainty. Namely, they receive five pieces of information they can use in their strategy, 3 pieces of social information, and 2 pieces of private information. They observe (1) the allocation of the most successful individual in the previous generation, (2) the group affiliation of the individual learners whose choice they observe, and (3) the probability of the optimal project to alternate from one generation to the next, $\phi \in \{0.05, 0.5, 0.95\}$. ϕ can also be interpreted as information on the value of the group affiliation information. Aside from social information, they receive (4) a private signal, indicating which project is optimal, and they observe (5) the probability that this signal is correct $\sigma \in \{0.5, 0.6, 0.9\}$. Table 2.1 summarizes the parameters' functions and values.

Genetic model Agents are characterized by a genotype with seven variables, c_1, c_2, d_1, d_2, γ_1, γ_2, γ_3. The strategy function includes 7 inherited quantities and two information pieces, the most successful allocation x, and the changing rate of the environment ϕ.

$$\begin{aligned} a &= c_1 + \frac{(d_1 - c_1)\phi^{\gamma_2}}{\phi^{\gamma_2} + (1-\phi)^{\gamma_2}} \\ b &= c_2 + \frac{(d_2 - c_2)\phi^{\gamma_3}}{\phi^{\gamma_3} + (1-\phi)^{\gamma_3}} \\ y &= a + \frac{(b-a)x^{\gamma_1}}{x^{\gamma_1} + (1-x)^{\gamma_1}} \end{aligned} \tag{2.1}$$

Individuals can face four different scenarios, observing either an ingroup or an outgroup demonstrator and a private signal equal to $s = 0$ or $s = 1$. The quantities can be different in each scenario; that is, individuals can adjust their strategy to group membership and private signal, too.

Gene and culture interaction Descendants come to the world with a particular strategy stored in their genes. To make their decisions, individuals use their genetically inherited strategies to process information and produce decisions. They can adjust to a maximum of five pieces of information. In each generation, individuals receive some information and adopt a behavior consistent with the information received and their inherited strategy.

As generations pass, the distribution of inherited strategies and behaviors evolves. Different behaviors produce different payoffs. As the strategies that yield higher payoffs have higher chances of reproducing, the distribution of behaviors can change and change the relative value of genetically encoded strategies. Some strategies are selected over others, and the distribution of strategies changes. The distribution of behaviors and strategies evolve together in a gene-culture coevolutionary system. The system potentially stabilizes on a set of behaviors and strategies. The equilibrium strategies represent predictions for settings involving the same information. We test these predictions in a companion experiment.

Procedure

The simulation starts by setting initial parameters: environment change probabilities ($\phi \in \{0.05, 0.5, 0.95\}$), signal reliability ($\sigma \in \{0.5, 0.6, 0.9\}$), payoff per unit invested in the optimal project ($\pi_h = 10$), payoff per unit invested in the suboptimal project$\pi_l = 0.1$), mutation rate ($\mu = 0.02$), population size ($n = 1000$), number of generations ($t_{\max} = 1000$), and number of independent runs ($r_{\max} = 50$). Parameters and initial conditions are summarized in Table 2.1. Agents' genotypes are first initialized with random values between 0 and 1 for 4 of the 7 variables ($c1, c2, d1, d2$ and with 1 for the factorial values $\gamma 1, \gamma 2, \gamma 3$). Agents are assigned to the circle or triangle group, each with a designated optimal project.

In the first generation ($t = 1$), agents observe private signals indicating the environment state ($s = 0$ or $s = 1$) and its reliability ($\sigma \in \{0.5, 0.6, 0.9\}$). They do not receive social information. Agents make decisions based on their genotypes and the private information received. Payoffs are calculated based on each agent's optimal project and allocation decision. Fitness is computed as a combination of these payoffs, endogenous fitness, and an exogenous fitness component. The best-performing agent in each group (triangle and circle) is selected as the demonstrator. We record population-level statistics, such as mean allocation, payoff, and fitness. Agents reproduced based on relative fitness. A higher fitness increases the likelihood of passing on genotypes. Offspring genotypes are subject to mutations.

In subsequent generations ($t > 1$), a round proceeds as follows. First, the environment changes with probability ϕ. When ϕ is low, the environment is stable, and optimal projects will likely be the same as in $t - 1$. When ϕ is high, the environment is unstable, and optimal projects are likely to differ from $t - 1$. Private signals are generated for each agent, reflecting the updated environment state. Additionally, agents receive social information. They observe the allocation of the most successful individual in the previous generation, its group affiliation (triangle or circle), and the probability of the optimal project to alternate from one generation to the next, $\phi \in \{0.05, 0.5, 0.95\}$. Agents make allocation decisions based on their genotypes and the observed behaviors. Payoffs are calculated based on the optimal project and the agent's decision. Fitness values are computed, and agents reproduce based on relative fitness. Mutations can occur. The best-performing agents in each group are selected as the new demonstrators.

This process is repeated for 3,000 generations (t_{max}). Due to variability and drift, the system does not entirely stabilize. However, we observe that genotype variables under selection stabilize within a specific range from about generation 1000 onward. To allow for some margin, we decide to stop at 3,000 generations. Each simulation includes 1000 individuals and is repeated 50 times. Some parameters vary within a simulation. The observed allocation depends on the most successful allocation in the previous generation and varies from generation to generation. The demonstrator group alternates as agents observed alternatively a circle or a triangle demonstrator. ϕ and σ vary across simulations. We ran separate simulation for the 9 parameters combinations ($\phi \in \{0.05, 0.5, 0.95\}, \sigma \in \{0.5, 0.6, 0.9\}$.

The gene-culture coevolutionary simulation explores the evolution of social learning strategies by modeling how agents observe and learn from successful agents. By observing behaviors and genotypes when the system stabilizes, we can understand if agents adjust to multiple pieces of information, exhibit biases, and show heterogeneity in their strategies. These benchmarks provide a valuable basis for comparing experimental results and understanding the potential evolution of social learning cognition complexity.

2.3.2 Simulation results

Figure 2.1: The average strategies when the system stabilizes

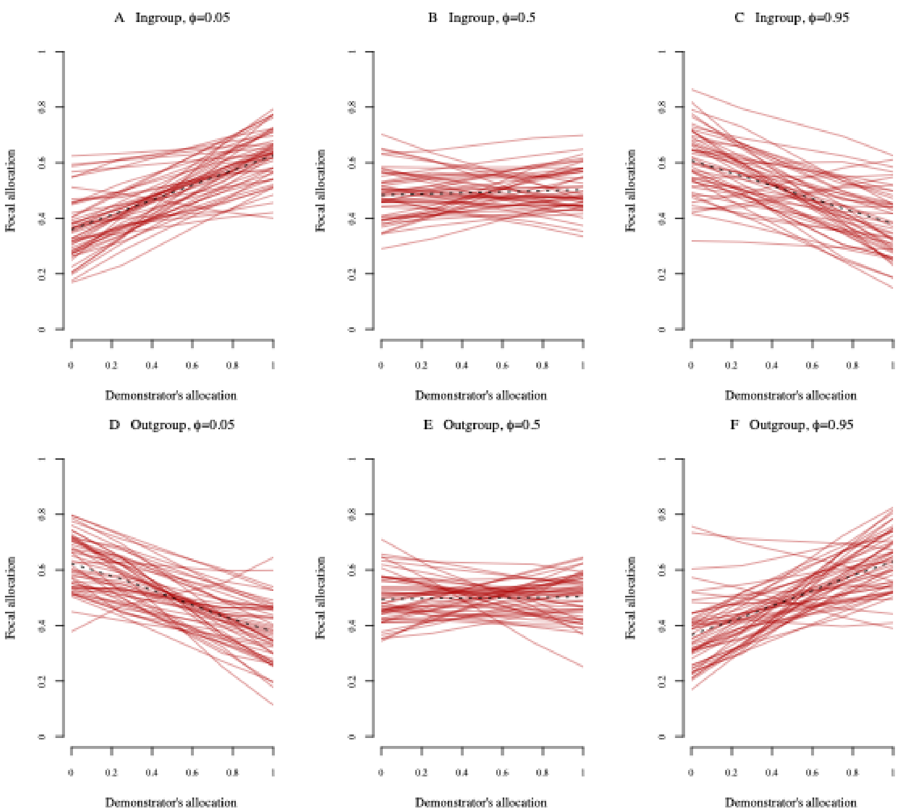

Notes: The average strategies when the system stabilizes. For each treatment combination and each simulation, we computed the average strategy. Each red line represents the average strategy in the last generation of a simulation. The black dotted line is the average of all simulations.

Multidimensional social learning strategies evolve. Below, we focus on scenarios that match our experimental design, when agents can only learn socially ($\sigma = 0.5$). Private information has no value. The curious reader can find the results for scenarios involving both social learning and individual learning in appendix 2.7.1. The results suggest that evolved strategies include both social information and private information. The valence of one type of information compared to the other depends on the context, specifically how reliable the private signal is.

Regarding scenarios where only social learning is available, two key findings follow. First, when social information is valuable, the system stabilizes on complex social learning strategies, that is, strategies that include all pieces of information available. Agents adjust how they respond to a successful demonstrator based on the demonstrator's group affiliation and the meaning of group affiliations. The set of strategies includes using a successful ingroup member as a negative example. Figure 2.1 shows the equilibrium strategies in each scenario. Each red line represents the average strategy at the end of one simulation.

When social information is valuable ($\phi = \{0.05; 0.95\}$), social learners' allocation varies according to the observed demonstrator's allocation, affiliation, and ϕ the meaning of the affiliation information. In figure 2.1 panel A and F, agents likely imitate the demonstrator; the bigger the proportion of the endowment the demonstrator allocates to Project 1, the bigger the proportion of their own endowment they allocate to Project 1 as well. The system stabilizes on the opposite strategy in figure 2.1 panels C and D. The bigger the proportion of the endowment the outgroup demonstrator allocates to Project 1, the bigger the proportion of their own endowment they allocate to Project 0. These findings suggest that agents adjust to affiliation information (ingroup versus outgroup) and ϕ the meaning of the group affiliation, too. Adjustments include the use of a successful demonstrator as a negative example. In figure 2.1 panels C and D, agents do the opposite of a successful demonstrator; they allocate most of their endowment to the project, whereas demonstrators do not.

Adjustments are symmetric. Second, these adjustments are symmetric. Agents perform equally well in all scenarios, keeping social information's value constant. They can learn similarly from ingroup and outgroup members and adjust perfectly to ϕ, the likelihood of the environment changing from generation to generation. Figure 2.2 describes the average performance over agents and simulations across treatments. We considered the last generations of each simulation when the system had stabilized. The graph shows no difference between the average payoffs in the four informationally equivalent scenarios: ingroup 0.95, outgroup 0.95, ingroup 0.05, and outgroup 0.05. We observe similar performance across ingroup and outgroup treatments and ϕ levels. The

Figure 2.2: Average performance across scenarios

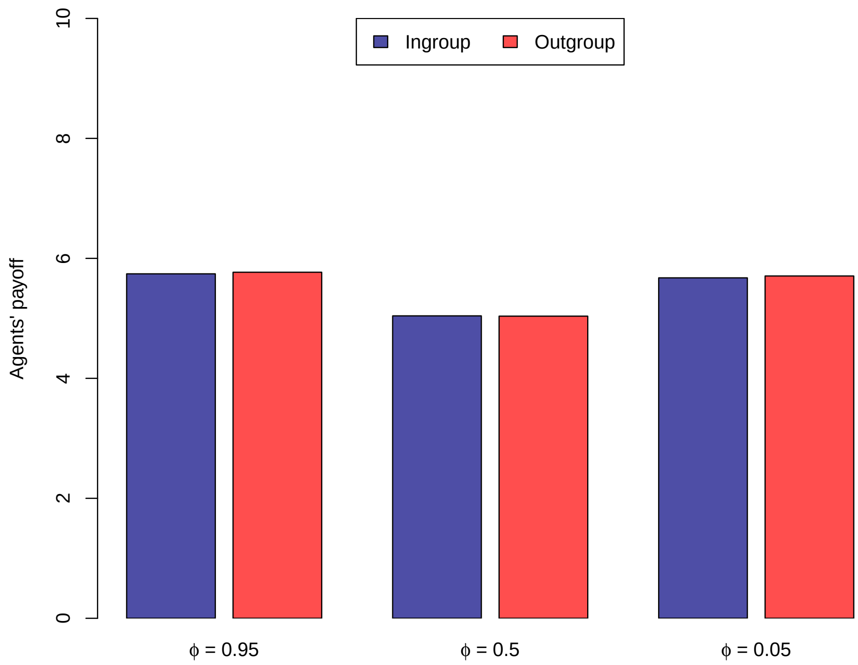

Average performance over agents and simulations across treatments. We average agents' performance across treatments in the 100 last generations of each simulation. Agent performance is the number of points obtained by a single agent. Each agent receives 10 points per unit invested in their optimal project and 0.1 points per unit invested in their suboptimal project. The blue bars represent treatments where agents observe an ingroup demonstrator, while the red bars represent outgroup treatment. ϕ indicates the likelihood of the optimal project changing from generation to generation. Results show no significant performance difference in information equivalent treatments

results suggest that the agents show no cognitive biases. Agents were equally able to learn from ingroup and outgroup members in stable or changing environments.

Heterogenous genotypes evolve. Having discussed the phenotypic behaviors and performance outcomes, we now turn to the underlying genotypes. The advantage of a gene-culture coevolutionary simulation is that we can analyze both behaviors and the inherited strategies driving them. By examining the evolution of these genotypes, we explore the mechanisms behind the observed behaviors and variability within the population. To this end, we analyze the genotypic variance over simulation rounds. Specifically, we look at the seven inherited quantities that define the inherited social learning strategy (c_1, c_2, d_1, d_2, γ_1, γ_2, γ_3). We aim to observe how these genetic parameters evolve and potentially stabilize over time.

Figure 2.3 displays the variance in all quantities included in the genotypes of the agents

Figure 2.3: Variance of Genotypic Quantities in Agents

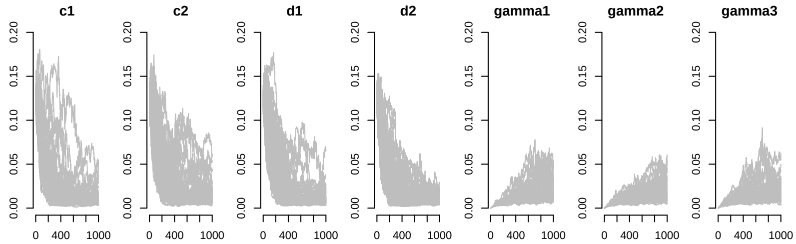

Notes: This graph displays the variance in all quantities included in the genotypes of the agents over simulation rounds. Each panel plots variance for different parameters (e.g., c1, c2, d1, d2, gamma1, gamma2, gamma3) against simulation rounds (0 to 1000). Some variables exhibit increased variance due to a lack of selection pressure and are subject to drift, while others are more stable, indicating different levels of selection pressure.

over simulation rounds. Diverse evolutionary patterns emerge for different genetic parameters. Some variables, c_1, c_2, c_3, d_1, and d_2, exhibit a decreasing then relatively stable variance over the simulation rounds. This pattern suggests that some selection pressures constrain these values within a specific range. These parameters are under stronger selection pressures than *gamma* variables. Despite the stronger pressures, c and d parameters still exhibit some variability. This variability indicates that multiple strategies can produce outcomes that are payoff-equivalent. Thus, we expect a degree of heterogeneity to persist within the population, even though these variables face some selection pressures.

Conversely, other variables, γ_1, γ_2, and γ_3, show increasing variance over time. These variables become increasingly diverse, suggesting random fluctuations due to drift rather than adaptive changes. Genetic drift is a process where certain traits change randomly over generations. Traits can take on multiple values without being strongly selected for or against, increasing traits' variance over time. This randomness contributes further to the heterogeneity of strategies within the population.

The simulation results suggest, first, that evolution selects the most complex learning strategies. The strategies present when the system stabilizes integrate the three pieces of social information. We test these predictions in a lab experiment. The lab experiment involves the same three pieces of social information. The simulation can accommodate up to five pieces of information; three pieces are social information, and two pieces are private information. In the experiment, we omitted the two private pieces of information to concentrate on social information and isolate social learning strategies. This method counters the reflection problem by separating individual learning from social learning. In the experiment, we test whether participants use social learning strategies of the same complexity level and whether they adapt their strategies to the three pieces of social information.

Second, in the simulation, adjustments are symmetric. That makes sense, as our simulation relies on two critical assumptions. First, agents inherit strategies that can process all the relevant variables. Second, there is no cost for associated cognitive complexity. However, human cognition could easily be different. For example, humans could have evolved better learning abilities in some scenarios compared to others due to more frequent exposure. Learning in some scenarios could be more costly than in other scenarios. Our current model does not incorporate cognitive costs or varying exposures to different scenarios, which could potentially lead to asymmetric adjustments. While this is a limitation of our simulation, we address this aspect in our experiment. We tested the presence of cognitive biases in our experiment by exposing the participants to four scenarios where the information value is the same but where the framing differs. Because the scenarios carry the same amount of information, any significant difference in performance will tell us that adjustments are not symmetric and that humans have evolved cognitive biases to learn more easily in some scenarios than others.

Third, our findings indicate that some genetic parameters are loosely regulated by selection pressures, leading to convergence in those traits, while others are subject to drift, fostering greater variability. We also anticipate some heterogeneity in social learning strategies to manifest in the experiment.

2.4 Experimental methods

2.4.1 Experiment overview

We conducted the following experiment with 133 students at the University of Lausanne. In each round, each participant had to allocate an endowment of 100 tokens to two projects, "Project A" and "Project 1." One project is optimal as the return per unit is greater than one. The other project is suboptimal, and the return per token invested is less than one. Participants maximize their profit by allocating their tokens to the optimal project. The optimal project is randomly selected from a uniform distribution at the beginning of each round. Participants do not know which project is optimal with certainty but can have some information. Information varies across treatments.

2.4.2 Task

Sessions last for 90 rounds, and all participants make one choice per round. In each round, participants have to make an investment decision about how to allocate 100 tokens to two projects, "Project A" and "Project 1." They have to allocate all of their endowment. To allocate the tokens, participants use a slider. They can decide to allocate any amount between 0 and 100 to project 1, in 1 increments. Figure 2.9 in the Appendix illustrates the decision screen. Between these two projects, one is optimal and yields a return per token greater than one, 2. For each token invested in that project, the participant receives 2 points. The other project is suboptimal, and the return per token invested is less than one, 0.5. For each token invested, the participant receives 0.5 points. The optimal project is randomly selected from a uniform distribution at the beginning of each round. Participants might have some information about which project is optimal, but they do not know it with certainty. The information provided is part of our treatments.

The total number of points for one round and one participant is the sum of the points yielded by the investment in the optimal project and the sum of points yielded by the investment in the suboptimal project. Therefore, the result for a round depends on the allocation decision and the fate of the two projects, whether the optimal project is Project 1 or Project A. Participants can earn each round between 50 ($100 * 0.5$) and 200 points ($100 * 2$). The points are translated into CHF at the end of the experiment. Because the decision-making domain involves risk, and because we want risk to matter, we randomly select 2 rounds to pay. Thus, participants cannot spread risk over all 90 rounds and essentially eliminate risk from consideration.

2.4.3 Groups and roles

Before the session starts, participants are assigned a role, individual learner or social learner, and a group, triangle or circle. Thus, participants can be of four types: triangle individual learners, circle individual learners, triangle social learners, and circle social learners. Participants are randomly assigned to the four types according to the following rule.

- 3 participants are circle individual learners.
- 3 participants are triangle individual learners.
- half of the remaining participants are circle social learners.

- half of the remaining participants are triangle social learners.

The roles and groups remain constant throughout the session.

Both individual and social learners are divided into two groups, triangles and circles. The optimal project for circle individual learners is randomly selected at the beginning of each round according to a uniform distribution. The optimal project for circle individual learners is the opposite of the optimal project for triangle individual learners. To illustrate, if the optimal project for triangle individual learners in round t is Project A, then the optimal project for circle individual learners in the same round is Project 1.

Roles, individual learner or social learner, differ according to the timing of the decision and the information they receive. Each round is divided into two parts. The individual learners play during the first part. The social learners play during the second part. Individual learners, playing first, do not receive any information about the optimal project. They make choices and receive immediate private feedback about their choices' payoff consequences. The decision screen for individual learners is in Appendix, Figure 2.9. In our experiment, individual learners only serve as successful examples to social learners. Although we are very grateful for individual learners to join for the experiment, we are only interested in the decisions of the social learners.

Social learners and social information

Social learners play in the second part of the round. Before making their choices, they receive three pieces of social information. The decision screen for social learners, including the social information display, can be found in the Appendix, Figure 2.10. First, social learners observe the allocation of a successful demonstrator. At the beginning of the experiment, the computer randomly draw one of the group, triangle or circle. Imagine, for example, that the computer chooses triangle. In each of the first 45 rounds, we will take the most successful individual learner (i.e., highest payoff) from the triangle group. We call this successful individual learner the "demonstrator". All social learners, whether triangle or circle social learners, observe the demonstrator's allocation choice. By extension, some social learners observe the allocation choice of the most successful ingroup individual learner. In contrast, other social learners observe the allocation choice of the most successful outgroup individual learner. We refer to the allocation choice itself as "first-order" social information. After 45 rounds, midway through the session, we switch groups and select for the remaining 45 rounds the most successful individual learner from the other group. From a social learner perspective, the demonstrator is an ingroup (outgroup) member for the first 45 rounds and an outgroup (ingroup) member for the last 45 rounds.

Second, social learners know the group affiliation (i.e., triangle or circle) of the demonstrator whose choice they observe. As mentioned above, the demonstrator is an ingroup member for some social learners and an outgroup member for others. We refer to the group affiliation of the reported demonstrator as "second-order" social information.

Third, social learners know the ex-ante probability with which their optimal project is the same as the optimal project of ingroup individual learners. This probability can be conceived as the similarly with ingroup members. Let this probability be ϕ. ϕ can take values from the set $\{0.1, 0.5, 0.9\}$. $\phi = 0.9$ is the most "natural" scenario and indicates a high probability of sharing the optimal project with ingroup demonstrators. Because triangle and circle individual learners always have opposite optimal projects, a high probability of sharing the optimal project with ingroup demonstrators indicates a low probability of sharing the optimal project with outgroup demonstrators $(1 - \phi)$. $\phi = 0.1$ indicates a low probability of sharing the optimal project with ingroup members demonstrators and, conversely, a high probability of sharing the optimal project with outgroup demonstrators. Finally, $\phi = 0.5$ provides no information. Social learners are randomly assigned to the ϕ values at the beginning of the experiment. Once assigned a value of ϕ, they retain this value for the session. We refer to ϕ as “third-"der” socia"information.

After receiving social information, social learners make a decision; however, they do not receive private feedback about their choices' consequences. Instead, they only see their total payoffs at the very end of the session. Thus, social learners can only learn socially. This distinction is crucial as it is impossible to disentangle individual learning from social learning if both occur simultaneously (Manski 2000, Efferson et al. 2016).

2.4.4 Two-by-three design based on social information

Table 2.2: Treatments based on social information

	$\phi = 0.9$	$\phi = 0.5$	$\phi = 0.1$
Ingroup demonstrator	Ingroup, 0.9	Ingroup, 0.5	Ingroup, 0.1
Outgroup demonstrator	Outgroup, 0.9	Outgroup, 0.5	Outgroup, 0.1

Treatments are based on social information and vary regarding second-order information, observing an ingroup or outgroup demonstrator, and third-order social information, ϕ, the probability of having the same (opposite) optimal project as ingroup (outgroup) demonstrator. The first treatment dimension is observing an ingroup or outgroup demonstrator. A social learner observes either the most successful ingroup individual learner or the most successful outgroup individual learner. Specifically, each social learner completes one block of 45 rounds in which she observes an ingroup demonstrator each round, and each completes another block of 45 rounds in which she observes an outgroup demonstrator each round. We counterbalance the order of these two within-subject treatments across social learners. The second treatment dimension is ϕ. ϕ can takes values from the set $\{0.1, 0.5, 0.9\}$. At the beginning of the session, each social learner is randomly assigned one of the 3 ϕ values and retains this value for the entire session. Therefore, a social learner with $\phi = 0.9$ will always have 90% chance to share the same optimal project as ingroup demonstrators, independent of the round. This dimension is between subjects but within sessions. In each session, all ϕ values are represented.

Optimal strategies

Table 2.3: Optimal strategies across treatments

	$\phi = 0.9$	$\phi = 0.5$	$\phi = 0.1$
Ingroup demonstrator	*Imitate*	-	*Do the opposite*
Outgroup demonstrator	*Do the opposite*	-	*Imitate*

Social learners' optimal strategies vary depending on the second-order and third-order social information they receive, whether they observe an ingroup or outgroup demonstrator and the probability ϕ of having the same optimal project as the ingroup demonstrator. Table 2.3 summarizes these strategies. Two treatment groups should imitate the demonstrator. Note we use "imitate" her in the sense "allocate the majority of the tokens to the same project as the demonstrator," rather that in a stricter sense "deciding on the exact same allocation as the demonstrator." Two groups maximize expected payoffs by imitating demonstrators and allocating the majority of their tokens to the same project as the observed demonstrator. First, social learners observing an ingroup demonstrator with a high probability ($\phi = 0.9$) of sharing the same optimal project as their ingroup demonstrators. They should imitate the ingroup demonstrator they are observing. Second, social learners observing an outgroup demonstrator with a low probability ($\phi = 0.1$) of sharing the same optimal project as their ingroup demonstrators. They should imitate the outgroup demonstrator they are observing.

On the contrary, two other groups maximize expected payoffs by "doing the opposite" of the demonstrator. Again, by "doing the opposite", we mean "allocating the majority of their tokens to the opposite project as the demonstrator." First, social learners observing an outgroup demonstrator with a high probability ($\phi = 0.9$) of sharing the same optimal project as their ingroup demonstrators. They should do the opposite of the outgroup demonstrator they are observing. Second, social learners observing an ingroup demonstrator with a low probability ($\phi = 0.1$) of sharing the same optimal project as their ingroup demonstrators. They should do the opposite of the ingroup demonstrator they are observing.

Identifying social learning complexity Notice that our treatments do not vary in terms of complexity, defined as the number of variables that feed into social learning strategies. In all treatments, social learners have access to three pieces of information: information about the observed allocation, the demonstrator's group membership, and the relevance of this information. The question is whether they actually process that information. By having treatments systematically varying in terms of optimal strategies, we can identify how many and which pieces of information are integrated into their social learning strategies.

To illustrate, imagine that participants use simple, first-order, social learning strategies in the sense that they process a single variable. For example, social learners only respond to the observed allocation. Then, we would observe different allocations across observed allocations, but similar allocations across ingroup and outgroup treatments, and across ϕ levels. In the 6 treatments, the allocation pattern would look the same. Probably a small amount when the observed allocation is low and a bigger amount when the observed allocation is higher. Now imagine social laerners use second-order strategies, in the sense that they process two pieces of information. For example, they pay attention to the observed allocation and the group membership. We would observe different allocation patterns across allocation levels and ingroup/outgroup treatments but similar allocations across ϕ levels. We would observe one allocation pattern across the ingroup treatments (Ingroup 0.9, Ingroup 0.5, Ingroup 0.1). And another allocation pattern across the outgroup treatments (Outgroup 0.9, Outgroup 0.5, Outgroup 0.1). Observing different allocation patterns across ingroup and outgroup treatments would indicate that social learners are processing at least two pieces of social information. Finally, imagine that social learners use social learning strategies of the third-order, that is, the maximum level of complexity allowed by our setting. We would observe different allocation patterns across the 6 treatments. Social learners would express different strategies depending on the allocation, the group membership and the relevance of the information. Observing different allocation patterns across the different ϕ levels in addition to adjustment to group membership and observed allocation would show that social learners respond to the three variables.

By systematically varing optimal strategies, we can identify how many and which pieces of information are integrated into social learning strategies. The allocation patterns across treatments tell us whether learners are employing first-order, second-order, or third-order strategies and allow us to identify the complexity of social learning processes.

Identifying cognitive biases These four treatments, (i) observing an ingroup individual learner under $\phi = 0.9$, (ii) observing an ingroup individual learner under $\phi = 0.1$, (iii) observing an outgroup individual learner under $\phi = 0.9$, and (iv) observing an outgroup individual learner under $\phi = 0.1$ are informationally equivalent. They are similar in terms of the value of the information available to subjects. By extension, social learners can earn equivalent amounts of money on average in these four treatments. However, these treatments are different in terms of the framing of the information. Comparing informationally equivalent treatments allows us to identify cognitive biases,if any. If social learners systematically perform better or worse in one treatment over another, despite having access to the same value of information, this would suggest the presence of cognitive biases.

In all four cases, the amount participants choose to allocate to the expected optimal project depends on their individual risk preferences. Risk-averse participants would potentially spread their tokens more evenly. Risk-seeking participants would make more pronounced allocations. A risk-neutral participant would maximize their expected payoff by allocating all of her endowment to the expected optimal project. We expect this heterogeneity to create variability in the decisions.

Finally, treatments with $\phi = 0.5$ are different as information has no value. All social learning strategies are equivalent in expected payoffs. Social learners cannot use social information to improve their payoffs. That said, if social learners have any intrinsic preference for one social learning strategy over another, for example, follow success, this preference should be most likely to manifest itself in treatments with $\phi = 0.5$ precisely because the treatment removes material incentives from consideration.

The experiment was pre-registered (Faessler & Efferson 2019). We collected data by running four 2-hour sessions on November 28th and 29th, 2019, in the behavioral laboratory at HEC Lausanne on a computer network using oTree (Chen et al. 2016). We used the online recruiting software ORSEE (Greiner 2004). A total of 133 participants were recruited (mean age = 21.20, SD = 2.78, males = 50.38%). Participants were mainly students from the University of Lausanne and the Swiss Federal Institute of Technology in Lausanne (EPFL). The most represented countries of birth were Switzerland (60 participants) and France (32 participants). A total of 27 countries were represented. The instructions were given in French. Before the experiment, we collected explicit consent from the participants. Participants were free to leave the experiment at any point, although none did it. Participants had to pass a comprehension check to start the experiment. We collected sociodemographic data at the beginning of the experiment and asked for impressions at the end of the experiment. Each participant received CHF 10.- for her participation plus a bonus varying between CHF 12.- and CHF 48.- depending on her success for the two randomly selected rounds to pay out.

2.5 Experimental Results

Our objective is to explore the complexity and heterogeneity of participants' social learning strategies. First, we examine the complexity of social learning strategies (see section 2.5.1). Using descriptive statistics and regression models, we analyze how participants allocate tokens and whether they adjust their behaviors based on the demonstrator's actions, group membership, and the relevance of this information. The results suggest that participants use third-order strategies, which is the maximum complexity allowed by our setting.

Second, we test for cognitive biases by analyzing performance across informationally equivalent treatments (see section 2.5.2). By identifying systematic differences, we learn that participants are more comfortable learning in some scenarios than others.

Finally, we explore the heterogeneity in social learning strategies (see section 2.5.3). We assess the variability of these strategies among individuals in the same situation and within the same individual across different contexts. This analysis reveals that participants do not rely on a single strategy but adapt their approach based on the context.

2.5.1 Social learning strategies are multi-dimensional.

Social learners expressed complex social learning strategies. Participants behaved differently in the four informationally equivalent treatments and did not mindlessly follow first-order information. On the contrary, they adapted their strategies according to second and third-order information.

Figure 2.4: Social learners' average allocations in the six treatment conditions

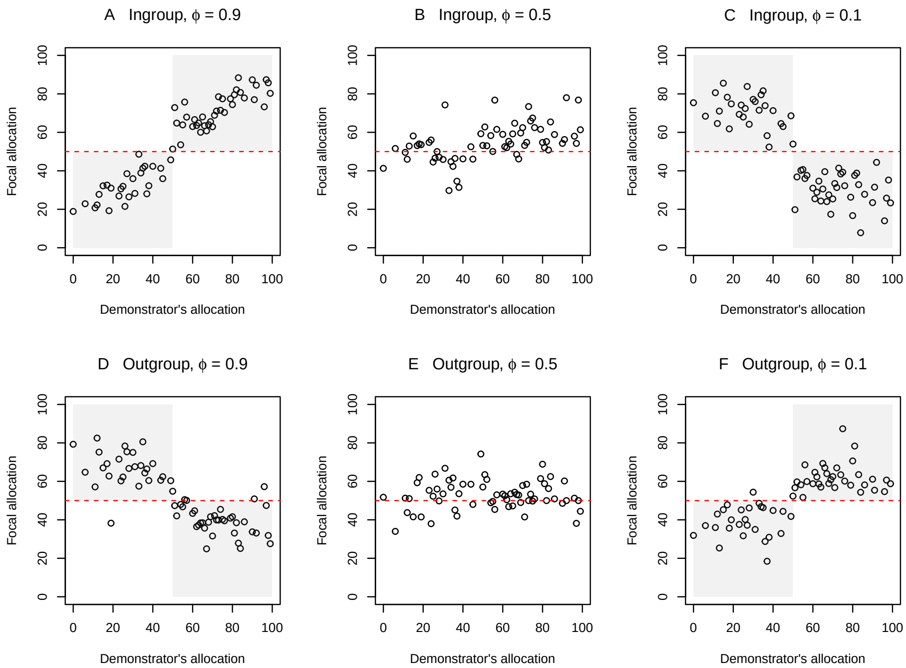

The average allocation to project 1 given the demonstrator's allocation to project 1. The panels show the average allocation over social learners for levels of demonstrator's allocation in the six treatments' combinations. The red lines represent an allocation of 50. The areas shaded in grey depict where an average would fall if they were consistent with an optimal strategy. Panels B and E are not shaded as social learners have no valuable information; there is no optimal strategy for these treatments' combinations.

Figure 2.4 visually describes the participants' average allocation across treatment, that is, the average number of tokens allocated to project 1. The average allocation varies in each treatment condition. For now, we leave aside the special cases when social learners have no valuable information, when the demonstrator allocates 50 tokens to project 1, or when $\phi = 0.5$, that is when social learners do not know if their information is reliable. We focus on scenarios where social learners have valuable social information. In figure 2.4 panel A, ingroup 0.9, participants likely imitate the demonstrator when observing a successful ingroup member likely to share the same optimal project. The more tokens the demonstrator allocates to Project 1, the more they allocate to Project 1. In figure 2.4 panel D, outgroup 0.9, when participants observe an outgroup member in an environment where they are likely to share the optimal project with ingroup members, participants adopt the opposite strategy: the more tokens the successful outgroup demonstrator allocates to project 1, the more they allocate tokens to project 0. They use a demonstrator as a negative example. Similarly, in figure 2.4 panel C, ingroup 0.1, when participants observe an ingroup member unlikely to share the same optimal project, participants do the opposite as the demonstrator. They use a successful ingroup demonstrator as a negative example. Finally, in figure 2.4 panel F, outgroup 0.1, when participants observe an outgroup member likely to share the same optimal project, participants are likely to imitate the demonstrator. The more tokens the demonstrator allocates to Project 1, the more they allocate to Project 1. They adopt a similar strategy to in panel figure 2.4 A. Figure 2.4 provides some first descriptive evidence that participants behave differently in the four informationally equivalent treatments and adapt their strategies according to the demonstrator's allocation, group membership information, and the meaning of the group membership information.

Table 2.4 presents the regression results with y the social learners' allocation to project 1 as the response variable. Results show that social learners use strategies that include the three pieces of information. Social learners adjust their allocation to the demonstrator's allocation (estimate = 0.906, $p < 0.001$), to group membership (estimate = 83.709, $p < 0.001$), and the relevance of the group membership information ϕ ($\phi = 0.5$ estimate = 35.725, $p < 0.001$; $\phi = 0.1$ estimate = 75.031, $p < 0.001$). Further, they adjust to any combination of the information pieces ($p < 0.001$ for all interaction terms). The regression results confirm the visual description provided in figure 2.4; social learners adjust their strategies according to the three pieces of information available.

We computed linear combinations for a subset of treatment combinations to investigate how social learners switch strategies from one scenario to another. Tables 2.5 and 2.6 present the linear combinations estimates when social learners observe the demonstrator allocating either 0 or 100 tokens to project 1. Table 2.5 shows the impact on social learners' allocation of a switch from observing an ingroup to an outgroup demonstrator. The table 2.5 first part presents the estimates of when social learners are likely to share the same optimal project as the demonstrator, $\phi = 0.9$. When demonstrators allocate 0 tokens to project 1, a switch from observing an ingroup to an outgroup member leads to a change of 84 tokens (estimate = 83.73, $p < 0.001$). Similarly, when demonstrators allocate 100 tokens to Project 1, a switch to an outgroup demonstrator reduces the allocation to Project 1 by 85 tokens (estimate = -84.98, $p < 0.001$). Social learners adjust their strategy to the demonstrator's allocation and group membership. The table 2.5

Table 2.4: The regressions modeling how many tokens social learners invested in project 1

Parameters	(1)	(2) Baseline
Demonstrator's allocation x	0.906***	0.906***
	(0.052)	(0.053)
Outgroup	83.797***	83.709***
	(5.530)	(5.552)
$\phi = 0.5$	35.748***	35.725***
	(4.360)	(4.344)
$\phi = 0.1$	74.886***	75.031***
	(5.479)	(5.509)
$x \cdot$ Outgroup	−1.688***	−1.687***
	(0.115)	(0.116)
$x \cdot \phi$=0.5	−0.702***	−0.700***
	(0.079)	(0.079)
$x \cdot \phi$=0.1	−1.539***	−1.539***
	(0.106)	(0.106)
Outgroup $\cdot \phi$=0.5	−73.359***	−73.004***
	(7.032)	(7.006)
Outgroup $\cdot \phi$=0.1	−134.759***	−134.557***
	(10.593)	(10.602)
$x \cdot$ Outgroup $\cdot \phi$=0.5	1.510***	1.503***
	(0.143)	(0.142)
$x \cdot$ Outgroup $\cdot \phi$=0.1	2.750***	2.746***
	(0.215)	(0.215)
Gender		0.921
		(1.103)
Faculty		−0.270
		(0.263)
Country of birth		−0.072
		(0.099)
Observations	9,810	9,810
Adjusted R^2	0.351	0.352
Note:	*p<0.05; **p<0.01; ***p<0.001	

The regression models how many tokens social learners invested in project 1. Independent variables include (a) x the demonstrator's allocation to project 1, (b) whether the social learner observes an outgroup member, (c) ϕ the meaning of the group affiliation information, and (d) the interactions between the demonstrator's allocation and each of these dummies. The omitted category of the regression is a demonstrator's allocation of 0, an ingroup demonstrator, and $\phi = 0.9$, that is, having a 90% chance for the social learner to share optimal projects with ingroup demonstrators. Robust standard errors in parentheses are clustered at the social learner level to reflect the multiple observations per social learner. Both models include fixed effects for the session. Model (2) also includes controls for the social learner's gender, faculty, and country of birth.

second part shows similar results for $\phi = 0.1$ when social learners are likely to have a different optimal project than ingroup demonstrators. A change from observing an ingroup to observing an outgroup modifies the social learners' allocation significantly by about 53 tokens ($x = 0$ estimate = -50.89, $p < 0.001$; ($x = 100$ estimate = 55.13). These estimates are smaller than the ones for $\phi = 0.9$ when social learners are likely to share their optimal project with the demonstrator.

Table 2.5: The linear combinations switching from ingroup to outgroup

	Switching from ingroup to outgroup
$\phi = 0.9$	
Demonstrator's allocation $x = 0$	83.737***
	(5.547)
Demonstrator's allocation $x = 100$	-84.983***
	(6.197)
$\phi = 0.1$	
Demonstrator's allocation $x = 0$	-50.891***
	(9.028)
Demonstrator's allocation $x = 100$	55.131***
	(9.378)

Note: *p<0.05; **p<0.01; ***p<0.001

Table 2.6 shows the impact on social learners' allocation of a switch from an ingroup demonstrator likely to have the same optimal project, $\phi = 0.9$, to an ingroup demonstrator likely to have a different optimal project, $\phi = 0.1$. Results are similar to a switch between observing an ingroup and an outgroup member. A change in ϕ leads to an average difference of 80 tokens allocated to project 1 when observing an ingroup member ($x = 0$ estimate = -75.19, $p < 0.001$; ($x = 100$ estimate = -78.73), and 60 tokens when observing an outgroup member ($x = 0$ estimate = -59.43, $p < 0.001$; $x = 100$ estimate = 61.38). Similarly, the average token differences are different across treatment combinations: the difference is bigger when observing an ingroup member than when observing an outgroup member. We investigate this difference further when looking at potential biases and social learners' performance. The linear combination results confirm that social learners respond to success and group membership and the relevance of the group membership information. Social learners' strategies can and do account for higher-order forms of information. The three available variables feed into social learning strategies, creating strategy functions with three dimensions: the maximum complexity allowed by the setting.

We now investigate special cases when social learners have no valuable information. When

Table 2.6: The linear combinations switching from $\phi = 0.9$ to $\phi = 0.1$

Switching from $\phi = 0.9$ to $\phi = 0.1$	
Ingroup	
Demonstrator's allocation $x = 0$	75.191***
	(5.530)
Demonstrator's allocation $x = 100$	-78.734***
	(5.317)
Outgroup	
Demonstrator's allocation $x = 0$	-59.437***
	(6.275)
Demonstrator's allocation $x = 100$	61.380***
	(7.330)

Note: *p<0.05; **p<0.01; ***p<0.001

social information is meaningless, social learners do not express an intrinsic preference for a particular strategy. They do not adjust either to success or to group affiliation. Linear combinations results in table 2.7 show that when social learners observe a demonstrator's allocation of 50 tokens to project 1, that is, when they have no valuable information about which project is the optimal one, participants did not vary their strategy across group membership or ϕ levels. Switching from observing an ingroup member to an outgroup member leads to an average difference of 1 token allocated to project 1, and this difference is non-significant ($\phi = 0.9$ estimate = -0.62, $p = 0.55$; $\phi = 0.1$ estimate = 2.12, $p = 0.151$). Social learners adopt similar strategies whether they observe an ingroup or an outgroup member. For example, they do not prefer following an ingroup member. Similarly, switching from $\phi = 0.9$ to $\phi = 0.1$ leads to no significant difference in the number of tokens allocated to project 1 (ingroup estimate = -1.77, $p = 0.158$; outgroup estimate = 0.97, $p = 0.599$). These results show that social learners did not jump on a generic strategy, for example, follow success or an ingroup member, when they had no information about its validity. These observations challenge the common assumption in social learning literature that individuals inherently follow successful individuals or ingroup members. Without meaningful information, social learners did not demonstrate a bias towards such generic strategies. This deviation from expected success or ingroup bias not only underscores the absence of inherent preferences but also helps us to rule out demand effects. These null results likely indicate that participants' choices were informed by the available information rather than assumptions about the experiment's expectations.

Table 2.7: The linear combinations when social learners have no valuable information

Demonstrator's allocation $x = 50$	
Switching from ingroup to outgroup	
$\phi = 0.9$	-0.623
	(1.053)
$\phi = 0.1$	2.120
	(1.466)
Switching from $\phi = 0.9$ to $\phi = 0.1$	
Ingroup	-1.771
	(1.245)
Outgroup	0.971
	(1.840)

Note: *p<0.05; **p<0.01; ***p<0.001

2.5.2 Asymmetric adjustments: social learning is biased.

While social learners adjust their strategies to second and third-order information, they do not adjust perfectly. They performed better in some treatment combinations than others. Namely, they failed to adjust perfectly to third-order information and performed better when $\phi = 0.9$ compared to $\phi = 0.1$. We analyzed performance across the four informationally equivalent treatments to test for asymmetric adjustments across treatments. Social learners can, in principle, earn equal amounts of money on average in these four treatments, observing an ingroup or outgroup member under $\phi = \{0.1, 0.9\}$.

Figure 2.5 visually describes the average performance in each treatment condition. The social learner's performance is the number of points a social learner obtains. Each social learner receives 2 points per token invested in their optimal project and 0.5 points per token invested in their suboptimal project. Thus, the performance of each social learner can range from 50 points to 200 points. We focus on treatments where social learners have valuable social information, when $\phi = \{0.1, 0.9\}$. Table 2.8 provides the linear combination coefficients. First, performance did not vary across group membership treatments; social learners obtained similar payoffs in the ingroup and outgroup treatments ($\phi = 0.9$ estimate = -5.54, $p = 0.52$; $\phi = 0.1$ estimate = -6.17, $p = 0.599$). Participants perfectly adjusted their strategy to second-order information; they did not perform better when presented with an ingroup demonstrator than an outgroup demonstrator. These results suggest that participants were just as capable of learning from an ingroup member as an outgroup member.

Figure 2.5: Average performance across treatments

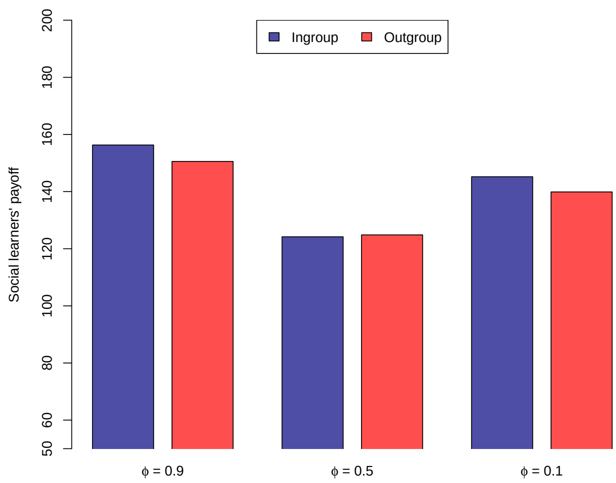

Average performance over social learners across treatments. The social learner's performance is the number of points a social learner obtains. Each social learner receives 2 points per token invested in their optimal project and 0.5 points per token invested in their suboptimal project; social learners' performance can range from 50 to 200 points. The blue bars represent treatments where the social learner observes an ingroup demonstrator, while the red bars represent outgroup treatment. ϕ indicates the likelihood of being similar to an ingroup demonstrator. Social learners in phi $= 0.9$ treatment have a 90% chance of sharing the optimal project with ingroup demonstrators. Results show significant difference between $\phi =$ 0.9 treatments and $\phi = 0.1$ treatments. Performance in ingroup and outgroup treatments are similar.

Second, participants obtained a significantly better average payoff under $\phi = 0.9$ than under $\phi = 0.1$, keeping second-order information constant. Switching from $\phi = 0.9$ to $\phi = 0.1$ decreases performance by 35 points on average (ingroup estimate = -34.94, $p < 0.001$; outgroup estimate = -35.56, $p < 0.05$). Participants did not adjust their strategy perfectly to third-order information; they performed better when ingroup members were likely to share the same optimal project (i.e., $\phi = 0.9$) rather than when ingroup members were unlikely to share the same optimal project (i.e., $\phi = 0.1$). These findings imply that participants did not similarly process information across these two learning settings. They were more comfortable learning in a setting where ingroup members share their environment compared to a setting where outgroup members share their environment.

Table 2.8: Social learners' performance across treatments

Switching from ingroup to outgroup	
$\phi = 0.9$	-5.542
	(8.628)
$\phi = 0.1$	-6.168
	(11.696)
Switching from $\phi = 0.9$ to $\phi = 0.1$	
Ingroup	-34.938***
	(12.133)
Outgroup	-35.565*
	(13.935)

Note: *p<0.05; **p<0.01; ***p<0.001

2.5.3 Social learning strategies are heterogeneous.

Heterogenity across social learners in the same situation

In the same situation, social learners did not rely on the same strategy but expressed various strategies. To describe social learning heterogeneity across participants in the same situation, we constructed a first-order regression model for each social learner in each treatment condition.

$$y_{ij} = \beta_0 + \beta_1 x_i + \epsilon_i \tag{2.2}$$

We then extracted and compared β_0 and β_1 across social learners.

Figure 2.6 shows the distribution of β_0 and β_1 in each treatment. Based on equation 2.2, β_0 is the estimated allocation when $x_i = 0$. β_1 measures how many tokens the social learner invests in project 1 for each token invested in project 1 by the demonstrator. When β_1 is positive, the social learner favors the same project as the demonstrator; when β_1 is negative, she favors the opposite project. β_0 and β_1 are diverse, suggesting various social learning strategies.

β_0 and β_1 are diverse but not random; they follow a negatively sloped pattern in the four informative treatments. β_0 and β_1 compensate for each other. The higher the β_0, the lower the β_1, and vice-versa. These results suggest that individuals employ diverse strategies leading to similar behaviors. For instance, a high β_0 coupled with a low β_1 indicates a high basic allocation with less reliance on social cues. In contrast, a low β_0 alongside a high β_1 suggests greater emphasis on social information and less on basic allocation. Different combinations can achieve comparable decision-making outcomes. This negative relationship between b1 and b0 highlights

Figure 2.6: Social learning heterogeneity across individuals

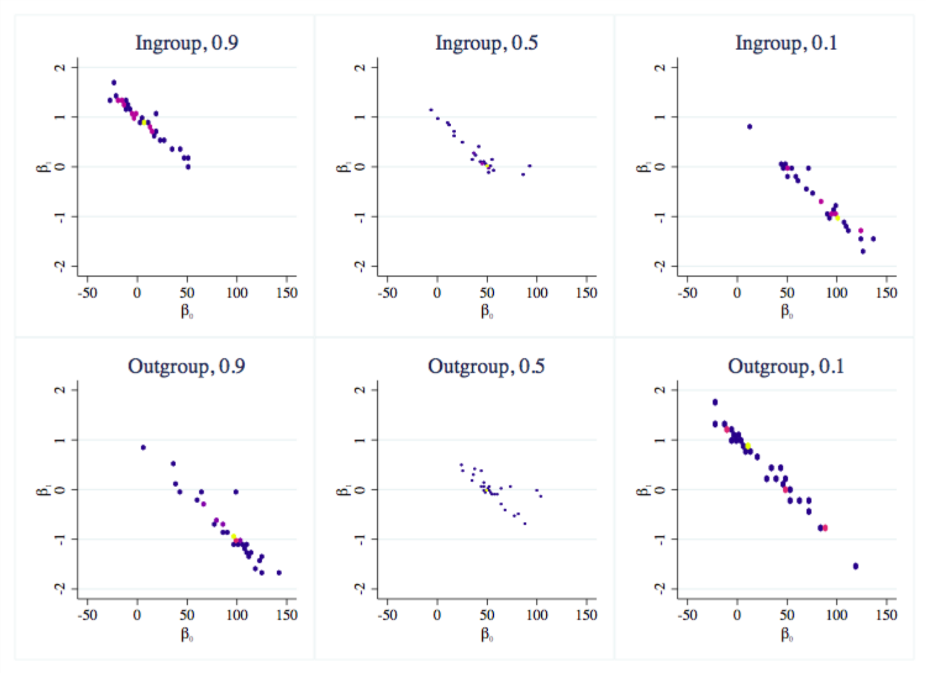

the flexibility and adaptability in social learning. There is no dominant strategy but a spectrum of approaches tailored to individual differences and contexts.

Most, but not all, of the strategies expressed are consistent with the optimal strategy in each treatment. In the Ingroup 0.9 and Outgroup 0.1 treatments, the optimal strategy was to follow the demonstrator's example and allocate the majority of the tokens to the project to which the demonstrator allocated most of her tokens. Namely, any strategy in the upper left corner of the square, $\beta_0 < 50$ and β_1 is positive, is consistent with the optimal strategy. In the Ingroup 0.1 and Outgroup 0.9 treatments, the optimal strategy was the opposite: not following the example of the demonstrator and allocating the majority of the tokens to the project to which the demonstrator allocated the least of her tokens. Strategies consistent with the optimal strategy fall in the bottom right corner, $\beta_0 > 50$, and β_1 is negative. Further, some treatments count more inconsistent strategies than others. The Ingroup 0.9 treatment counts the least inconsistent strategies, followed by the Ingroup 0.1 treatment, the Outgroup 0.9 treatment, and finally, the Outgroup 0.1 treatment. These results suggest that finding the optimal strategy in the Outgroup 0.1 treatment was more challenging for social learners than finding the optimal strategy in other treatments.

Heterogenity across situations for the same social learner

Social learning strategies also vary across situations for the same participant. The same social learner changes her strategy when she moves from one decision setting to another. Most social learners expressed a different strategy when observing an ingroup member than an outgroup member. We built a regression model for each social learner according to equation 2.3. ϕ is a between-subjects dimension and thus does not enter the regression.

$$y_{ij} = \beta_0 + \beta_1 x_i + \beta_2 d_{ij,\text{out}} + \beta_3 x_i * d_{ij,\text{out}} + \epsilon_i \tag{2.3}$$

For each social learner, we compared the allocation across group membership treatments through linear combination 2.4.

$$-\beta_2 d_{ij,\text{out}} - \beta_3 x_i * d_{ij,\text{out}} = 0 \tag{2.4}$$

We computed the linear combination for three levels of x, $x = \{0, 50, 100\}$. The linear combinations produce a β and its associated p-value. β indicates how the allocation varies across group membership treatments for a single individual learner.

Figure 2.7: Social learning heterogeneity across situations for the same individual

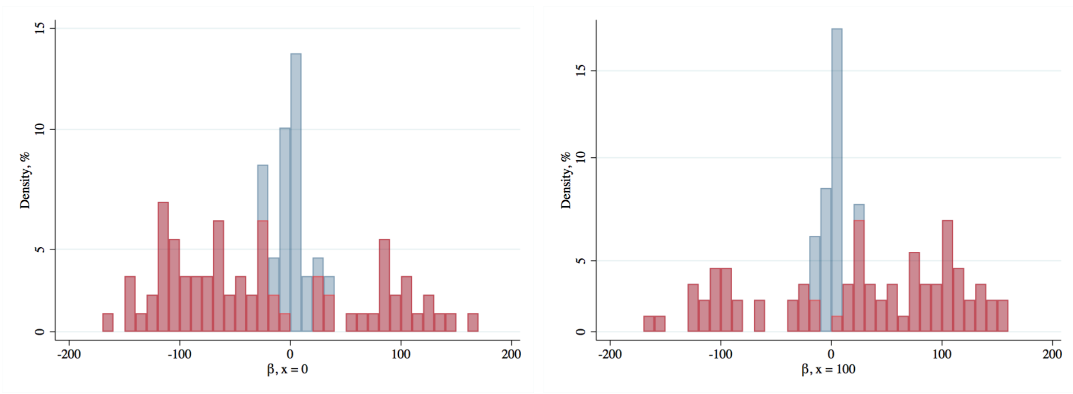

Social learning heterogeneity across situations for the same individual. β can be interpreted as a measure of allocation variation across group membership treatments for a single individual learner. β significant at the 0.05 level are displayed in red.

Figure 2.7 summarizes β for $x = 0$ and $x = 100$. When $x = 0$, 69.81% of the β are significant; these results suggest that 69.81% of the social learners use a different strategy when observing an ingroup member compared to an outgroup member. When $x = 100$, 72.64% of the β are significant; 72.64% of the social learners varied their strategy across group membership treatments.

These data include social learners with all ϕ levels, including under treatment $\phi = 0.5$. Social learners demonstrated facultative adjustment. Namely, they demonstrate strategies rich enough to accommodate all three information orders. Our results suggest tremendous heterogeneity across individuals and situations for the same individual.

We observe an extensive heterogeneity in social learning strategies among participants in the same situation and within the same participant across different situations. This heterogeneity suggests that social learners do not rely on a single, uniform strategy but instead employ various strategies tailored to specific contexts. Social learning strategies are adaptable and flexible.

The observed heterogeneity also implies that multiple strategies can achieve similar outcomes, indicating a compensatory mechanism across different dimensions of social learning. For instance, a high basic allocation with less reliance on social cues (high β_0, low β_1) can be behaviorally equivalent to a strategy that places greater emphasis on social information and less on basic allocation (low β_0, high β_1). This flexibility highlights the complexity of social learning and the capacity of individuals to adjust their strategies based on contextual cues.

Both the heterogeneity observed in our experimental results and the gene-culture coevolutionary simulation suggest an essential role for genetic drift in maintaining diversity in social learning strategies. Drift, through random fluctuations, can lead to the persistence of multiple strategies within a population in the absence of strong selection pressures. Simulation results suggest that several genotype variables are subject to drift while others were selected to be in a particular range without stabilizing on a single value. The experiment's social learning strategies that allow for compensating effects across dimensions further support this idea. Different strategies can produce similar behaviors through compensating effects, allowing drift to maintain polymorphisms in social learning strategies within a gene-culture system in equilibrium.

2.6 Discussion

2.6.1 Third-order social learning strategies

These findings indicate that social learning strategies' current assumed complexity level is far too simple. We should consider integrating multidimensional social learning functions, considering cognitive biases in belief updating, and allowing for heterogeneity in social learning. Assuming simple social learning strategies implies that a social learner observes a successful demonstrator's behavior and chooses based on this observation. The learner can respond to variations in demonstrator choices, but she cannot do more. Although many theories on social learning evolution assume this complexity level, our results join other recent studies indicating that social learning strategies are more complex than this (Efferson et al. 2016). Suppose a social learner can only respond to variation in the distribution of choices. In that case, she cannot change how she responds to any specific distribution given some other variable's value. Social learners, however, can actually do this. We added this kind of complexity to our simulation and our experiment by allowing learners to process information about the demonstrator's group membership and relevance. Social learners can and did respond to the three variables: the observed behavior, the group membership, and the demonstrator's relevance. Both the simulation and experiment show social learning strategies of the third-order, that is, the maximum level of complexity allowed by our setting. People adjust to success-dependent social information in complex ways, including using successful people as negative examples.

Social learning strategies functions can have at least three dimensions and could be even more complex (Efferson et al. 2020). The question is to what extent. Could an infinite number of variables feed into social learning strategies? For example, social learners may also observe the distribution of demonstrators' choices in addition to responding to successful behavior, group membership, and membership relevance. Strategies of this kind would be "fourth-order" and would include both payoff-dependent and frequency-dependent social learning. However, social learners' ability to adjust may not be unlimited. Contrary to our simulation prediction, the participants in our experiments did not fully adjust to "third-order" information. These results could suggest that three information pieces are the limit to adjustment and that social learning strategies would not integrate more than three variables.

2.6.2 Cognitive biases

Second, we should consider cognitive biases in belief updating. Participants appear to process information differently across informationally equivalent treatments. We compared four situations that were equivalent in terms of the value of the information available to subjects. Two situations were consistent with the hypothesized past environment: observing an ingroup member and having a high probability of sharing the environment with ingroup members (i.e., $\phi = 0.9$). The experimental results show that social learners performed similarly when observing an ingroup or compared to an outgroup member. However, they did not fully adjust to third-order information; they performed better under $\phi = 0.9$ than $\phi = 0.1$ even though the two settings carry the same information. The key idea is that persistent asymmetric exposure to specific decision-making settings in the past shaped the evolution of human cognition (Barrett 2014, Haselton et al. 2015). Evolution has retained a higher level of learning complexity only for frequent environments in our ancestral past. If the contemporary settings are similar to ancestral ones, the learner will demonstrate a higher level of complexity in her social learning strategy; she will better integrate information and perform better. Moreover, she will do so in a way that is distinct from the effects of explicit contemporary incentives (Cosmides & Tooby 2013). The asymmetric adjustment across ϕ values is coherent with this interpretation. By asymmetric exposure to certain learning settings compared to others, we evolved a cognition better equipped to deal with these frequent settings.

One significant limitation of our current model is its inability to replicate these experimentally observed cognitive biases. Our model assumes no extra cognitive costs and equal exposure to learning scenarios, failing to account for the evolutionary influences of various exposure to learning scenarios. Incorporating these biases into social learning models would enhance models' realism and predictive power. Ideally, future models should integrate asymmetric exposure to learning situations or cognitive costs. Concretely, this could involve creating models where participants encounter some learning scenarios more frequently than others, thereby simulating asymmetric exposure. For instance, based on our experiment results, agents could be exposed to contexts where ingroup members share optimal projects more often than contexts where ingroup members do not share optimal projects. The asymmetric exposure would reflect the frequency of these scenarios in ancestral environments. A second approach would be adding cognitive costs related to processing certain types of information or information framing to simulate the increased effort required for less familiar or more complex scenarios. These approaches would help capture the evolved cognitive biases, providing a more realistic picture of human social learning strategies.

2.6.3 Social learning heterogeneity

Finally, social learning strategies appeared to be much more heterogeneous than we thought. Participants expressed a tremendous variety of social learning strategies in the same situations, and the vast majority switched strategies across situations. Our gene-culture coevolutionary simulation sheds light on this diversity. Genotypic variance over time illustrates different patterns. Some parameters show decreasing but stable variance, indicating selection pressures that keep these values within a specific range while allowing for some variability. As multiple strategies yield similar outcomes, heterogeneity remains in the population. Other parameters had increasing variance due to genetic drift, which contributes to maintaining diverse strategies. These findings show how both selection and drift preserve strategy diversity. Solutions already exist to integrate part of this heterogeneity in simulations and allow more realistic cultural evolution predictions. Our simulation allows strategies of the third order; other simulations may even enable any strategy coherent with statistical inferences rules (Perreault et al. 2012). The advantage of such simulations is that we do not have to define arbitrary social learning strategies, as we know that we often underestimate their complexity. Such simulations allow for social learning strategies of any complexity level and heterogeneity.

However, our experimental design does not consider correlates such as cognitive abilities or risk preferences, which could influence social learning strategies and help us understand the heterogeneity structure. This omission is a limitation, as these factors could provide insights into the underlying mechanisms driving the observed heterogeneity. Future research should explore how these proximate factors interact with social learning strategies to offer a more comprehensive understanding of social learning variability. Experimental designs could manipulate factors such as cognitive load or risk preferences to observe their impact on social learning behaviors. Exploring these interactions will help build a more accurate understanding of the complexity and heterogeneity of social learning strategies.

Implications for leadership Our research challenges prevailing assumptions in the social learning and leadership literature, which often propose generic strategies like following success, leaders, or ingroup members. We demonstrate that social learning strategies are far more nuanced. Our results show the absence of inherent biases in decision-making when information is lacking and the possibility of adopting strategies that run counter to following the leader. This implies that leadership should not be seen as a one-directional influence. Leaders sometimes serve as negative examples, guiding what not to do. Our research underscores the complexity of social learning in organizational settings and calls for a shift away from 'follow-the-leader' models towards a more dynamic understanding of social learning, where individual and contextual factors play significant roles.

The above findings about social learning complexity have significant implications for management and organizational practices. Understanding that followers may not always imitate successful leaders if they perceive those leaders as irrelevant to their context can help managers design more effective leadership strategies. For example, when introducing new tools or practices, leaders should consider the relevance of their demonstrations to their followers' specific situations. Additionally, the heterogeneity in social learning strategies suggests that a one-size-fits-all approach to leadership may be ineffective. Leaders should tailor their strategies to accommodate the diverse ways followers learn and adapt. This study highlights the importance of considering multiple variables in social learning, such as success, group membership, and relevance, to enhance organizational performance.

2.6.4 Limitations and avenues for future research

While our findings emphasize social learning strategies' complexity and diversity, our study has several limitations. First, our study relies on a sample of European university students, that is, people from Western, Educated, Industrialized, Rich, and Democratic countries (WEIRD; Henrich et al. (2010)). Concerns exist about how these results would translate to non-WEIRD populations. Research indicates that organizational theory findings from WEIRD samples may not accurately reflect behaviors in more diverse populations (Banks 2023, Pitesa & Gelfand 2023). Additionally, large-scale studies have shown that while some psychological phenomena are consistent across different settings, others exhibit variability influenced by cultural contexts (Schimmelpfennig et al. 2024). While we primarily explored evolved biases that are theoretically expected to be consistent across diverse cultures, Schimmelpfennig et al. (2024) demonstrates that without a robust theoretical framework, predicting which aspects of social learning are universally consistent and which are culturally specific remains challenging. To address these concerns, future research should aim to replicate this study in varied cultural settings. Such cross-cultural studies would provide a deeper understanding of the universality or specificity of these strategies, refining our theoretical frameworks and enhancing the generalizability of our findings.

Second, we focus on success bias in isolation and do not consider interaction with other biases, such as conformity. These biases often intersect in real-world scenarios and could influence each other in complex ways. For instance, how success bias interacts with the tendency to conform to majority behavior could significantly alter the observed learning strategies. Success bias might drive individuals to imitate the most successful members, while conformity bias could lead them to adopt more frequent behaviors within their group. The interplay between these biases could result in nuanced strategies where individuals weigh the relative success of behaviors against their prevalence within the group. For example, in environments where successful behaviors are also the majority behaviors, the biases might reinforce each other, leading to strong cultural norms. Conversely, when successful behaviors are not widely adopted, individuals might face a trade-off between imitating success and conforming to the majority, potentially leading to diverse and context-dependent learning strategies. Future studies integrating multiple biases could offer more comprehensive insights into social learning flexibility and complexity. Methodologically, this could involve developing agent-based models that simulate environments with both payoff-dependent and frequency-dependent strategies available. Additionally, experimental studies could be designed to test hypotheses about how different biases interact, such as creating conditions where the success of behaviors and their frequency within the group are manipulated independently.

Conclusion Altogether, given the limitations of our approach, our contribution is threefold. First, our study advocates multidimensional social learning strategies, moving beyond traditional assumptions of simple, success-based imitation. Social learners do not simply imitate successful behavior. They also consider additional information like group membership and relevance. Second, our research highlights the importance of cognitive biases in belief updating. In our study, participants performed differently across information-equivalent situations. Despite similar information value, social learners responded differently based on whether they were similar to ingroup members. These systematic differences suggest that asymmetric exposure to specific decision-making settings has shaped our cognitive evolution. This finding challenges traditional views on decision-making and underscores the need to consider evolutionary influences and cognitive biases when developing theories and simulations in social learning. Finally, our findings underscore the remarkable heterogeneity in social learning strategies. We observed various strategies used by participants in identical situations and by different participants in the same context. This heterogeneity challenges simulations that rely on simple, generic social learning strategies. Our research supports the use of more sophisticated social learning simulations and leadership success-based techniques.

Bibliography

Argote, L., Ingram, P., Levine, J. M. & Moreland, R. L. (2000), 'Knowledge transfer in organizations: Learning from the experience of others', *Organizational behavior and human decision processes* **82**(1), 1–8.

Argote, L. & Levine, J. M. (2020), *The Oxford handbook of group and organizational learning*, Oxford University Press.

Bandura, A. & Walters, R. H. (1977), *Social learning theory*, Vol. 1, Englewood cliffs Prentice Hall.

Banks, G. C. (2023), 'Eight puzzles of leadership science', *The Leadership Quarterly* **34**(4), 101710.

Barrett, H. C. (2014), *The shape of thought: How mental adaptations evolve*, Oxford University Press.

Bellamy, A., McKay, R., Vogt, S. & Efferson, C. (2022), 'What is the extent of a frequency-dependent social learning strategy space?', *Evolutionary Human Sciences* **4**, e13.

Boyd, R. & Richerson, P. J. (1985), *Culture and the Evolutionary Process*, University of Chicago Press.

Cavalli-Sforza, L. L. & Feldman, M. W. (1981), *Cultural transmission and evolution: A quantitative approach*, Princeton University Press.

Chen, D. L., Schonger, M. & Wickens, C. (2016), 'otree—an open-source platform for laboratory, online, and field experiments', *Journal of Behavioral and Experimental Finance* **9**, 88–97.

Cosmides, L. & Tooby, J. (2013), 'Evolutionary psychology: New perspectives on cognition and motivation', *Annual review of psychology* **64**, 201–229.

Efferson, C., Lalive, R., Cacault, M. P. & Kistler, D. (2016), 'The evolution of facultative conformity based on similarity', *PLOS ONE* **11**(12), 1–28.
URL: *https://doi.org/10.1371/journal.pone.0168551*

Efferson, C., Lalive, R., Richerson, P. J., McElreath, R. & Lubell, M. (2008), 'Conformists and mavericks: the empirics of frequency-dependent cultural transmission', *Evolution and Human Behavior* **29**(1), 56–64.

Efferson, C., Vogt, S. & Fehr, E. (2020), 'The promise and the peril of using social influence to reverse harmful traditions', *Nature human behaviour* **4**(1), 55–68.

Ehret, S., Vogt, S., Hefti, A. & Efferson, C. (2021), 'Leading with the (recently) successful? performance visibility and the evolution of risk taking', *University of Zurich, Department of Economics, Working Paper* (382).

Faessler, L. & Efferson, C. (2019), 'Heterogeneity in social learning', `https://osf.io/ukt2c/`. Preregistration.

Granovetter, M. (1978), 'Threshold models of collective behavior', *American journal of sociology* **83**(6), 1420–1443.

Greiner, B. (2004), 'An online recruitment system for economic experiments'.

Haselton, M. G., Nettle, D. & Murray, D. R. (2015), *The Evolution of Cognitive Bias*, John Wiley Sons, Ltd, chapter 41, pp. 1–20.
URL: *https://onlinelibrary.wiley.com/doi/abs/10.1002/9781119125563.evpsych241*

Henrich, J. & Gil-White, F. J. (2001), 'The evolution of prestige: Freely conferred deference as a mechanism for enhancing the benefits of cultural transmission', *Evolution and human behavior* **22**(3), 165–196.

Henrich, J., Heine, S. J. & Norenzayan, A. (2010), 'Most people are not weird', *Nature* **466**(7302), 29–29.

Kendal, R. L., Boogert, N. J., Rendell, L., Laland, K. N., Webster, M. & Jones, P. L. (2018), 'Social learning strategies: Bridge-building between fields', *Trends in cognitive sciences* **22**(7), 651–665.

Laland, K. N. (2004), 'Social learning strategies', *Animal Learning & Behavior* **32**(1), 4–14.

Levitt, B. & March, J. G. (1988), 'Organizational learning', *Annual review of sociology* **14**(1), 319–338.

Manski, C. F. (2000), 'Economic analysis of social interactions', *Journal of economic perspectives* **14**(3), 115–136.

Manz, C. C. & Sims Jr, H. P. (1981), 'Vicarious learning: The influence of modeling on organizational behavior', *Academy of Management Review* **6**(1), 105–113.

Matthews, S. H., Kelemen, T. K. & Bolino, M. C. (2021), 'How follower traits and cultural values influence the effects of leadership', *The Leadership Quarterly* **32**(1), 101497.

McElreath, R., Bell, A. V., Efferson, C., Lubell, M., Richerson, P. J. & Waring, T. (2008), 'Beyond existence and aiming outside the laboratory: estimating frequency-dependent and pay-off-biased social learning strategies', *Philosophical Transactions of the Royal Society B: Biological Sciences* **363**(1509), 3515–3528.

McKay, R. & Efferson, C. (2010), 'The subtleties of error management', *Evolution and human behavior* **31**(5), 309–319.

Mesoudi, A., Chang, L., Dall, S. R. & Thornton, A. (2016), 'The evolution of individual and cultural variation in social learning', *Trends in ecology & evolution* **31**(3), 215–225.

Morgan, T., Rendell, L., Ehn, M., Hoppitt, W. & Laland, K. (2012), 'The evolutionary basis of human social learning', *Proceedings of the Royal Society B: Biological Sciences* **279**(1729), 653–662.

Muthukrishna, M., Morgan, T. J. & Henrich, J. (2016), 'The when and who of social learning and conformist transmission', *Evolution and Human Behavior* **37**(1), 10–20.

Myers, C. G. (2018), 'Coactive vicarious learning: Toward a relational theory of vicarious learning in organizations', *Academy of Management review* **43**(4), 610–634.

Oc, B., Chintakananda, K., Bashshur, M. R. & Day, D. V. (2023), 'The study of followers in leadership research: A systematic and critical review', *The Leadership Quarterly* **34**(1), 101674.

Offerman, T. & Sonnemans, J. (1998), 'Learning by experience and learning by imitating successful others', *Journal of economic behavior & organization* **34**(4), 559–575.

Perreault, C., Moya, C. & Boyd, R. (2012), 'A bayesian approach to the evolution of social learning', *Evolution and Human Behavior* **33**(5), 449–459.

Pitesa, M. & Gelfand, M. J. (2023), 'Going beyond western, educated, industrialized, rich, and democratic (weird) samples and problems in organizational research', *Organizational Behavior and Human Decision Processes* **174**, 104212.

Schimmelpfennig, R., Spicer, R., White, C. J., Gervais, W., Norenzayan, A., Heine, S., Henrich, J. & Muthukrishna, M. (2024), 'The moderating role of culture in the generalizability of psychological phenomena', *Advances in Methods and Practices in Psychological Science* **7**(1), 25152459231225163.

Uhl-Bien, M., Riggio, R. E., Lowe, K. B. & Carsten, M. K. (2014), 'Followership theory: A review and research agenda', *The leadership quarterly* **25**(1), 83–104.

Young, H. P. (2009), 'Innovation diffusion in heterogeneous populations: Contagion, social influence, and social learning', *American economic review* **99**(5), 1899–1924.

2.7 Appendix

2.7.1 Gene-culture coevolutionary simulation

How private information modifies strategies. In the experiment, we distinguish social learning from individual learning. This separation is essential if we want to distinguish the effect of social learning from the effect of individual learning. However, outside this controlled setting, both happen at the same time. To understand the combined effects, we include individual learning in the simulation. We examine here how private information influences social learning strategies, which we now call strategies as social learning mixes with individual learning.

Figure 2.8 illustrates the average strategies across the different parameters' combinations. Each red line in the graph represents the average strategy observed in the last generation of a simulation. The black dotted line shows the overall average strategy across all simulations. The reliability of the private signal can take the following values $\{0.5, 0.6, and 0.9\}$. A reliability of 0.5 means that the signal carries no information. Agents can only learn socially. We focus here on scenarios where the private signal is valuable ($\sigma \in \{0.6, and 0.9\}$) and how the allocation strategies differ when the private signal is valuable compared to the situations where the signal is meaningless. Focusing first on the first row, when the environment is stable, and the private signal indicates that the optimal project is Project 0 ($\phi = 0.1, s = 0$), we observe that allocations decrease as the reliability of the private signal increases. Conversely, in the second row, when the private signal indicates that the optimal project is Project 1 ($\phi = 0.1, 1 = 0$), allocations increase as the reliability of the private signal increases. However, even in the presence of meaningful private information, allocations still vary according to the social information pieces, the observed allocations, and whether the demonstrator is an ingroup or outgroup member. Even when the private signal is highly reliable ($\sigma = 0.9$), allocations differ across Ingroup and Outgroup treatments and the observed allocation levels. These results suggest that agents combine social and private information in a stable environment. Even when reliable, private information modifies the social learning strategy rather than replacing it.

Moving to the third and fourth rows, when the environment is unpredictable and social information has no value ($\phi = 0.5$), we observe that allocations are shifted in the direction indicated by the private signal as the reliability of the private signal increases. When the private signal indicates that the optimal project is Project 0 ($s = 0$), allocations decrease. When the private signal indicates that the optimal project is Project 1 ($s = 1$), allocations increase. Notice

Figure 2.8: Average strategies across treatments and simulations

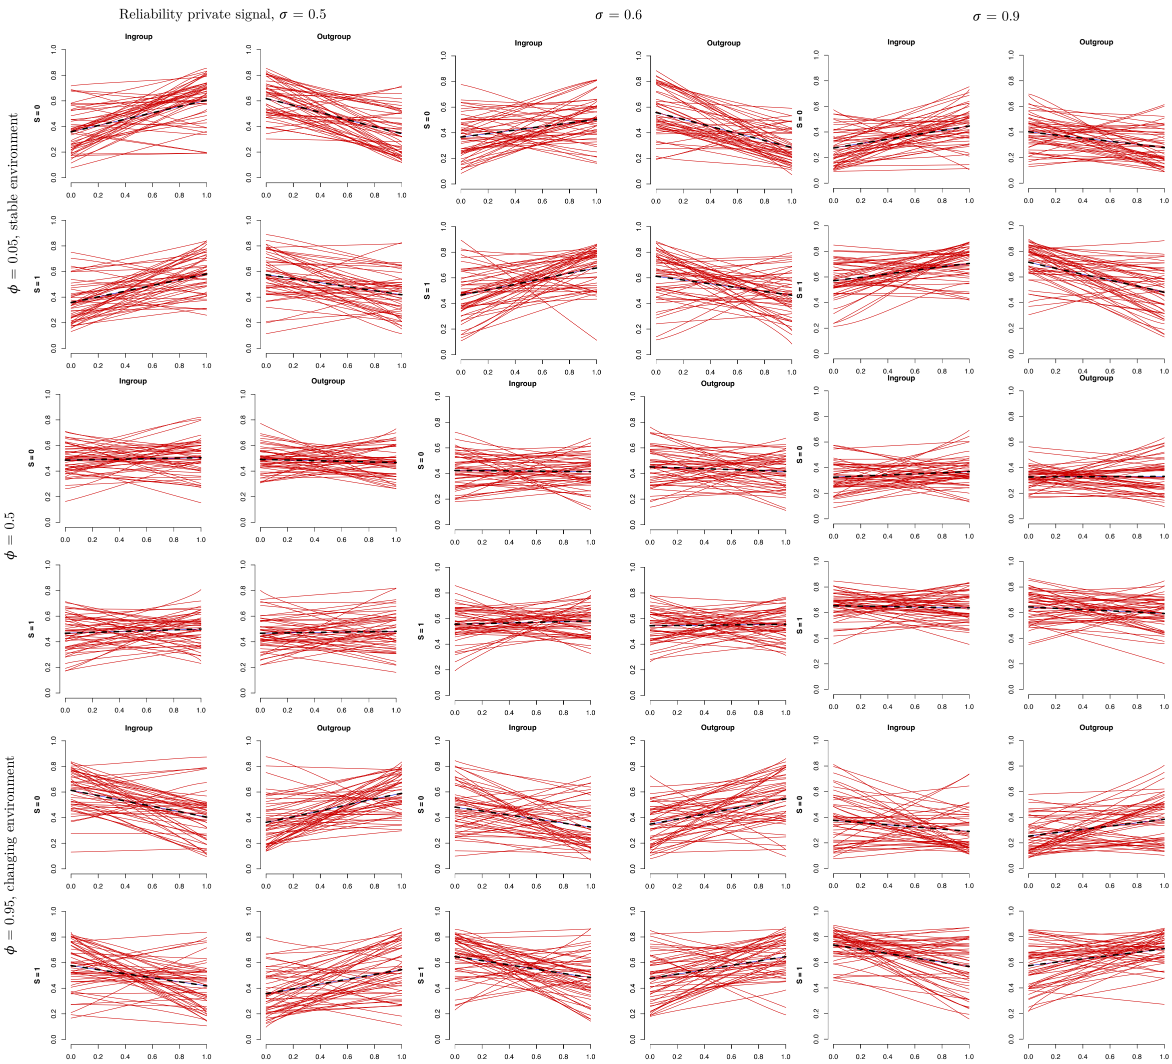

Notes: The graph shows the average strategies when the system stabilizes. The observed allocation from the successful agent to Project 1 is on the x-axis. The allocation of the focal agent to Project 1 is on the y-axis. Each red line represents the average strategy in the last generation of a simulation. The black dotted line represents the overall average strategy across all simulations. ϕ indicates the stability of the environment. $\phi = 0.95$ is a highly changing environment. $\phi = 0.05$ is a stable environment. $\phi = 0.5$ indicates a moderately changing environment. The reliability of the private signal can take the following values: 0.5, 0.6, and 0.9. When private signal reliability = 0, private information has no value. s is the private signal and indicates whether the winning project is 1 ($s = 1$) or project 0 ($s = 0$). Ingroup (outgroup) indicates whether the agent observed a successful demonstrator from ingroup (outgroup).

that even in that scenario, where social information has no value and private information is the only information available, we do not observe an immediate shift in the strategies. Instead, allocations increase (decrease) in a degree that seems related to the reliability of the private signal, a small increase (decrease) when $\sigma = 0.6$ and a bigger increase (decrease) when $\sigma = 0.9$.

Last, we observe similar patterns when the environment is highly changing ($\phi = 0.95$). Allocations are shifted in the direction indicated by the private signal as the reliability of the private signal increases. As in the stable environment ($\phi = 0.05$), allocations also vary according to observed allocations and whether the demonstrator is an ingroup or outgroup member. Although we could have imagined that private information would be more valuable in a changing environment, the observed strategies in stable and unstable environments do not suggest that agents balanced social information and private information differently across these environments.

The results suggest that even when reliable, private signals do not replace social learning but adjust it. In all environments, social learning strategies are modified in the direction indicated by the private signal. Further, the weight given to private information increases as its reliability increases. This smooth shift suggests that agents blend individual and social learning. They adapt their strategies based on the reliability of private information and the observed social context rather than switching from one form of learning to another as the information becomes reliable. Again, the evolution selected the most complex version of learning strategies, where agents can adjust to social and private information.

2.7.2 Experiment

Screenshots

Figure 2.9: Decision screen for individual learners

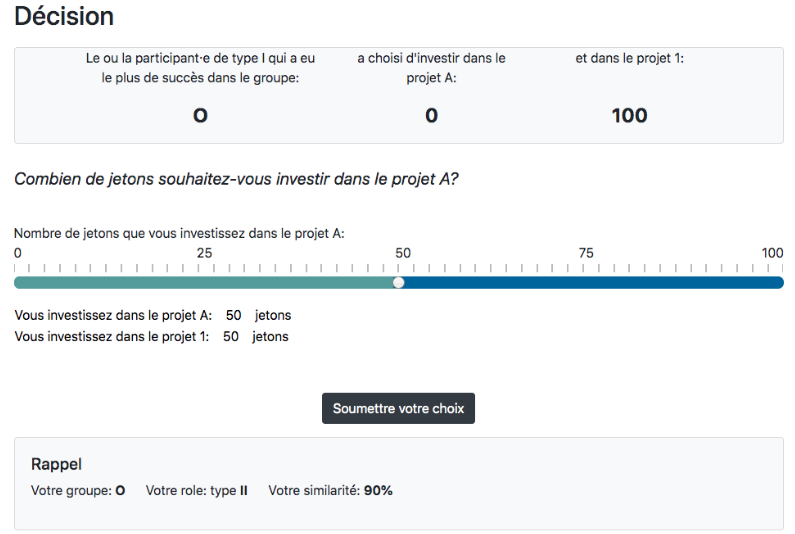

Notes: This figure illustrates what the decision screen looks like for an individual learner. The slider and a numerical summary of the decision are displayed in the middle of the screen. Using the slider, participants have to decide how many tokens they would like to allocate to Project A. Below, they can click "Soumettre votre choix" to submit their decision. At the bottom of the screen, participants find a reminder of their group membership and type (individual learner). These information pieces are the same throughout the session.

Figure 2.10: Decision screen for social learners

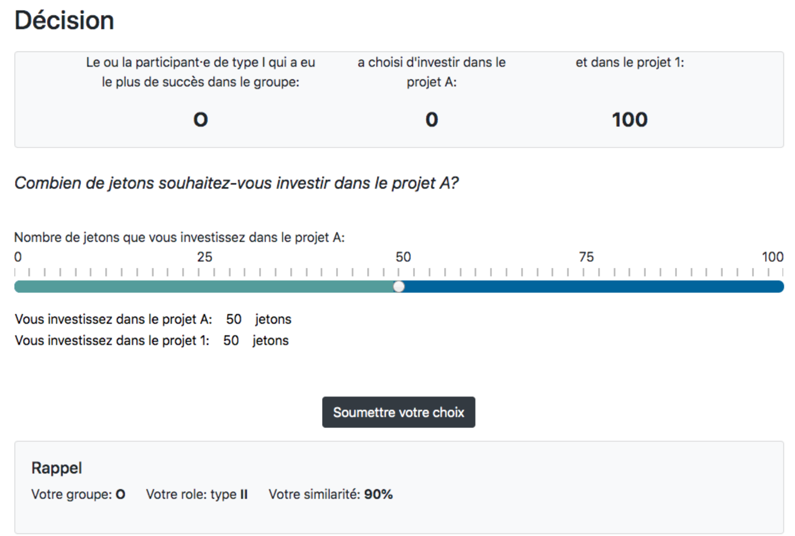

Notes: This figure illustrates the decision screen for a social learner. On top of the screen, social information is displayed: the demonstrator's group and allocation decision. In this case, the participant observes a demonstrator from the circle group who decided to invest all of their endowment in Project 1. Participants know the demonstrator is the most successful person in the group. The slider and a numerical summary of the decision are displayed in the middle of the screen. Below, participants should click on "Soumettre voter choix" to submit their decision. At the bottom of the screen, participants find a reminder of additional information: their group membership, their type (social learner), and the probability that they share the same winning project as ingroup members. These information pieces are the same throughout the session.

Chapter 3

Collective problem solving:

Turnover is not sufficient to avoid over-exploitation

Lisa Faessler

Faculty of Business and Economics, University of Lausanne, Switzerland

Abstract

Teams in various organizational settings often rapidly transition from exploration to the exploitation of known solutions, potentially at the cost of more innovative or optimal outcomes. Further, fully connected teams perform worse than moderately connected teams. This experiment aims to determine whether turnover and the arrival of new members can offset these tendencies and facilitate the creation of more diverse solutions. In a lab experiment, participants solved a complex combinatorial task. Two forms of disruption were introduced. First, some teams experienced changes in team composition, while others remained stable. Second, among teams that had not reached the highest-performing solution, some were selectively informed about the existence of superior solutions, while others were not. The results indicate an increase in search distance following the disruptions in all treatment conditions. However, the increase was modest and temporary. Interestingly, changes in team composition did not create a more significant disruption compared to changes in reference points. Contrary to expectations, the increased exploration was insufficient to create alternative solutions and improve performance. Instead, payoffs decrease significantly after the treatment in all conditions before returning to their initial improving trend. These findings suggest that disruptions can trigger short-term exploration but do not reliably enhance overall performance.

Keywords: social learning, team decision-making, problem-solving, exploration versus exploitation, turnover, changing teams

3.1 Introduction

Teams in various organizational settings often rapidly transition from exploration to the exploitation of known solutions, potentially at the cost of more innovative or optimal outcomes. To solve complex problems, firms, teams, and individuals engage in search behaviors (Simon 1957). They alternate between exploitation, which means applying known and tested solutions to achieve stability and efficiency, and exploration, which involves experimenting with novel solutions to enhance innovation and adaptability (March 1991). Exploitation and exploration often compete for the same limited resources, thus implying a tradeoff. Exploitation uses these resources to improve what is already known for immediate returns, while exploration invests them in new, untested ideas for potential long-term gains. While the optimal balance between the two behaviors depends on the context, a pervasive finding is that firms and teams tend to over-emphasize exploitation (Denrell & March 2001). Further, fully connected teams perform worse than moderately connected teams (Derex & Boyd 2016).

At first, an increased connectivity leading to worse performance may seem counterintuitive. Yet, this paradox resolves when we recognize that innovation arises from balancing three critical levers: group size and connectivity, transmission fidelity, and cultural trait diversity (Muthukrishna & Henrich 2016, Schimmelpfennig et al. 2022). Each lever can enhance or hinder innovation. First, group size must surpass a certain threshold to ensure a variety of learning models and sufficient diversity (Kline & Boyd 2010, Muthukrishna et al. 2014). However, if the team is too connected, people tend to over-rely on social information and homogeneity (Derex & Boyd 2016). This loss of diversity stifles creativity and innovation. The over-reliance on social cues leads the highly connected teams to underperform.

The second lever, transmission fidelity, refers to the accuracy and quality of information flow within a team. High transmission fidelity enhances teaching and information dissemination, leading to better learning outcomes. However, accurate transmission also means fewer opportunities for novel variations and an increase in homogeneity. Some distortion in transmission is beneficial for serendipity and innovation.

The third lever, cultural trait diversity, involves the variety of ideas, values, and perspectives within a team (Schimmelpfennig et al. 2022). More variety leads to more opportunities for recombination and novel ideas. However, cultural trait diversity also brings coordination problems and inequality of outcomes (Muthukrishna 2020). Despite these challenges, diversity provides the raw material for creative solutions and remains the best potential fuel for innovation.

In this study, I asked whether introducing diversity in teams could help mitigate the downsides of highly connected teams. By increasing diversity, newcomers could reduce conformity and overreliance on social cues. In that sense, turnover, while generally negative due to its disruptive nature, may have potential virtues. For example, a specific turnover rate increases efficiency in bird population (Chimento et al. 2021). Turnover might also be helpful in human populations.

Turnover is often, quite rightly, viewed negatively. Turnover disrupts operations, destabilizes organizational routines, and slows organizational learning. Further, it depletes human and social capital, typically leading to loss of knowledge, lower productivity, and reduced profits. Effective communication and coordination can become more complex as new team members integrate and adjust to the team's existing processes and culture. However, turnover might introduce new perspectives and knowledge into a team. Both disruption and fresh knowledge could potentially lead to increased exploration. For instance, research in the biotech industry has demonstrated that turnover among top scientists correlates with increased exploration (Tzabbar & Kehoe 2014). However, the underlying mechanism remains unknown. The increase in exploration related to turnover could be due to many factors.

First, turnover disrupts operations and destabilizes organizational routines, creating opportunities for the team to rethink and innovate. Further, this period of disruption can sometimes increase risk-taking behavior. Second, losing a team member or a leader can alter team dynamics by increasing uncertainty. The perceived uncertainty can then increase risk-taking and exploration. Third, the introduction of new team members adds diversity to the group. These new members bring novel experiences and viewpoints. They can challenge existing norms and inspire alternative solutions. This diversity in the learning models fosters an environment where exploration is more likely to occur.

In this study, I test whether turnover can increase exploration in connected teams and help them break off from this over-exploitation equilibrium. Further, additional treatments allow me to disentangle the mechanical effect of turnover from (1) changes in reference points and (2) psychological effects related to the departure of the top-performer. A "reference point" is a standard or baseline individuals use to compare their current situation. Shifts in this point can influence decision-making behaviors significantly (Tversky & Kahneman 1992, Bromiley 2010). Previous studies confirm this pattern of success-induced exploitation and failure-induced exploration in solving complex tasks (Billinger et al. 2014, Giannoccaro et al. 2020). By shifting the reference point to the highest achievable performance, I aim to create a perceived failure situation for the participants, similar to the perceived uncertainty at the departure from the leader, thereby increasing their willingness to explore (Koop & Johnson 2012). The second treatment involves the departure of a top-performer to investigate whether knowing that a high-performing member is leaving prompts the team to explore more.

The results suggest that turnover and the other treatments lead to a small and brief increase in exploration. Surprisingly, the effect of turnover does not significantly differ from the effect of the more straightforward disruption, the change in reference point. However, in all treatments, teams quickly revert to the exploitation equilibrium. Further, the treatments do not influence performance in the long run. The comparison of the human participant's performance with benchmarks simulated via Monte Carlo simulation suggests that human participants achieved performance similar to connected teams doing a local search. The simulation results suggest that increased exploration could have led to a higher performance. The exploration period induced by the experiment's treatments was too short to lead the teams to discover better alternative solutions. Moving away briefly from over-exploitation is relatively easy, but maintaining exploration for a sufficient duration to enhance performance is challenging.

3.2 Solving complex problems in teams

3.2.1 Search behaviors: an alternate of exploitation and exploration

Complex problems involve multiple interdependent decisions, and the best solution is often unknown. To tackle these problems, individuals must engage in search behaviors and alternate between exploitation and exploration (Simon 1957, March 1991). The search distance measures how different a new solution is from the previous one and distinguishes exploitation from exploration. In exploitation, the distance is short. Individuals reuse or modify slightly the previous solution. Since this solution is similar to the previous one, its payoff is known, making it a safer choice. On the other hand, exploration involves deviating further from the last solution by changing more components. The new solution is markedly different from the original, and its payoff is unknown. The solution could be a breakthrough or a setback, depending on the complexity of the problem, the landscape of possible solutions, and luck. Exploration, due to its uncertain payoff, typically entails greater risk. Exploitation and exploration are generally seen as mutually exclusive. At each decision point, individuals choose whether to exploit or explore.

For instance, in software development, the decision at each stage is whether to keep the current version of the software or modify it and to what extent.

3.2.2 Success-induced exploitation and failure-induced exploration

When individuals tackle problems alone instead of in a team, success induces exploitation, while failure induces exploration (Billinger et al. 2014). If the current solution outperforms the previous one, individuals tend to exploit it. Conversely, they tend to explore if the current solution is less effective. Considering the previous solution as the reference point, a current, less effective solution places the individual in the loss domain, making them more likely to take risks and explore. Conversely, if the current solution is more effective, the individual is prone to less risk-taking and more likely to exploit this success. Exploration is perceived as riskier because the payoff is unknown, whereas exploitation promises a known reward. However, suppose the individual currently has the worst possible combination. In that case, exploration becomes the only way to improve their payoff and, thus, less risky than sticking with and exploiting the current combination. Apart from this specific scenario, exploration consistently involves more risk than exploitation. Facing complex tasks alone, individuals typically exploit during success and explore during failure, regardless of the task's complexity level.

3.2.3 Connected teams over-exploit.

Complex tasks are often assigned to teams, not just individuals. Understanding how teams tackle complex tasks and their search behaviors is crucial (Wagner III et al. 2012, Yoon & Kayes 2016). Similar search behavior patterns emerge in teams as with individuals: success induces exploitation, and failure induces exploration (Kostopoulos & Bozionelos 2011, Goldstone et al. 2013, Døjbak Håkonsson et al. 2016, Giannoccaro et al. 2020). However, team dynamics also influence these behaviors. Highly connected teams, where information exchange among members is frequent, tend to explore less (Derex & Boyd 2016). These connected teams quickly reach an over-exploitation equilibrium and miss on better solutions. This raises a significant challenge for innovation. Working in teams might hinder the generation of more creative and effective solutions. How can connected teams avoid getting stuck in this over-exploitation equilibrium and instead explore alternative solutions?

3.2.4 Introducing diversity through turnover

Team composition changes, or turnover, offer a potential tool to counteract over-exploitation. Turnover, especially turnover of top-performing individuals, entails significant costs for businesses, including recruitment expenses, knowledge loss, and the time required for newcomers to achieve full productivity (Abelson & Baysinger 1984, Aime et al. 2010, Kwon & Rupp 2013). However, turnover could also present an opportunity. Newcomers might introduce diversity, innovative ideas, and varied expertise that can counteract the conformity induced by team connectivity (Schimmelpfennig et al. 2022). In the context of search behaviors, newcomers could help foster the exploration of alternative solutions and facilitate the shift away from over-exploitation. Tzabbar & Kehoe (2014) analyze longitudinal data about star scientist turnover in the biotechnology industry. They show that star scientist turnover creates opportunities for the firm to search beyond existing knowledge boundaries and increases exploration.

In this laboratory-based study, I test if team turnover leads to increased exploration and can offset the expected detrimental tendency of highly connected teams to over-exploit. Additionally, I disentangle which turnover component leads to increased exploration. I thus differentiate between two conditions: one where team members are informed about the departure of the top-performer and another where they are not. The constant across both conditions is the departure of the top-performer, but the critical variable is the team's awareness of this change. The objective is to understand whether the knowledge of the top-performer's exit psychologically impacts the team's search behaviors. Given the controlled setting of a laboratory task, where the departure of a top-performer might seemingly have a lower impact compared to real-world teams, detecting any behavioral differences under these conditions would provide compelling evidence of the psychological effect related to the departure of a top-performer.

Hypothesis 1a. *Turnover increases search distance compared to conditions where groups are stable.*

Hypothesis 1b. *By extension, turnover increases performance compared to conditions where groups are stable. The assumption is that teams operate at suboptimal over-exploitation levels, and any shift towards more exploration is expected to increase their overall performance.*

Hypothesis 2. *Information about the departure of the top-performer increases search distance compared to similar scenarios where the information is not provided.*

3.2.5 Artificially alter reference point to increase exploration.

To further distinguish the turnover components involved in increased exploration, I consider whether changing the reference point could lead individuals to explore other solutions. The idea is to artificially alter the reference point, creating a perceived situation of failure similar to the one that turnover could induce and encourage risk-taking behaviors and exploration (Koop & Johnson 2012). In previous studies, researchers assume that the reference point is the payoff of the last solution or the best payoff obtained in the sequence (Billinger et al. 2014, Giannoccaro et al. 2020). In this study, I introduce a new, higher reference point to motivate participants to explore. I use the highest achievable performance as the new reference point. Participants are informed that they can attain solutions yielding higher payoffs. This shift in reference points could lead participants to see their current performance as a failure and induce exploration. They could break away from the sub-optimal exploitation equilibrium and find more efficient solutions. I expect individuals and teams to adjust search behaviors following the new reference point disclosure.

Hypothesis 3a. *Information about the existence of better solutions increases search distance compared to conditions where no such information is provided.*

Hypothesis 3b. *By extension, information about the existence of better solutions increases performance compared to conditions where no such information is provided.*

3.2.6 Studying complex-problem solving in the lab

This study explores whether team turnover can help connected teams break the expected over-exploitation equilibrium typically seen in complex problem-solving. For this purpose, I employ the NK landscape task framework (Csaszar 2018). The NK landscape task is a widely used method for studying complex problem-solving in a laboratory setting (for review, see, for example, Baumann et al. (2019)). Participants are challenged to discover the optimal combination of components, like cracking a secret code. An NK algorithm determines the value of each combination. Participants engage in trials to identify the most effective combination possible, navigating through various configurations to uncover the optimal solution. Optimal strategies in such tasks are complex and often ex-ante unknown. Therefore, I built a Monte Carlo simulation to benchmark teams' performance in the experiment.

3.3 Predictive benchmarks using Monte Carlo simulations

This study explores how turnover influences search behaviors in a controlled lab setting. To set performance benchmarks, I conducted a Monte Carlo simulation. This simulation models individual search behaviors under different conditions by varying (1) the likelihood of exploration (ϕ) and (2) the possibility of learning from others (γ). Participants face different probabilities of exploring new strategies versus sticking with known ones ($\phi \in \{0, 0.1, 0.5, 1\}$). Additionally, agents experience two scenarios: they can be isolated, without access to social learning (γ=0), or they can be connected to team members and learn from their solutions (γ=1). By manipulating exploration likelihood and connectedness, I produce three primary performance benchmarks for the experimental results. The primary benchmarks of interest include:

- **The low benchmark** represents a random walk, where agents only search locally and are not connected ($\phi = 0, \gamma = 0$).
- **The medium benchmark** involves random walks within connected teams, where agents only search locally but can learn from each other ($\phi = 0, \gamma = 1$).
- **The high benchmark** represents the optimal scenario, where connected teams balance local search and exploration ($\phi = 0.5, \gamma = 1$).

These benchmarks provide a comprehensive range of expected outcomes, offering a clear basis for comparison with our experimental results.

3.3.1 Simulation setup

Initial conditions An NK landscape is generated and assigns a specific payoff to each possible combination of components. A new NK landscape is created at the beginning of each simulation to introduce variability and ensure robustness in the results. Each landscape is characterized by 10 elements ($N = 10$) and 5 interactions ($K = 5$), matching the complexity of the experimental setup. As in the experimental setup, each team in the simulation consists of 3 individuals ($n = 3$), and the simulation runs over 20 decision-making periods, referred to as rounds ($t_max = 20$). At $t = 1$, each agent is assigned a random combination of components. The payoffs for these combinations are calculated based on the generated NK landscape. This setup establishes the initial conditions.

Table 3.1: Simulation parameters and role

Parameters	Role	Data condition
ϕ	**Exploration probability**. Controls the probability that agents engage in exploration, defined as changing two or more combination components. $\phi = 0$ implies local search each round and $\phi = 1$ exploration each round.	$\phi \in \{0, 0.1, 0.5, 1\}$
γ	**Team connectedness**. Controls whether agents have access to the combinations and payoff of others. $\gamma = 0$ means teams are disconnected. Social learning is not available. $\gamma = 1$, teams are connected. Social learning is available.	$\gamma \in \{0, 1\}$
n	Number of individuals in the team	n = 3
t_max	Maximum number of rounds or trials	t_max = 20
r_max	Maximum number of simulation runs	r_max = 100
N	Number of elements in NK landscape	N = 10
K	Number of interactions in NK landscape	K = 5

Decision-making rule For subsequent rounds, agents update their strategies based on two key factors: the exploration probability (ϕ) and the level of team connectedness (γ). Each agent starts a new round with the combination they had in the previous round. The decision-making process begins with applying the Lazer and Friedman rule (Lazer & Friedman 2007). Decisions proceed in two stages: first, social learning, and second, individual learning. If the agents are connected ($\gamma = 1$), agents start by learning socially. In this phase, agents observe the strategy combination that yielded the highest payoff in their group in $t - 1$. If this observed payoff exceeds their current payoff, they adopt this more successful combination.

If social learning is unsuccessful, or the agents are not connected ($\gamma = 0$), the process shifts to individual learning. During individual learning, the agents balance exploration and exploitation depending on the exploration probability ($\phi \in \{0, 0.1, 0.5, 1\}$). When $\phi = 0$, agents can only exploit and change one random component of their combination. When $\phi = 1$, agents fully explore. They can change between 2 and 10 components of their combination. The specific components and the number of changes are selected randomly. When $\phi > 0 and < 1$, agents probabilistically decide between exploration and exploitation based on the value of ϕ. For example, when $\phi = 0.1$, agents have a 90% chance to change one component of their combination and a 10% chance to change between 2 and 10 components. This probabilistic approach allows agents to balance the advantages of refining known successful strategies (exploitation) with the potential benefits of discovering new strategies (exploration).

After making changes, agents observe the new combinations and their corresponding payoffs. If the new payoff is higher than the previous one, agents adopt the new combination. This process ensures that agents continuously refine their strategies. After each round, the combinations and payoffs are updated accordingly, and each round's mean and maximum payoffs are recorded.

Stoping criteria The simulation runs for a fixed number of trials, set at 20 rounds ($t_max = 20$). This duration was chosen to align with the experimental design and to provide a performance benchmark for participant behavior in the final round. I conducted 100 replications for each of the 8 parameter combinations to ensure statistical reliability. These combinations result from varying the exploration probability ($\phi \in \{0, 0.1, 0.5, 1\}$) and the level of team connectedness ($\gamma \in \{0, 1\}$).

3.3.2 Assumptions

The parameters and setup chosen for the simulation are designed to mirror the experimental conditions closely. These constraints lead to several assumptions and related limitations. First, I operationalize exploitation as changing one random component of the combination, a local search or random walk strategy. Although I do not allow for changing zero components, this occurs in practice. If the new combination produces a lower payoff than the previous one, agents will retain their old combination. Exploration, on the other hand, is defined as changing 2 or more components. The distinction between exploitation and exploration in this simulation is arbitrary.

Second, the selection of which variables to change and the number of changes made during exploration are determined randomly. This randomness introduces variability and reflects the unpredictability of the experiment and some real-world decision-making processes. However, we can imagine other settings where people do not randomly decide which components to change for various reasons. These reasons could include previous knowledge, biases, or heuristics. This information could orient the search and thus change strategies and payoff distribution. Our simulation does not reflect these scenarios. Instead, it focuses on situations where such insights are unavailable.

Third, the simulation assumes that social learning and individual learning occur in a two-step process, with social learning always preceding individual learning (Lazer & Friedman 2007).

This sequential approach simplifies the decision-making process but may not fully capture the complexity of real-world behavior, where individuals often integrate social cues and personal insights simultaneously. In reality, people may combine both types of learning and adjust their strategies accordingly.

Despite these constraints, the Monte Carlo simulation benchmarks participant performance by modeling search behaviors under different exploration and social learning conditions. These benchmarks provide a valuable basis for comparing experimental results and understanding the impact of turnover and team connectedness on strategy and performance.

3.3.3 Simulation results

Descriptive results The Monte Carlo simulation results reveal three key insights. First, connected teams with access to social learning consistently outperform disconnected teams that rely solely on individual learning. Second, incorporating some exploration into the decision-making process is beneficial for performance. The results suggest that any level of exploration is better than none. Teams with $\phi > 0$ perform systematically better than those with $\phi = 0$. A moderate level ($\phi = 0.5$) proved the most effective among the different levels of exploration probability tested. Third, exploration can compensate partially for the absence of social learning. Isolated individuals exploring moderately ($\phi = 0.5, \gamma = 0$) achieve a mean payoff close to the payoff of connected teams doing random walks ($\phi = 0, \gamma = 1$).

Figure 3.1 illustrates the mean payoff over 20 decision-making rounds for all combinations of exploration probability (ϕ) and team connectedness (γ). The lines represent different parameter combinations based on the values of ϕ and γ. Connected teams ($\gamma = 1$) perform better over the rounds than disconnected teams ($\gamma = 0$). However, some exceptions appear as teams approach round 20. For connected teams only exploiting ($\phi = 0$), the payoff does not increase as much as it does for other levels of exploration. Similarly, disconnected teams with ($\phi = 0.5$) show relatively higher payoffs than others. Table 3.2 provides the mean payoff across parameters' combinations in round 20. The mean payoff for connected teams only exploiting ($\phi = 0, \gamma = 1$) is 685.451, close to 681.603, and the mean payoff for isolated individuals moderately exploring ($\phi = 0.5, \gamma = 1$). These results suggest that a random walk, doing only local search in teams, can produce results similar to individual exploring in some rounds. In a sense, exploration could compensate for the lack of social learning.

Figure 3.2 shows the mean payoff for different exploration probabilities (ϕ) for connected

Figure 3.1: Mean agent's payoff over rounds by team connectedness

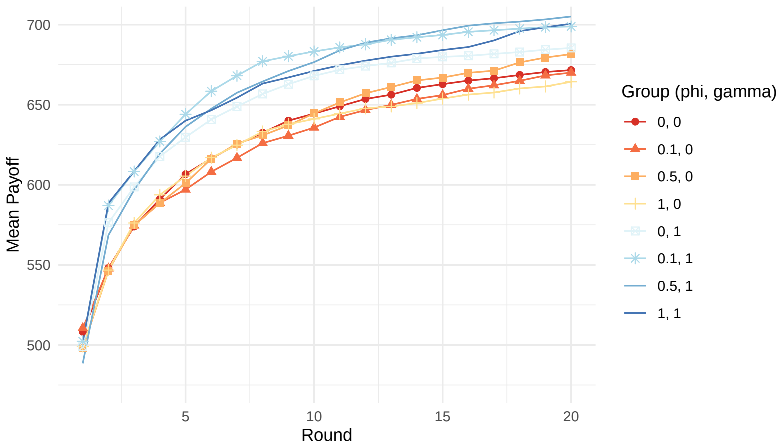

Notes: This graph illustrates the mean payoff over 20 decision-making rounds for all combinations of exploration probability (ϕ) and team connectedness (γ). The lines represent different parameters' combinations based on the values of ϕ and γ. Connected teams ($\gamma = 1$, in blue colors) have access to social learning and can learn from the combinations and payoffs of team members. Disconnected teams ($\gamma = 0$, in yellow-red colors) cannot access social learning and rely only on individual learning. (ϕ) is the probability of exploring and changing two or more components in the combination. Payoffs are derived from the NK landscapes. The results demonstrate the impact of exploration probability and team connectedness on the mean payoff, with connected teams generally achieving higher payoffs over time than non-connected teams.

Table 3.2: Mean payoff across parameters combination and comparison with human participants

		Computational agents (ϕ = probability of random search)			
Connected	Human participants	$\phi = 0$	$\phi = 0.1$	$\phi = 0.5$	$\phi = 1$
$\gamma = 0$	-	671.623	670.04	681.60	664.31
$\gamma = 1$	676.63	685.45	698.87	705.08	700.57

Note: Payoffs are averaged over 1,200 computational agents. Connected teams ($\gamma = 1$) have access to social learning and combinations and payoff of team members. Disconnected teams ($\gamma = 0$) cannot access social learning and rely only on individual learning. Human participants mean payoff is provided for comparison. Human participants did not experiment with the asocial condition ($gamma = 0$); only the payoff for the social learning condition is provided ($gamma = 1$).

Figure 3.2: Mean agent's payoff in round 20 and connected teams, by exploration probability

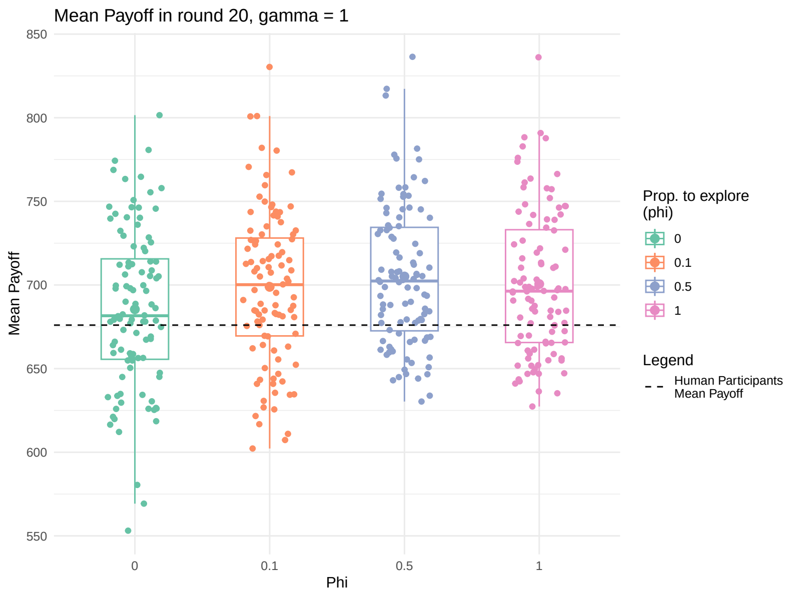

Notes: This boxplot illustrates the mean payoff in the 20th round for different exploration probabilities (ϕ) when teams are connected ($\gamma = 1$). The boxes display the distribution of agent payoffs across the four levels of ($\phi \in \{0, 0.1, 0.5, 1\}$). ($\phi$) is the probability of exploring and changing two or more components in the combination. The black dotted line indicates the mean payoff in the same round, the 20th, for human participants in the experiment. Payoffs are derived from the NK landscapes. The results suggest some levels of exploration are beneficial for performance. Agents with $\phi > 0$ perform on average better than agents only exploiting $\phi = 0$.

teams in round 20. The boxplot illustrates the distribution of payoffs across four levels of ϕ $(0, 0.1, 0.5, 1)$. In line with figure 3.1 and table 3.2, plots suggest that any level of exploration is beneficial compared to no exploration ($\phi = 0$) in connected teams. Teams with some level of exploration ($\phi > 0$) achieve higher mean payoffs. Among these, $\phi = 0.5$ which produces the highest mean payoff. A moderate amount of exploration could be the best balance between exploiting known strategies and exploring new ones in that specific setting.

These results highlight two key findings. First, connected teams consistently achieve higher payoffs compared to disconnected teams. Second, some level of exploration is beneficial for performance, with a moderate level ($\phi = 0.5$) performing best. Third, exploration can compensate partially for the absence of social learning.

The higher performance of connected teams might appear straightforward, as being in a connected team provides more information through social learning. The additional information could turn into enhanced performance. However, previous studies suggest that being in connected teams reduces the propensity to explore, which could lead to decreased performance. Fully connected teams would then perform worse than moderately connected teams. Our results support this idea. When teams rely only on exploitation, their performance is close to those of individuals who explore independently. If being connected reduces the propensity to explore, then the performance of connected teams would decrease. Individuals who explore independently can achieve similar results to those in connected teams performing local searches ($\phi = 0$). The ideal scenario would involve both connected teams and some exploration. Fostering an environment encouraging moderate exploration can lead to better outcomes, even within connected teams.

Performance benchmarks I now focus on establishing performance benchmarks for the experiment. These benchmarks serve as different reference points for comparing experimental outcomes. The Monte Carlo simulations produce three primary benchmarks. First, the low benchmark represents a random walk, where agents only search locally and are not connected ($\phi = 0, \gamma = 0$). This benchmark provides a baseline for performance and represents the simplest scenario. Agents rely solely on individual learning and the exploitation of local solutions. This benchmark highlights the minimum expected performance and serves as a worst-case scenario.

Second, the medium benchmark involves random walks within connected teams, where agents only search locally but can learn from each other ($\phi = 0, \gamma = 1$). This benchmark serves as an intermediate point between the low and high benchmarks. Similarly to the low benchmark, agents rely only on exploitation. However, contrary to the low benchmark, agents now have access to the combinations of others and can learn socially. As human participants are in connected teams, this benchmark represents the minimum expected performance, assuming participants use available social information.

Third, the high benchmark represents the optimal scenario, where connected teams balance local search and exploration ($\phi = 0.5, \gamma = 1$). In this scenario, agents not only exploit local solutions but also explore new possibilities. This benchmark highlights the maximum expected performance. By comparing human participants' performance to this high benchmark, I assess their ability to balance exploitation and exploration within a connected team setting.

Figure 3.3 presents the mean payoff for the three benchmark scenarios in round 20. The boxplot illustrates the distribution of payoffs the low ($\gamma = 0, \phi = 0$), medium ($\gamma = 1, \phi = 0$), and high ($\gamma = 1, \phi = 0.5$) benchmarks. The results show a clear performance progression from the low to the high benchmarks. Disconnected teams with no exploration ($\gamma = 0, \phi = 0$) achieve the lowest payoffs. Connected teams with no exploration medium ($\gamma = 1, \phi = 0$) perform better than those not connected. Connected teams with moderate exploration achieve the highest payoffs. In this scenario, teams exploit both social information and the value of balancing exploitation

Figure 3.3: Mean agent's payoff in round 20 for different benchmark scenarios

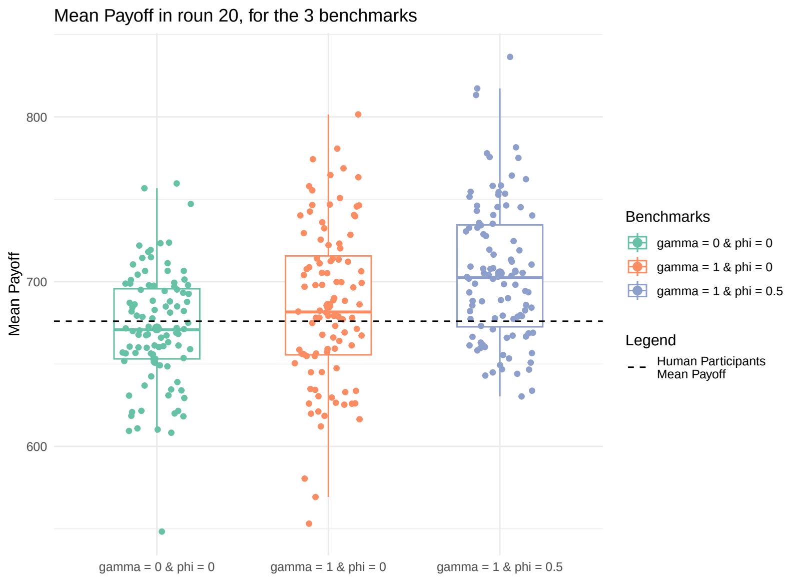

Notes: This boxplot illustrates the mean payoff in the 20th round for the three benchmark scenarios: (1) the low benchmark with disconnected teams and no exploration ($\gamma = 0$, $\phi = 0$), (2) the medium benchmark, with connected teams but no exploration ($\gamma = 1$, $\phi = 0$), and (3) the high benchmark with connected teams and moderate exploration ($\gamma = 1$, $\phi = 0.5$). The boxes display the distribution of agent payoffs in round 20 for each scenario. The black dotted line indicates the mean payoff in the same round for human participants in the experiment. Payoffs are derived from the NK landscapes.

with exploration. The mean payoff of human participants is provided for information purposes. Together, these three benchmarks provide a comprehensive range of expected outcomes and allow me to assess the performance of human participants in the experiment accurately. In the experimental results section, I discuss in detail the comparison between these computational agents' benchmarks and the human participant's performance.

3.4 Experimental methods

3.4.1 Experiment overview

This experiment investigates whether team disruptions can foster exploration and improve performance in complex tasks. The study involved 258 participants from the University of Lausanne, working in teams of 3. The teams engage in a combinatorial task involving 10 binary decisions based on an NK landscape framework. The experiment tests the effects of two types of disruptions: changes in team composition (turnover) and shifts in reference points. Turnover is implemented by replacing the top-performing member of a team with a newcomer. Reference points were shifted by informing participants of better possible solutions. The goal is to determine if these disruptions can counteract the tendency of teams to quickly transition from exploration to exploitation and potentially improve overall performance.

3.4.2 Framed NK landscape task

From a participant's point of view, the task is to design a painting for aliens. To do so, each participant has to select or deselect 10 different geometric shapes that compose the painting. As participants engage in the task, they do not know what combinations would yield higher payoffs. To create a painting, participants select or deselect 10 different geometric shapes. The combination of these shapes creates the artistic creation. After making their selections, they submit their design and receive immediate feedback through a payoff, representing the aliens' willingness to pay for that particular combination. The payoff related to each combination is derived from the NK algorithm (Csaszar 2018). The framing matches the alien task framing from Billinger et al. (2014), which aims to ensure that participants do not rely on their previous experiences to guide their decisions. The payoffs for each combination of shapes were determined by the NK algorithm and were not based on real-life experience. Payoffs are given in alien currency to preserve the belief that participants cannot rely on previous knowledge.

The task is based on an NK landscape, a model that simulates complex problem-solving. N represents the number of binary choice variables, and K indicates the interdependencies among these variables. In the experiment, N equals 10, matching previous experiments. Each participant has to make 10 binary decisions in each round. $K = 5$ indicates a moderate level of complexity. To generate the NK landscape, I first create a matrix of random numbers between 0 and 1. Second, I combine this first matrix with an interaction matrix to capture the interdependencies (5) between the variables. Finally, combining these two matrices leads to the payoff matrix, which captures the randomly generated fitness contributions for each possible combination of the 10 variables, considering the interaction between the 5 variables and reflecting the landscape's complexity. The payoff structure is directly derived from this NK payoff matrix, where each combination of the 10 binary decision variables resulted in a specific payoff.

From the NK task, I obtain two outcome measures: payoffs and search distance. Payoffs are directly derived from the payoff landscape. Search distance is the number of attributes that are different in the focal combination than in the last round. Because the task includes 10 components, the search distance can range from 0 to 10. 0 means the combination is the same as in the last round. 10 means that all attributes are different. All the attributes selected in the focal

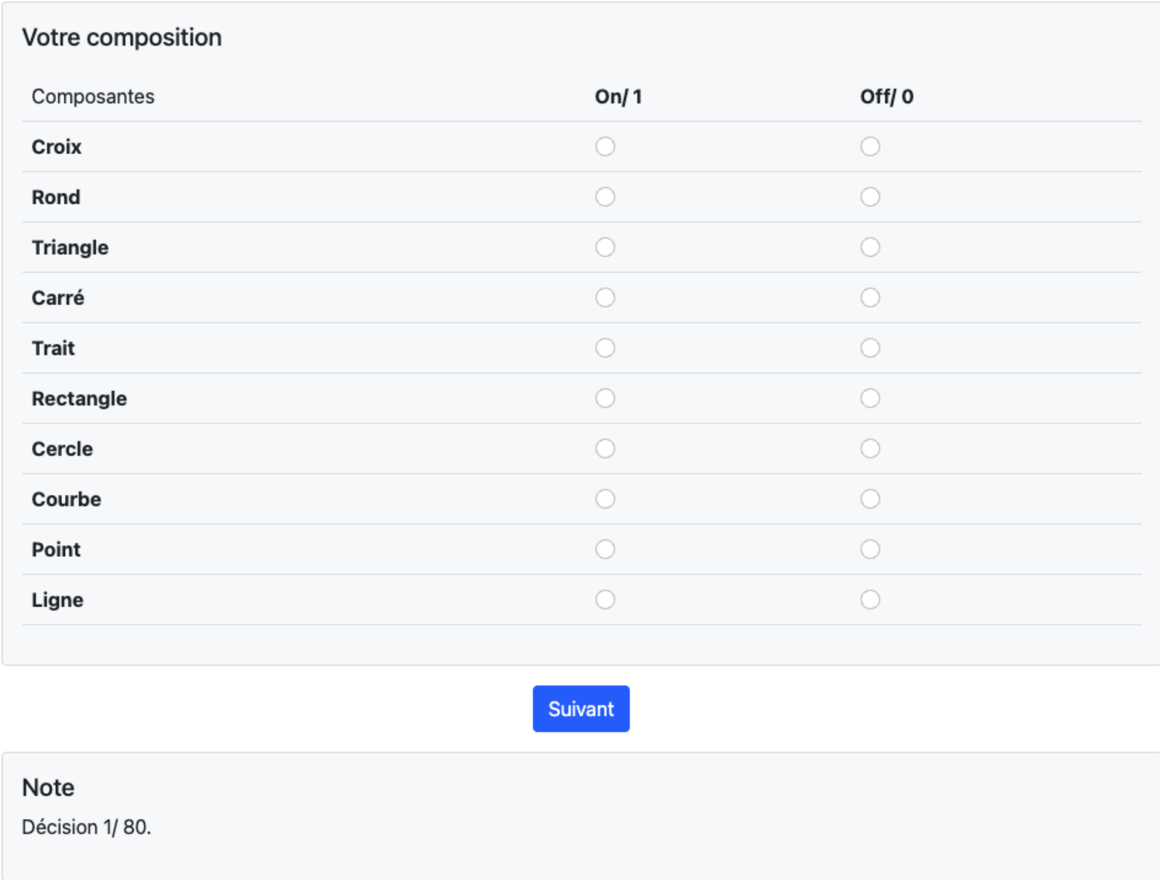

Figure 3.4: Decision screen in $t = 1$. The participant has to decide for each of the 10 components if she would like to add it (On/1) or not (Off/0).

combination were not selected in the previous combination, and all the deselected attributes were selected. Search distance is used to pinpoint exploitation from exploration behaviors. A small search distance is related to exploitation, while a high search distance is related to exploration. Both payoffs and search distance are analyzed to assess the impact of the treatments.

3.4.3 Procedure and design

Participants are randomly divided into teams of 3 players, and the teams remain the same throughout the experiment. Participants in the same team share the same payoff landscapes throughout the game. In other words, they explore the same environment. In the first round, no combination is displayed to avoid any anchoring effect on this initial combination. Participants start with a blank page where attributes are neither selected nor deselected. They select or deselect the 10 different geometric shapes. Once they submit their combination, they receive immediate private feedback. They observe their combination and the payoff related to that combination given by the NK landscape. Additionally, information is shared within teams. Participants observe team members' combinations and related payoffs. Because team members share the same landscape, social information is valuable. Imagine that a participant observes a team member's combination and related payoff in the last round; she can be sure to obtain the same payoff if she tries the same combination. The next round starts once all team members have submitted their combination and received feedback. In the following rounds, each participant starts with the combination they chose in the previous round.

Midway through the sequence, after 10 rounds, treatments are introduced.

Participants go through 20 decision-making rounds. They decide whether to keep or change each of the 10 binary variables in each round. 20 rounds form a sequence. In a sequence, participants can try sequentially 20 different combinations. Given that participants are divided into teams of 3, the maximum number of different combinations a team can explore during a sequence is 60, which is a small fraction of the total number of possible combinations, 1024 (2_{10}, each of the 10 attributes having 2 possible states). After 20 rounds, the sequence stops. Each participant goes through 4 sequences in a session. Before each new sequence, initial conditions are reset, and a new NK landscape is generated. Participants are informed that the sequence is finished, and they start a new one. They are told that knowledge about successful combinations acquired during the previous sequence is no longer helpful. To make that information salient to the participants, they are informed that they are now facing a new alien population unrelated to the previous population. To prevent participants from inferring their current level of performance based on the payoffs they received in the previous sequence, the payoff multiplier changes after each sequence, as well as the name of the alien currency. For example, in sequence 1, payoffs range around 34 galax. In sequence 2, around 865 fluxi.

At the very end of the session, participants see their total payoff translated into CHF.

3.4.4 Two-by-three design based on turnover and reference point

Table 3.3: Treatment Conditions

	No turnover	Turnover	Turnover + "Best performer is leaving"
No information about better solutions	*Control*	*Turnover*	*Turnover + Know*
"Better Solutions Exist"	*Reference*	*Reference + Turnover*	R*eference + Turnover + Know*

The experiment is a two-by-three design, and both dimensions are between subjects within sessions. Specifically, each session will involve the six treatments. Treatments are based on disruptions and vary regarding changes in reference points, turnover, and turnover information. All treatments appear midway through the sequence, after round 10 and before round 11 in each sequence. All participants get to observe a screen announcing they are midway through the sequence. Treated groups receive additional information.

First, participants can receive the information that better solutions exist and that they can do better. Treated groups are randomly selected at the beginning of the experiment, but the program only shows the message to those groups that have not yet reached the global optimum. Because the number of possible solutions is significantly higher than the number of trials, reaching the global optimum is unlikely. Indeed, in our experiment, no group reached the global optimum. Telling participants they can improve aims to shift their reference point and place them in the loss domain, encouraging more risk-taking and exploration.

Second, groups can experience turnover. Specifically, in the treated groups, the top-performer leaves the group while a newcomer enters the group. Newcomers for each sequence and each group are randomly selected at the beginning of the experiment. During the first 10 rounds of a sequence, newcomers wait and do not observe other team members' combinations and related payoffs. The participants waiting have the same expected payoff as the participants playing. During waiting rounds, participants who wait to join a group are compensated. The goal is to equalize the experimental conditions for all participants, whether actively engaged in decision-making or waiting. By doing so, I prevent any disadvantage in expected earnings due to the timing of joining the teams. When entering the team, newcomers observe the combinations and associated payoffs of the two remaining team members in the last round and make their first decision. Mirroring the initial conditions of the other participants, newcomers start with a blank combination with all attributes neither selected nor deselected.

In the turnover conditions, turnover always involves the highest-performing member. However, some teams know this detail, while others do not. Specifically, participants in the aware condition see on their screen that the top-performer is leaving the team. In the unaware condition, the message only indicates that a participant is leaving the team. The goal is to study the influence of turnover on search behaviors. Further, the two different treatments allow us to disentangle the mechanical effect of turnover from psychological effects related to the departure of the top-performer.

The experiment was pre-registered (Faessler 2023). I collected data by running 9 sessions from November 23 to 30, 2023, in the behavioral laboratory at HEC Lausanne on a computer network using oTree (Chen et al. 2016). Participants were recruited using the online recruiting software ORSEE (Greiner 2004). A total of 258 participants were recruited (males = 52.96%). Participants were mainly students from the University of Lausanne and the Swiss Federal Institute of Technology in Lausanne (EPFL). Most of the students were born in Switzerland (97 participants), followed by France (84 participants) and Tunisia (19 participants). The sample includes a total of 31 different nationalities. The sessions lasted 1 hour and a half. Before the experiment, I collected explicit consent from the participants. Participants were free to leave the experiment at any point, although none took that decision. Participants had to pass a comprehension check to start the experiment. I collected sociodemographic data and risk preferences at the beginning of the experiment. Participants received CHF 20.- for their participation plus a bonus varying between CHF 10.- and CHF 50.- depending on their average performance.

3.5 Results

The analysis is divided into four main parts. First, descriptive statistics present the evolution of search distance and performance over time across treatments. In the second part, I contrast these descriptive observations with treatment comparison. I compare in a regression analysis search distance and performance across treatments in the 3 rounds following treatments. In addition, I conducted a diff-and-diff analysis to compare each treatment search distance and performance before and after the treatment. The third part contrasts participants' performance with simulated benchmarks. Finally, I focus on the newcomers and explore the differences between newcomers and oldtimers.

3.5.1 Descriptive Results

I provide first descriptive data about the evolution of (1) search distance and (2) payoff. Figure 3.5 shows the treatment groups' average search distance over time. At first, participants explore widely and show a larger average search distance, deviating from a pure local search strategy (search distance = 1). The search distance quickly decreases, mirroring patterns observed in previous studies about search behaviors when alone and in teams (Billinger et al. 2014,

Figure 3.5: Search behavior over rounds across treatments

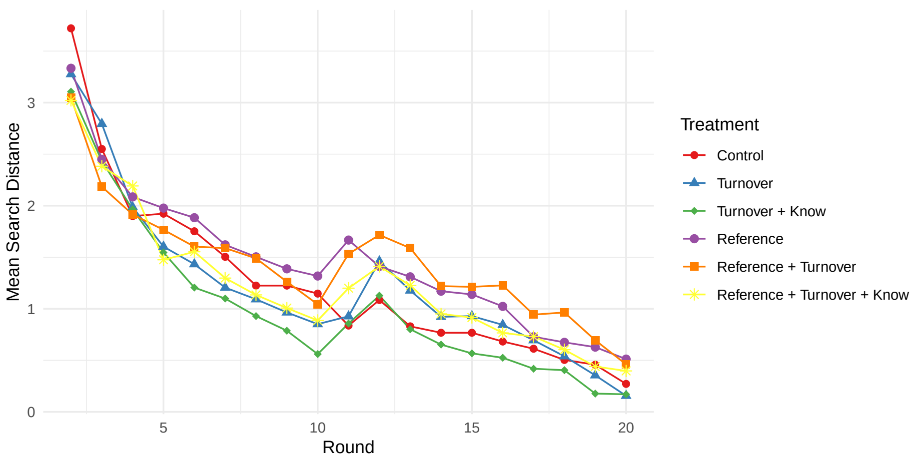

Notes: The search distance is measured as the number of attributes changed in a round relative to the combination in the previous round for a participant. Search distance ranges from 0 to 10. Treatments happen after round 10 and before round 11.

Giannoccaro et al. 2020, Billinger et al. 2021).

The treatments are introduced between rounds 10 and 11. Post-treatment, a brief increase in search distance occurs across all treatments, including the control group. The control group does not experiment with turnover or receive information about their performance. However, the participants are still exposed to a screen, indicating they are midway through a sequence. The sudden increase in all treatments suggests that the treatments work and motivate the participants to increase search distance. The participants temporarily explore more. This sudden increase in search distance deviates from the steady decrease seen in other studies without midway treatments (Billinger et al. 2014, Giannoccaro et al. 2020, Billinger et al. 2021). These observations suggest that inducing a brief change in search behaviors when the team exploits is surprisingly easy.

After a few rounds, however, the trend reverts to its initial trajectory, and the means search distance continues to decline. In the last rounds, the average search distance is lower than 1 across all treatments. The participants either submit the same attribute combination or change only one of the attributes. The decrease in search distance indicates a shift towards exploitation, as participants focus more on refining and capitalizing on their current solutions over rounds. Even though generating a brief moment of exploration is easy, maintaining this change in time seems challenging.

Turning to the evolution of payoff, Figure 3.6 shows the evolution of the mean payoff for

Figure 3.6: Performance over rounds across treatments

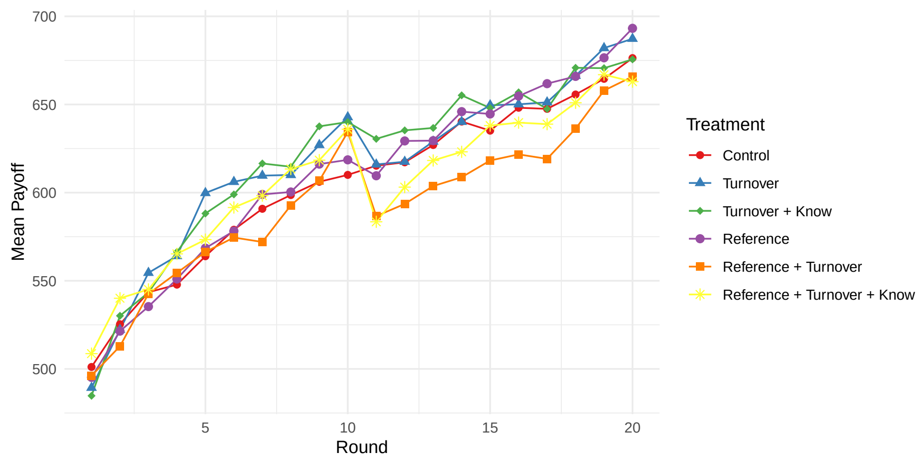

Notes: The payoff is derived from NK algorithms. Maximum payoff ranges from 725 to 874 and depends on each NK landscape.

different treatment groups over rounds. In the first round, all groups start with a payoff that reflects the expected outcome of making selections randomly, given that participants begin without prior knowledge and face a blank state as their initial condition. As the experiment advances, the mean payoff increases. Participants discover which combinations result in better payoffs and start refining them.

Interestingly, in round 11, which directly follows the treatment, all groups but the control group experience a decrease in payoff. This decrease deviates from the findings from previous studies, which did not include a midway treatment (Billinger et al. 2014, Giannoccaro et al. 2020, Billinger et al. 2021). These previous studies all report a steady increase in payoff over rounds. The decrease is moderate for some treatment groups, Turnover + Know and Reference. For the other groups, Turnover, Reference + Turnover, and Reference + Turnover + Know, the decrease is substantial. These groups are also the groups that show the bigger increase in search distance following the treatment. These results suggest that while the increase in search distance led to exploring more distant solutions, these alternative solutions yielded lower payoffs than previous ones. Deviations from local search strategies did not benefit participants. We know, however, that deviations from local search could have produced higher payoffs if only these deviations had been the correct ones.

Following this brief decrease in payoff, the trend reverts to its initial trajectory. Participants are back to exploiting their previous combination. Payoffs continue to increase for all groups. Notably, by the end of the sequence, all groups, including the groups that experience a drop in payoff in round 11, achieve a similar payoff between 650 and 700.

Overall, the analysis of the descriptive statistics points to two observations. First, the treatment works and generates a change in search behaviors. Search distance increases post-treatment for all groups. However, turnover does not seem to induce a bigger change than inducing a change in the reference point or simply telling participants they are midway through a sequence. In all groups, the increase in search distance is slight and brief. Second, this increase in search distance leads to a decrease in payoff. Although better solutions are available for most groups, the increase in search distance is too brief to lead to the discovery of the optimum.

3.5.2 Treatments comparison

Descriptive statistics suggest that all treatments increase the search distance, and some treatments decrease the payoff. I now turn to pre-registered regression analysis and examine how the treatments impact search distance and payoff (Faessler 2023). I distinguish between group and individual responses and immediate post-treatment effects and effects in all subsequent rounds.

Changes in reference point but not turnover influences search distance.

Figure 3.7: Comparison of group average search distance in rounds 11 to 13 across treatments

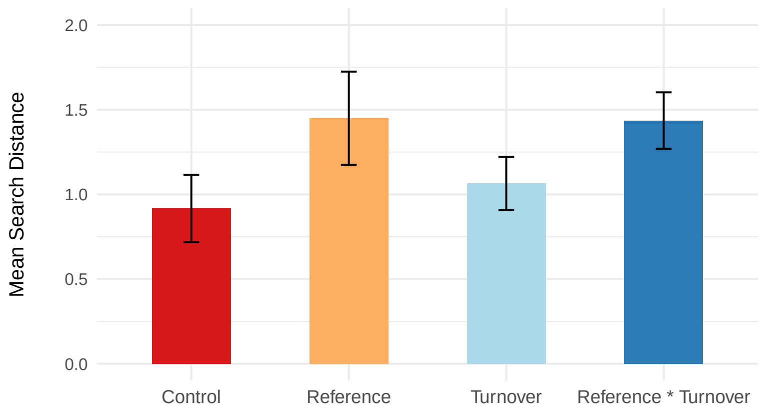

Notes: Search distance is averaged across individuals within each group for rounds 11 to 13. Search distance is measured as the number of attributes changed in a round relative to the combination in the previous round for a participant and ranges from 0 to 10. Error bars indicate 95% confidence intervals.

In this first part, I compare the average group search distance in the three rounds following the treatment across turnover and reference treatments. For this analysis, I do not differentiate between turnover treatments. I aggregate the Turnover and the Turnover + Know treatments, informing participants that the member who just left was the top-performer. I analyze the difference between the two treatments and the potential psychological effects in a dedicated analysis. I compare the search distance across treatments. Figure 3.7 suggests that the change in the reference point, but not turnover, significantly increases the search distance. The average search distance is higher in Reference and Reference * Turnover treatments compared to the control. On the other hand, the average search distance does not significantly differ in the turnover treatment compared to the control.

Table 3.4: Average group search distance in rounds 11 to 13: regression analysis

	Dependent variable:	
	Group average search distance	
	(1)	(2)
Reference	0.540*	0.541*
	(0.219)	(0.219)
Turnover	0.145	0.145
	(0.187)	(0.187)
Reference + Turnover	0.512**	0.512**
	(0.187)	(0.187)
Sequence		−0.192***
		(0.036)
Constant	0.919***	1.392***
	(0.155)	(0.178)
Observations	274	274
Log Likelihood	−319.066	−307.926

Notes: The four treatments are coded as a single variable: control (committed category), Reference, Turnover, Reference + Turnover. Model (2) includes a control variable for the sequence. Participants go through four sequences in a session, which can produce a learning or fatigue effect. Robust standard errors are clustered at the group level and are in parenthesis. *p<0.05; **p<0.01; ***p<0.001.

The companion regression confirms these results. Table 3.4 shows positive and significant Reference and Reference * Turnover coefficients. However, the Turnover coefficient is not significant. Changes in the reference point increase search distance, but Turnover alone has no significant effect. The sequence coefficient is negative and highly significant across models. Search distance decreases with sequences. As the session progresses, participants are less likely to explore. This effect can be due to a learning effect. Participants understand they have little chance to find a better solution and converge quicker to exploitation. An alternative explanation is a fatigue or “willingness to go home” effect. Participants understand that the quicker they play, the quicker they will be done. Although the sequence effect is significant and its origin unknown, I have no reason to believe that the effect was different across treatments.

Knowing that the top-performer is leaving makes no difference. So far, I have not differentiated turnover treatments, one in which team members are informed about the departure of the top-performer and another in which they are not. In both conditions, the top-performer is leaving the team. The difference lies in whether the teams are aware of this condition. The objective was to understand whether the knowledge of the top-performer's departure psychologically impacts the team's search behaviors. So far, field evidence suggests that top-performer turnover increases exploration (Tzabbar & Kehoe 2014). However, we do not know if the effect is linked to the characteristics of the person leaving the team or to turnover itself. In the following regression, I focus on turnover treatments and distinguish between Know treatments, where participants are made aware that the top-performer is leaving the team. Figure 3.8 shows that the difference in average search distance across the two groups is not significant. Regression analysis displayed in table 3.5 confirms these results. In this controlled lab setting, being aware of the identity of the person leaving the team makes no difference in search behaviors.

Figure 3.8: Comparison of group average search distance in rounds 11 to 13 across turnover treatments

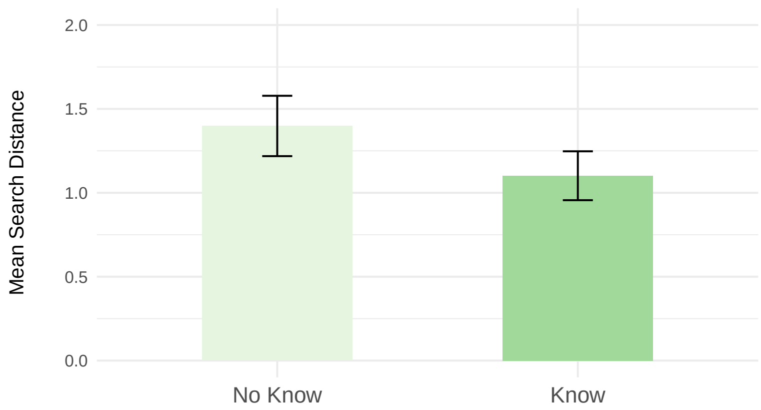

Notes: Search distance is averaged within groups experimenting turnover for rounds 11 to 13. Data include only groups encountering turnover. Search distance is measured as the number of attributes changed in a round relative to the combination in the previous round for a participant and ranges from 0 to 10. Error bars indicate 95% confidence intervals.

Table 3.5: Average group search distance in rounds 11 to 13 in turnover treatments: regression analysis

	Dependent variable:			
	Group search distance			
	(1)	(2)	(3)	(4)
Know	-0.295^{*}	-0.295^{*}	-0.253	-0.253
	(0.134)	(0.135)	(0.191)	(0.193)
Reference	0.368^{**}	0.368^{**}	0.410^{*}	0.410^{*}
	(0.134)	(0.135)	(0.191)	(0.193)
Reference * Know			-0.084	-0.084
			(0.270)	(0.272)
Sequence		-0.176^{***}		-0.176^{***}
		(0.046)		(0.046)
Constant	1.212^{***}	1.645^{***}	1.191^{***}	1.624^{***}
	(0.116)	(0.162)	(0.135)	(0.177)
Observations	188	188	188	188
Log Likelihood	-220.839	-215.976	-221.184	-216.315

Notes: Analyses include only groups encountering turnover. Treatments are coded as 2 binary variables: Reference and Know. Model (2) includes a control variable for the sequence. Participants go through four sequences in a session, which can produce a learning or fatigue effect. Model (3) includes interaction terms, and model (4) includes controls and interaction terms. Robust standard errors are clustered at the group level and are in parenthesis. *p<0.05; **p<0.01; ***p<0.001.

Changes in reference point and turnover do not influence performance.

Figure 3.9: Comparison of group average performance in rounds 12 to 20 across treatments

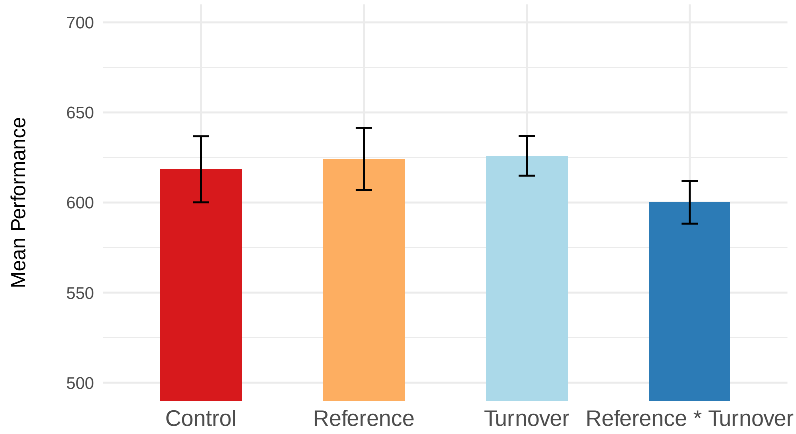

Notes: Performance is calculated as the average payoff across individuals within each group for rounds 11 to 13. Payoffs are derived from NK algorithms. Maximum payoff ranges from 725 to 874 and depends on each NK landscape. Error bars indicate 95% confidence intervals.

Table 3.6: Average group payoff in rounds 11 to 13: regression analysis

	Dependent variable:	
	Group average payoff	
	(1)	(2)
Reference	5.796	5.792
	(13.752)	(13.739)
Turnover	7.584	7.568
	(11.741)	(11.729)
Reference + Turnover	−17.987	−17.993
	(11.741)	(11.729)
Sequence		7.824**
		(2.900)
Constant	618.274***	598.998***
	(9.724)	(12.058)
Observations	274	274
Log Likelihood	−1,478.971	−1,473.395

Notes: The four treatments are coded as a single variable: control (committed category), Reference, Turnover, Reference + Turnover. Model (2) includes a control variable for the sequence. Participants go through four sequences in a session, which can produce a learning or fatigue effect. Robust standard errors are clustered at the group level and are in parenthesis. *p<0.05; **p<0.01; ***p<0.001.

In this second part, I switch to performance and compare the average payoff in the 3 rounds following the treatment across treatments. Figure 3.6 suggests that none of the treatments significantly influenced payoff. Even though the descriptive graphs suggest a drop in payoff in some treatments, the difference appears to be not significant. The average payoff does not differ across treatments. The companion regression 3.6 shows similar results. Both Reference and Turnover coefficients are not significant. The sequence coefficient is positive and significant. Payoffs increase with sequences. As sequences progress, search distance decreases, and payoffs increase.

Altogether, the analysis of the groups' average search distance right after the treatment suggests that Reference but not Turnover treatments have an effect. Search distance increases following the information that better solutions exist but does not significantly differ following a change in the team composition. The increase in search distance after the reference treatment does not translate into a significant difference in performance. None of the treatments significantly affect the groups' average payoffs. Interestingly, progressing through the sequences has a significant effect. As the participants progress in the experiment, their search distance decreases, and their performance increases.

On the one hand, the descriptive results suggest that search distance responds to all treatments. On the other hand, the regression analysis results show that only Reference treatments lead to a different increase in search distance than the control. The reason behind this inconsistency might be that even the control condition leads to an increase in search distance. Even though search distance increases following team turnover, the distance does not increase significantly compared to the control group. Like the other groups, the control group encounters a screen midway that provides information. In the control case, the information is that participants are midway through the sequence. It is possible that even this minor intervention, making participants aware that they are midway through a sequence, affects search distance. In the following analysis, I explore if there is a treatment effect for each treatment, comparing the trends before and after the treatment.

Comparing before and after for each treatment

In the following regressions, I compare the state before and after the treatment for each treatment condition. Search distance and payoff are measured at the individual level. Regressions include rounds 6 to 15, 5 rounds before, and 5 rounds after the treatment. Again, I focus on Reference and Turnover treatment and do not distinguish between the two different conditions within the

Turnover treatment. The regressions include only individuals in the group before and after the treatment. Newcomers and members who left after the treatment are excluded from the analysis as they do not have pre-treatment or post-treatment observations.

Table 3.7: Individual search distance in rounds 6 to 15: regression analysis

	Dependent variable:	
	Individual search distance	
	(1)	(2)
Post-treatment	-0.508^{***}	0.035
	(0.065)	(0.095)
Reference	0.106	0.139
	(0.189)	(0.189)
Turnover	-0.416^{*}	-0.403^{*}
	(0.166)	(0.165)
Reference + Turnover	−0.117	−0.096
	(0.166)	(0.165)
Post * Reference	0.296^{**}	0.301^{***}
	(0.092)	(0.090)
Post * Turnover	0.245^{**}	0.246^{**}
	(0.082)	(0.080)
Post * Reference + Turnover	0.349^{***}	0.353^{***}
	(0.082)	(0.080)
Round		−0.475
		(0.327)
Round^2		0.029
		(0.033)
Round^3		−0.001
		(0.001)
Sequence		-0.152^{***}
		(0.013)
Constant	1.414^{***}	4.041^{***}
	(0.138)	(1.026)
Observations	7,280	7,280
Log Likelihood	−11,711.940	−11,573.000

Notes: Search distance is measured as the number of decisions that differ from the decisions in the previous combination for each individual. Post-treatment is a dummy variable that indicates whether the observation is before (rounds 6 to 10) or after the treatment (rounds 11 to 15). The four treatments are coded as a single variable: control (committed category), Reference, Turnover, Reference + Turnover. Time effect is included via the *round* variable, up to the power 3. Model (2) includes a sequence and time effect control variable. Robust standard errors are clustered at the group and individual levels in parenthesis. *p<0.05; **p<0.01; ***p<0.001.

Search Distance Table 3.7 shows the results for search distance. All treatment coefficients, when interacting with the Post-treatment variable, are positive and significant. Participants are likely to break off from exploiting and briefly increase search distance following information about the existence of better solutions or changes in team composition. Search distance significantly increases following Reference and Turnover treatments. The sequence effect remains. As sequences progress, search distance decreases in all treatment conditions.

Surprisingly, the Turnover coefficient not interacting with the Post-treatment dummy is negative and significant. This result suggests that the Turnover participants have a lower search distance even before the treatment happens. Because participants are randomly assigned to treatments, I did not expect any significant difference across groups before treatment. Ex-post analyses of treatment groups' characteristics show no significant difference in the distribution of risk preferences, gender, age, or origin. Even though some individuals perform well in this treatment and can be considered outliers, excluding outliers from the analysis does not change the results.

Payoff Switching to the performance analysis, Table 3.8 shows the results for individual payoffs. The Post-treatment coefficient is negative and significant. Participant's payoff significantly decreases following the treatment. However, this decrease is not significantly different across treatments, but the Reference * Turnover interactions. The coefficient related to the interaction of the Post-treatment variable and Reference * Turnover is negative and significant. Participants in the Reference * Turnover group experience a significant drop in performance after the treatment.

Like in the search distance analysis, the Turnover alone coefficient, not interacting with the Post-treatment variable, is significant, which indicates a significant difference across treatment groups even before the treatment. This time, the coefficient is positive. Participants in the Turnover treatment are more likely to exploit and have a smaller search distance before and after the treatment. Conversely, they perform better than the other groups. Because there is no significant difference across groups and potential outliers do not influence the results, the reason behind this finding remains unclear.

Altogether, these regression analysis results confirm the previous findings from the descriptive analysis. Treatments, whether altering the reference point by informing participants that better solutions exist or inducing group composition changes, increase the search distance. Participants increase their search distance and explore more following the treatment. In that sense, the goal

Table 3.8: Individual payoff in rounds 6 to 15: regression analysis

	Dependent variable:	
	Individual payoff	
	(1)	(2)
Post-treatment	29.638***	−18.809**
	(4.668)	(6.870)
Reference	−0.253	−2.617
	(12.286)	(12.300)
Turnover	24.150*	23.414*
	(10.680)	(10.704)
Reference + Turnover	1.715	0.049
	(10.680)	(10.705)
Post * Reference	0.169	−0.164
	(6.616)	(6.515)
Post * Turnover	−4.432	−4.466
	(5.896)	(5.807)
Post * Reference + Turnover	−14.174*	−14.429*
	(5.911)	(5.822)
Round		−12.397
		(23.723)
Round^2		2.542
		(2.412)
Round^3		−0.091
		(0.077)
Sequence		5.768***
		(0.909)
Constant	599.290***	568.756***
	(8.907)	(74.360)
Observations	7,280	7,280
Log Likelihood	−42,776.560	−42,669.100

Notes: Post-treatment is a dummy variable that indicates whether the observation is before (rounds 6 to 10) or after the treatment (rounds 11 to 15). The four treatments are coded as a single variable: control (committed category), Reference, Turnover, Reference + Turnover. Time effect is included via the *round* variable, up to the power 3. Model (2) includes a control variable for the sequence and time effect. Robust standard errors are clustered at the group and individual levels and are in parenthesis. $^{*}p<0.05$; $^{**}p<0.01$; $^{***}p<0.001$.

is reached. The teams do not stick to the exploitation equilibrium and can break off to search for alternative solutions. However, the induced exploration does not lead to higher payoffs. On the contrary, across all treatment groups, performance significantly decreases following treatments. Further, the Turnover group obtains a significantly lower search distance throughout the game and performs significantly higher.

3.5.3 Performance comparison with simulated benchmarks

I now compare human participants' performance in the experiment's final round (round 20) with performance benchmarks generated using the Monte Carlo simulation. I generated three primary benchmarks. First, the low benchmark represents agents searching locally and not connected ($\phi = 0, \gamma = 0$, mean payoff in round 20 = 671.63). This benchmark represents a baseline performance. Second, the medium benchmark involves agents searching locally but connected and having access to social learning ($\phi = 0, \gamma = 01$, mean payoff in round 20 = 685.45). Reaching this benchmark suggests that participants leverage social learning but exploit known solutions. Finally, the high benchmark represents agents balancing local search and exploration in connected teams ($\phi = 0.5, \gamma = 01$, mean payoff in round 20 = 705.08). This benchmark would indicate that human participants effectively balance exploitation and exploration while leveraging social information.

Table 3.9: Mean payoff across treatment groups in round 20

	Control	Turnover	Turnover + Know	Reference	Reference + Turnover	Reference + Turnover + Know
Mean Payoff	676.30	687.28	675.59	693.20	665.82	662.95

Note: Mean payoff for human participants in each treatment group in round 20. For comparisons with computational agents, benchmarks derived from the Monte Carlos simulations are low 671.63, medium 685.45, and high 705.08.

Table 3.9 and Figure 3.10 present the mean payoffs for human participants across the different treatment groups in round 20, compared to the simulated benchmarks. The performance of human participants varies across treatment groups. The mean payoffs of the two groups are below the low benchmark, the Reference + Turnover group (mean payoff = 665.82) and the Reference + Turnover + Know group (mean payoff = 662.95). The mean payoff of two other groups is slightly above the low benchmark and approaching the medium benchmark, the Control group (mean payoff = 676.30) and the Turnover + Know group (mean payoff = 675.59). The

Figure 3.10: Mean payoff in round 20 across treatments and comparison with simulated benchmarks

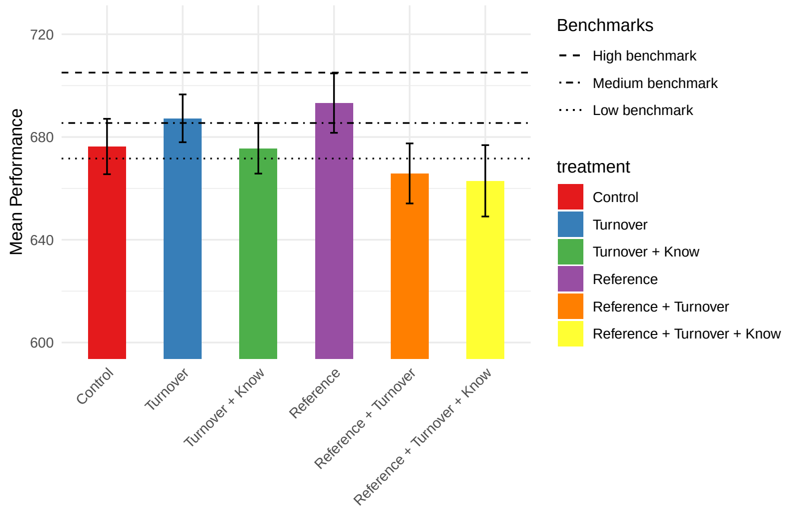

Notes: This graph illustrates the mean payoff in the last round, round 20. Payoffs are derived from the NK landscapes. Maximum payoff ranges from 725 to 874 and depends on each NK landscape. Error bars indicate 95% confidence intervals. Lines indicate simulated benchmarks. The low benchmark (dotted line) represents a random walk in isolation. The medium benchmark (dash-dotted line) represents a random walk with access to social learning. The high benchmark (dashed line) represents a balance between exploration and exploitation with access to social learning.

Turnover group's mean payoff exceeds the medium benchmark (mean payoff = 687.28). Finally, the Reference group's mean payoff is the highest among all groups and approaches the high benchmark (mean payoff = 687.28).

The performance comparison reveals that most treatment groups reach the medium benchmark, representing the performance of agents performing a random walk in connected teams with access to social learning. This finding aligns with the theory suggesting that being in connected teams increases performance compared to isolated individuals but can reduce exploration. The participants in these groups achieved the payoff expected from connected teams that mostly exploit known solutions. However, two treatment groups, Reference + Turnover and Reference + Turnover + Know, did not reach the medium benchmark. This result is somewhat puzzling. Their performance is similar to that of isolated individuals engaging in random walks. Two hypotheses are in order. First, this could suggest that these groups failed to leverage the benefits of being connected. Second, these groups could have used search strategies different from those modeled in the simulation, which appeared less effective.

On the other hand, one treatment group, the Reference group, performed exceptionally well, nearly reaching the high benchmark. This group also exhibited the highest search distance before the treatments were applied. This is pure speculation, but this group might have benefitted from enhanced exploration before the treatment, as the simulations suggest that any level of increased exploration can lead to higher performance. This group's ability to balance exploration and exploitation more effectively than others might have contributed to their superior performance, independent of the treatments.

Overall, the groups achieved medium performance compared to the benchmarks generated with the Monte Carlo simulation. These results indicate that participants performed at the level of connected teams, primarily exploiting known solutions. The simulation results suggest that exploration could have improved participants' payoffs, which was the goal of the treatments. However, none of the treatments led to a sustainable increased exploration, as the teams quickly reverted to exploitation. These results highlight the importance of finding solutions that can maintain higher levels of exploration in connected teams to enhance team performance.

3.5.4 Newcomers versus oldtimers

Search distance Do newcomers and oldtimers exhibit different search behaviors post-treatment? At the core of the study lies the idea that newcomers could bring novel ideas into the teams and lead to more exploration. I test this idea in the following analysis. Figure 3.11a plots the average search distance for oldtimers and newcomers over rounds. The average search distance for newcomers in round 12 is around 2, almost 1 point higher than the average search distance for oldtimers. However, the newcomers' average search distance quickly decreases to match the level of oldtimers. Even though newcomers might exhibit a higher search distance in the first rounds following their entry, they quickly conform to the group and adopt an exploiting strategy. Figure 3.11b plots the average search distance in the three rounds following newcomers' entry. The average search distance for newcomers is significantly higher. The regression results in table 3.10 confirm the higher search distance for newcomers in rounds 12 to 14. The newcomer coefficient is positive and significant. Nonetheless, the difference quickly vanishes. The negative and significant round coefficient confirms the trend observed in graph 3.11a. When they enter the game, newcomers start by exploring more than the oldtimers, but this difference diminishes as the newcomers explore less over rounds. By round 14, newcomers and oldtimers exhibit similar average search distances.

(a) Newcomers and oldtimers search distance evolution

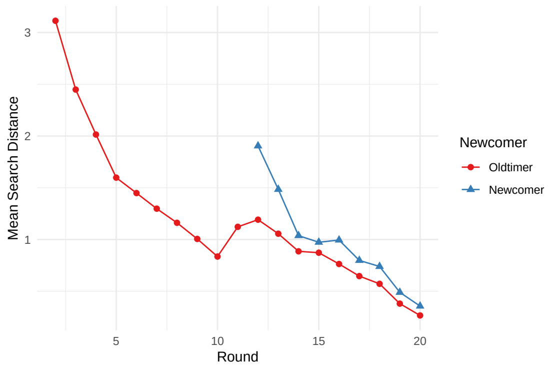

(b) Search distance across newcomers and oldtimers in rounds 12 to 15

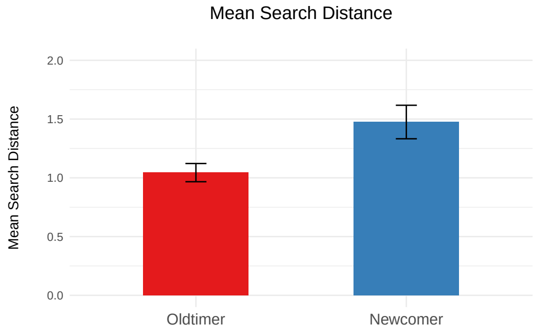

Figure 3.11: Search distance across newcomers and oldtimers. Notes: Search distance is measured as the number of decisions that differ from the decisions in the previous combination for each individual. Search distance is available only from round 12 on for newcomers. The analysis focuses on the three first search distance measures available: rounds 12 to 14. The analysis includes only Turnover treatments. Error bars indicate 95% confidence intervals.

(a) Newcomers and oldtimers payoff evolution

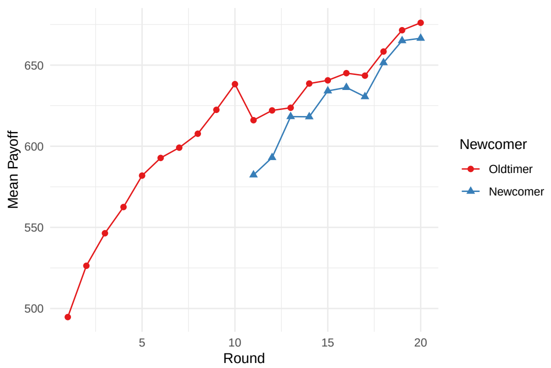

(b) Payoff across newcomers and oldtimers in rounds 12 to 14

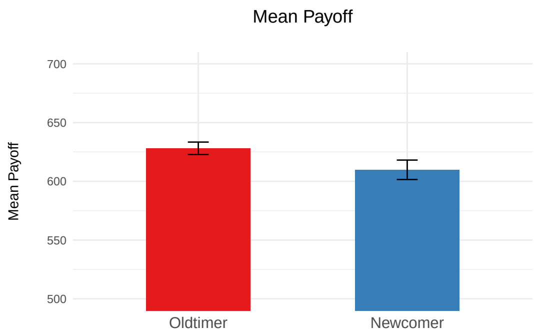

Figure 3.12: Payoff across newcomers and oldtimers Notes: The analysis focuses on the three first rounds where both search distance and payoff measures are available, that is, rounds 12 to 14. The analysis includes only Turnover treatments. Payoffs are derived from NK algorithms. Maximum payoff ranges from 725 to 874 and depends on each NK landscape. Error bars indicate 95% confidence intervals.

Table 3.10: Individual search distance in rounds 12 to 14: regression analysis

	Dependent variable:			
	Individual search distance			
	(1)	(2)	(3)	(4)
Newcomer	0.455***	0.457***	0.620***	0.616***
	(0.068)	(0.067)	(0.136)	(0.132)
Reference	0.318*	0.313*	0.499*	0.482*
	(0.134)	(0.135)	(0.204)	(0.205)
Know	−0.316*	−0.313*	−0.280	−0.286
	(0.134)	(0.135)	(0.204)	(0.205)
Newcomer ∗ Reference			−0.473*	−0.456*
			(0.193)	(0.188)
Newcomer ∗ Know			−0.080	−0.074
			(0.193)	(0.188)
Reference ∗ Know			−0.178	−0.157
			(0.288)	(0.290)
Newcomer * Reference * Know			0.439	0.419
			(0.273)	(0.266)
Round		−0.246***		−0.246***
		(0.037)		(0.037)
Sequence		−0.206***		−0.205***
		(0.029)		(0.029)
Constant	1.009***	4.718***	0.946***	4.659***
	(0.118)	(0.507)	(0.144)	(0.512)
Observations	1,692	1,692	1,692	1,692
Log Likelihood	−2,953.629	−2,912.725	−2,953.036	−2,912.260

Notes: Search distance is measured as the number of decisions that differ from the decisions in the previous combination for each individual. Search distance is available only from round 12 on for newcomers. The analysis focuses on the three first search distance measures available: rounds 12 to 14. The analysis includes only Turnover treatments. Treatments are coded as 2 binary variables: Reference and Know. Model (2) includes a control variable for round and sequence effects. Model (3) includes interaction terms, and model (4) includes controls and interaction terms. Robust standard errors are clustered at the group and individual levels and are in parenthesis. *p<0.05; **p<0.01; ***p<0.001.

Payoff Figure 3.12 indicate similar results regarding performance. The graph 3.12a shows that, even though the average oldtimers' payoff drops in round 11, the average payoff of newcomers is even lower. Similarly to search distance, the newcomers' average payoff quickly matches the oldimers'. Even if restraining the analysis to these first rounds, the difference in performance is not significant. In table 3.11, the Newcomers coefficient is negative and not significant. Being a newcomer compared to an oldtimer leads to no significant difference in performance in rounds 12 to 14.

Table 3.11: Individual payoff in rounds 12 to 14: regression analysis

	Dependent variable:			
	Individual payoff			
	(1)	(2)	(3)	(4)
Newcomer	-19.357^{***}	-19.454^{***}	-16.523	-16.353
	(4.392)	(4.334)	(8.738)	(8.621)
Reference	-26.887^{**}	-26.727^{**}	-21.623	-20.900
	(9.138)	(9.216)	(13.803)	(13.891)
Know	15.218	15.105	17.431	17.723
	(9.138)	(9.216)	(13.786)	(13.874)
Newcomer ∗ Reference			-10.841	-11.543
			(12.429)	(12.264)
Newcomer ∗ Know			-2.442	-2.709
			(12.409)	(12.244)
Reference ∗ Know			-8.128	-9.052
			(19.516)	(19.641)
Newcomer * Reference * Know			15.089	15.939
			(17.580)	(17.347)
Round		9.707^{***}		9.707^{***}
		(2.452)		(2.453)
Sequence		10.108^{***}		10.138^{***}
		(1.875)		(1.876)
Constant	635.609^{***}	484.615^{***}	633.937^{***}	482.623^{***}
	(8.060)	(33.212)	(9.727)	(33.679)
Observations	1,692	1,692	1,692	1,692
Log Likelihood	$-9{,}972.220$	$-9{,}946.774$	$-9{,}957.805$	$-9{,}932.302$

Notes: The analysis focuses on the three first rounds where search distance and payoff measures are available: rounds 12 to 14. The analysis includes only Turnover treatments. Treatments are coded as 2 binary variables: Reference and Know. Model (2) includes a control variable for round and sequence effects. Model (3) includes interaction terms, and model (4) includes controls and interaction terms. Robust standard errors are clustered at the group and individual levels and are in parenthesis. *p<0.05; **p<0.01; ***p<0.001.

When entering the game, newcomers adopt different search behaviors compared to oldtimers.

In the first rounds, they display a significantly higher search distance than oldtimers. However, this difference quickly disappears, and newcomers adopt similar search behaviors to oldtimers. This difference in search behaviors does not translate into a payoff difference. The bad news is that newcomers' entry does not lead to discovering alternative, better solutions. Newcomers quickly adopt their group's solution and search attitude and conform to group norms, but their entry does not lead to further exploration. However, the good news is that Turnover does not result in a performance difference either. T-tests report no significant payoff differences in round 20 across Reference ($t = 0.632, df = 592.34, p - value = 0.527$) and Turnover treatments ($t = 1.658, df = 551.59, p - value = 0.0977$). In this laboratory setting, turnover is not a double burden. Even though turnover does not increase exploration, it at least does not penalize the group in the long run.

3.6 Discussion

This study's goal was to explore whether two treatments, altering reference points and inducing team turnover, could help connected teams avoid over-exploration equilibrium and discover alternative solutions. Both treatments lead to a minor and short-lived increase in the search distance. Despite this change in search behavior, the increase in search distance is not substantial enough to lead to the discovery of better solutions. Moreover, the increase in search distance does not translate into improved payoff. Instead, payoffs decrease significantly after the treatment in all conditions before returning to their initial improving trend. Overall, the groups achieved medium performance compared to the benchmarks generated with the Monte Carlo simulation, indicating that participants performed at the level of connected teams primarily exploiting known solutions. This suggests that while the treatments temporarily increased exploration, they were insufficient to sustain higher exploration levels, which could have led to a performance matching the high benchmark.

In turnover treatments, newcomers initially exhibited a higher search distance than oldtimers, suggesting a potential for increased exploration. However, this difference is short-lived as newcomers quickly conform to their groups' existing search behaviors. Although newcomers' performance is lower than oldtimers' in the rounds following newcomers' entry, the payoff difference is not significant. In the long run, neither of the treatments hurts the teams' performance. At around 20, all groups express similar average payoff. This finding indicates that turnover does not substantially enhance exploration and leads to better solutions as hypothesized, but it does not detrimentally affect the groups' performance over time.

Alternative incentive schemes The experiment incentivizes participants based on their average performance across all rounds. I chose this approach for two main reasons. First, to simulate a realistic scenario where consistent performance is typically rewarded. Employees are regularly evaluated based on sustained performance rather than isolated successes. Second, I intentionally chose an incentive scheme that does not inherently encourage exploration. The study aims to examine if turnover could increase exploration. Using an incentive scheme that naturally hinders exploration, I set a high bar for the treatment. Any observed deviation towards increased exploration would be more convincing evidence of the treatment's effectiveness compared to a setting where the incentive scheme promotes exploration.

However, alternative incentive schemes could have been employed and would have probably significantly impacted participants' exploration behaviors. The first alternative is rewarding only the last round. Rewarding only final performance lowers the opportunity cost of exploring and experimenting with new strategies. Rewarding long-term success and tolerating early failures can better support innovative efforts (Ederer & Manso 2013). On the contrary, traditional pay-for-performance schemes hinder innovation by discouraging risk-taking. Applying this to our study, rewarding only the final round could maintain the initial exploration or the exploration triggered by treatments, as participants would have less to lose by exploring.

A second alternative is the use of collective incentives. This might be especially important in teams as the exploration efforts of one team member can benefit the entire team. In our current setting, individual incentives might have encouraged participants to avoid exploration and instead wait for others to take risks with the caveat that, in the end, nobody explores. In contrast, collective incentives could foster coordination and exploration. When groups are incentivized at the group rather than the individual level, they are less likely to over-exploit social information and engage in maladaptive scrounging behaviors (Deffner et al. 2024). Collective incentives encourage participants to be more selective in using social information, promoting a healthier balance between exploration and exploitation. Using collective incentives in the experiment could have mitigated the over-exploitation of known solutions and promoted broader exploration within the team.

A third alternative would be to compare groups against each other to introduce competitive pressure. In the experiment, groups operated in isolation, with some receiving information that better solutions exist but without the direct competitive pressure that could drive further exploration. Increased competition can foster cooperation and trust, which suggests that competitive pressure might similarly encourage exploration and innovation (Francois et al. 2018). In competitive environments, participants might feel more compelled to explore novel strategies to outperform their peers, potentially leading to the discovery of better solutions. In our setting, introducing incentives based on relative performance compared to other groups could have fostered increased exploration.

More broadly, I could have used alternative incentive schemes that align more closely with the behaviors I wanted to observe. In particular, I could have used incentive schemes promoting exploration by rewarding risk-taking and learning from failures. Participants might have been more inclined to explore if they knew their risk-taking efforts would be rewarded, even if immediate payoffs were not always positive. Future research should explore these alternative incentive schemes. For instance, comparing schemes such as rewarding the final round, using collective incentives, and introducing inter-group competition. These alternative incentive schemes could be combined with disruptions such as turnover to understand whether incentives promote exploration from the beginning or can sustain the exploration triggered by the treatments.

3.6.1 Turnover in the field and the lab

In management, turnover is correctly associated with significant costs, disruptions in operations, knowledge and skill losses, and performance decline. In this study, I aimed to explore whether the disruptive nature of turnover could have beneficial effects in the specific context of solving complex problems. Specifically, whether turnover could encourage exploration when it was necessary to find more effective solutions. The results show that while turnover can cause a brief increase in exploration, this effect is only temporary and insufficient to lead to discovering alternative solutions. Newcomers quickly adjust their behaviors to conform to the group.

The comparison between participants' performance and the simulated performance benchmarks suggests that the teams were underperforming. In that regard, the study reproduces known but counter-intuitive findings (Derex & Boyd 2016). Chances are that other highly connected teams are not performing at their full potential due to the over-reliance on social cues and conformity. Unfortunately, introducing diversity through turnover did not produce the expected increase in exploration and performance. A brief increase was relatively easy to trigger, disruptions could be strategically used to increase briefly exploration when needed. Alternative incentive schemes offer a promising avenue to obtain a stable and optimal level of exploration.

Nonetheless, turnover in the lab differs in several aspects from turnover that happens outside the lab. Contrary to many real-life scenarios, the chosen task does not require memorization. When a team member leaves, there is no loss of knowledge. One detrimental aspect of Turnover is the loss of knowledge, as reported in studies on Transactive Memory Systems (TMS). In this study, we avoid this complexity by using a task that does not require memorization. When a team member leaves, the loss of knowledge is negligible. In my study, turnover did not lead to a decrease in performance in the long run. This could be different in situations where the performance in the task depends on the knowledge acquired by the members.

Second, the turnover in this study involved the departure of the top-performer, not the team leader. In my experiment, every team member played an equal role, with uniform access to information and no designated leadership responsibilities. The observed effects of turnover are linked to the top-performer's departure but do not reflect the effects of the departure of a leader. The departure of a leader might lead to more detrimental effects than those observed in this study. Future studies could explore the consequences of leadership turnover and compare them with those of top-performer turnover.

The advantage is that performance is not significantly impacted. Newcomers do not help find more effective solutions but do not penalize the group's performance. With teams increasingly subject to change, this can be an encouraging result. In this laboratory situation of complex problem-solving, turnover did not lead to a performance decline. The question is how we can minimize its impact in other complex problem-solving situations outside the lab. Future studies could add layers of complexity to the task to approach a real-life setting while trying to preserve the non-significant influence of turnover on performance.

3.6.2 Limitations and avenues for future research

While these findings emphasize interesting avenues for reference points and team composition changes to encourage exploration, this study has several limitations. First, I worked with groups of size three, primarily due to budget constraints. While practical, this size has limitations. Small groups may not fully capture the effects of connectivity on problem-solving (Muthukrishna & Henrich 2016). In very small groups, like in the present setting, limited connections can restrict information flow and innovation. Larger groups can foster more innovation. They benefit from a wider pool of knowledge and skills, leading to better exploration and refinement of ideas. Larger groups might sustain exploration longer and achieve more breakthroughs. They balance cooperation and competition better, enhancing problem-solving. Future studies should use larger groups to understand whether our results would extend to larger teams.

Second, the studies rely exclusively on WEIRD (Western, Educated, Industrialized, Rich, and Democratic) subjects. WEIRD samples often exhibit cognitive, social, and behavioral patterns not representative of the global population (Henrich et al. 2010). Cultural differences significantly impact organizational behavior and decision-making processes (Banks 2023, Pitesa & Gelfand 2023). Consequently, our results might not apply to non-WEIRD populations. Moreover, students tend to be more homogeneous regarding age, socioeconomic status, and life experience compared to the general population (Peterson 2001). The over-reliance on student samples in organizational behavior research has been questioned for this reason (Hanel & Vione 2016). To address these limitations, I welcome future research that includes more diverse and representative samples, ensuring the broader applicability of the findings.

Third, I assumed that exploration did not generate costs beyond opportunity costs. However, the assumption could be relaxed. One can imagine that drastically changing a product, for example, generates additional costs. Including these costs in the experiment would probably further reduce the exploration. Generating an increase in exploration, even temporarily, might be even more challenging than in this setting, where exploration only generates opportunity costs.

In conclusion, this study explored whether altering reference points and inducing team turnover could help connected teams break from an over-exploitation equilibrium. Both interventions led to a brief increase in search distance. However, the increase in search distance did not increase performance. Despite initially displaying more exploratory behavior than oldtimers, newcomers quickly conformed to existing team norms. These results suggest that while new reference points and team turnover might introduce temporary variability in search behavior, the induced exploration does not necessarily translate into lasting changes or improved outcomes. Importantly, neither intervention had a long-term detrimental impact on team performance. Turnover did not lead to the hoped-for increase in exploration behaviors, but neither penalized the group performance.

Bibliography

Abelson, M. A. & Baysinger, B. D. (1984), 'Optimal and dysfunctional turnover: Toward an organizational level model', *Academy of management Review* **9**(2), 331–341.

Aime, F., Johnson, S., Ridge, J. W. & Hill, A. D. (2010), 'The routine may be stable but the advantage is not: Competitive implications of key employee mobility', *Strategic Management Journal* **31**(1), 75–87.

Banks, G. C. (2023), 'Eight puzzles of leadership science', *The Leadership Quarterly* **34**(4), 101710.

Baumann, O., Schmidt, J. & Stieglitz, N. (2019), 'Effective search in rugged performance landscapes: A review and outlook', *Journal of Management* **45**(1), 285–318.

Billinger, S., Srikanth, K., Stieglitz, N. & Schumacher, T. R. (2021), 'Exploration and exploitation in complex search tasks: How feedback influences whether and where human agents search', *Strategic Management Journal* **42**(2), 361–385.

Billinger, S., Stieglitz, N. & Schumacher, T. R. (2014), 'Search on Rugged Landscapes: An Experimental Study', *Organization Science* **25**(1), 93–108.

Bromiley, P. (2010), 'Looking at prospect theory', *Strategic Management Journal* **31**(12), 1357–1370.

Chen, D. L., Schonger, M. & Wickens, C. (2016), 'otree—an open-source platform for laboratory, online, and field experiments', *Journal of Behavioral and Experimental Finance* **9**, 88–97.

Chimento, M., Alarcón-Nieto, G. & Aplin, L. M. (2021), 'Population turnover facilitates cultural selection for efficiency in birds', *Current biology* **31**(11), 2477–2483.

Csaszar, F. A. (2018), 'A note on how NK landscapes work', *Journal of Organization Design* **7**(1), 15.

Deffner, D., Mezey, D., Kahl, B., Schakowski, A., Romanczuk, P., Wu, C. M. & Kurvers, R. H. (2024), 'Collective incentives reduce over-exploitation of social information in unconstrained human groups', *Nature Communications* **15**(1), 2683.

Denrell, J. & March, J. G. (2001), 'Adaptation as information restriction: The hot stove effect', *Organization science* **12**(5), 523–538.

Derex, M. & Boyd, R. (2016), 'Partial connectivity increases cultural accumulation within groups', *Proceedings of the National Academy of Sciences* **113**(11), 2982–2987.

Døjbak Håkonsson, D., Eskildsen, J. K., Argote, L., Mønster, D., Burton, R. M. & Obel, B. (2016), 'Exploration versus exploitation: Emotions and performance as antecedents and consequences of team decisions', *Strategic Management Journal* **37**(6), 985–1001.

Ederer, F. & Manso, G. (2013), 'Is pay for performance detrimental to innovation?', *Management Science* **59**(7), 1496–1513.

Faessler, L. (2023), 'Turnover effects on exploration and exploitation: An nk landscape experiment', `https://doi.org/10.17605/OSF.IO/AGVBY`. Preregistration.

Francois, P., Fujiwara, T. & Van Ypersele, T. (2018), 'The origins of human prosociality: Cultural group selection in the workplace and the laboratory', *Science advances* **4**(9), eaat2201.

Giannoccaro, I., Galesic, M., Massari, G. F., Barkoczi, D. & Carbone, G. (2020), 'Search behavior of individuals working in teams: A behavioral study on complex landscapes', *Journal of Business Research* **118**, 507–516.

Goldstone, R. L., Wisdom, T. N., Roberts, M. E. & Frey, S. (2013), Chapter One - Learning Along With Others, *in* B. H. Ross, ed., 'Psychology of Learning and Motivation', Vol. 58, Academic Press, pp. 1–45.

Greiner, B. (2004), 'An online recruitment system for economic experiments'.

Hanel, P. H. & Vione, K. C. (2016), 'Do student samples provide an accurate estimate of the general public?', *PloS one* **11**(12), e0168354.

Henrich, J., Heine, S. J. & Norenzayan, A. (2010), 'Most people are not weird', *Nature* **466**(7302), 29–29.

Kline, M. A. & Boyd, R. (2010), 'Population size predicts technological complexity in oceania', *Proceedings of the Royal Society B: Biological Sciences* **277**(1693), 2559–2564.

Koop, G. J. & Johnson, J. G. (2012), 'The use of multiple reference points in risky decision making', *Journal of Behavioral Decision Making* **25**(1), 49–62.

Kostopoulos, K. C. & Bozionelos, N. (2011), 'Team Exploratory and Exploitative Learning: Psychological Safety, Task Conflict, and Team Performance', *Group & Organization Management* **36**(3), 385–415.

Kwon, K. & Rupp, D. E. (2013), 'High-performer turnover and firm performance: The moderating role of human capital investment and firm reputation', *Journal of Organizational Behavior* **34**(1), 129–150.

Lazer, D. & Friedman, A. (2007), 'The Network Structure of Exploration and Exploitation', *Administrative Science Quarterly* **52**(4), 667–694.

March, J. G. (1991), 'Exploration and exploitation in organizational learning', *Organization science* **2**(1), 71–87.

Muthukrishna, M. (2020), 'Cultural evolution and the paradox of diversity', *National Academy of Engineering: The Bridge* **50**(4), 26–28.

Muthukrishna, M. & Henrich, J. (2016), 'Innovation in the collective brain', *Philosophical Transactions of the Royal Society B: Biological Sciences* **371**(1690), 20150192.

Muthukrishna, M., Shulman, B. W., Vasilescu, V. & Henrich, J. (2014), 'Sociality influences cultural complexity', *Proceedings of the Royal Society B: Biological Sciences* **281**(1774), 20132511.

Peterson, R. A. (2001), 'On the use of college students in social science research: Insights from a second-order meta-analysis', *Journal of consumer research* **28**(3), 450–461.

Pitesa, M. & Gelfand, M. J. (2023), 'Going beyond western, educated, industrialized, rich, and democratic (weird) samples and problems in organizational research', *Organizational Behavior and Human Decision Processes* **174**, 104212.

Schimmelpfennig, R., Razek, L., Schnell, E. & Muthukrishna, M. (2022), 'Paradox of diversity in the collective brain', *Philosophical Transactions of The Royal Society B* **377**(1843), 20200316.

Simon, H. A. (1957), 'Models of man; social and rational.'.

Tversky, A. & Kahneman, D. (1992), 'Advances in prospect theory: Cumulative representation of uncertainty', *Journal of Risk and uncertainty* **5**, 297–323.

Tzabbar, D. & Kehoe, R. R. (2014), ‘Can Opportunity Emerge From Disarray? An Examination of Exploration and Exploitation Following Star Scientist Turnover’, *Journal of Management* **40**(2), 449–482.

Wagner III, J. A., Humphrey, S. E., Meyer, C. J. & Hollenbeck, J. R. (2012), ‘Individualism–collectivism and team member performance: Another look’, *Journal of Organizational Behavior* **33**(7), 946–963.

Yoon, J. & Kayes, D. C. (2016), ‘Employees’ self-efficacy and perception of individual learning in teams: The cross-level moderating role of team-learning behavior’, *Journal of Organizational Behavior* **37**(7), 1044–1060.