Echo Chambers and Their Effects on Economic and Political Outcomes

Gilat Levy[1] and Ronny Razin[2]

[1]Department of Economics, London School of Economics, London WC2A 2AE, United Kingdom; email: g.levy1@lse.ac.uk
[2]Department of Economics, London School of Economics, London WC2A 2AE, United Kingdom; email: r.razin@lse.ac.uk

Abstract

We survey the economics literature on echo chambers.[1] We identify echo chambers as arising from a combination of two phenomena: (i) The choice of individuals to segregate with like-minded ones, i.e., the creation of *chambers*; (ii) Behavioural biases that induce polarisation when individuals exchange beliefs in these chambers, i.e., the *echo*. We summarise the literature on (i) and (ii) and suggest how to combine the two literatures to gain insights about the effects of echo chambers on economic and political outcomes. We end by suggesting pathways for future research and discuss policy interventions to alleviate echo chambers.

[1]This survey is not meant to be an exhaustive overview of the literature in economics about echo chambers. Rather it reflects both authors' reading of this literature as well as the echo chamber they work within.

Xxxx. Xxx. Xxx. Xxx. YYYY. AA:1–29

https://doi.org/10.1146/((please add article doi))

Contents

1. Introduction

1.1. Echo chambers: Motivation

Evoking the fate of Narcissus, and his estranged lover, Echo, echo chambers have been blamed in recent years for many of our society's ailments.[1] It is thought of as the engine behind phenomena such as political gridlock and constitutional crises, the rise in violence, extremism, populism, and polarisation, as well as economic outcomes such as lower social mobility and higher inequality. In politics, it has been repeatedly blamed for increased polarisation and political fragmentation. Bishop (2009), in his influential book "The Big Sort", claims that the clustering of like-minded Americans "is tearing us apart". He makes the connection between the segregation of Americans and the political polarisation and gridlock of recent decades. Barber & McCarty (2015) argue that the resulting polarisation undermines the legislative quality in the US. Sunstein (2001) suggests how echo chambers were behind the impeachment of President Clinton.

Turmoil in financial markets has also been attributed to echo chambers. In their book "Animal Spirits", Akerlof & Shiller (2009) argue that the business cycle is tied to feedback loops between speculative economic activity and the discussions that these activities incite. A downward movement in stock prices, for example, generates chatter and media response,

[1] The myth of Echo and Narcissus tells the story of a "talkative nymph" who is cursed by Juno, making her able to only finish others' sentences, and unable to say anything on her own. She falls in love with Narcissus, a young man who is unwelcoming to all around him. Narcissus rejects her love. Echo prays to Venus, who makes Echo disappear, until she "remains a voice" and "is heard by all".

Narcissus rests by a spring, and whilst drinking, falls deeply in love with his image reflected in the water. He then wastes away with love for himself, echoing the manner in which Echo did earlier.

and reminds people of long-standing pessimistic stories and theories. These stories, newly prominent in their minds, incline them toward gloomy intuitive assessments. As a result, the downward spiral can continue: declining prices cause the stories to spread, causing still more price declines and further reinforcement of the stories. This phenomenon is empirically examined in Jiao et al. (2016).

In this survey we explore the mechanism behind echo chambers; "chambers" arise when individuals segregate with like-minded. "Echos" happen when individuals fail to process the information that is correlated and repeated within the chambers. The combination of segregation and communication with those with similar beliefs can induce then extremism and polarisation in society. Our focus is on the two-way relationship between segregation and beliefs.

This recent interest in echo chambers stems to a large degree from the dramatic technological changes in communication and media in the last few decades. But environments and behaviour which enable echo chambers are not new and have been around for centuries. We will therefore focus the survey on the general tendency of people to segregate, both off-line and on-line. In fact, the evidence about the extent of echo chambers on-line is not conclusive, as on-line communication also facilitates communication of diverse opinions. While Quattrociocchi et al. (2016) find that there is very little communication between groups on Facebook, and Del Vicario et al. (2016) find that conspiracy theories and scientific news generate homogeneous and polarised communities, Dubois & Blank (2018) find on the other hand that those who are interested in politics and those with diverse media diets tend to avoid echo chambers. Moreover, while the internet is more segregated than off-line media, it is significantly less segregated than face-to-face interactions, as Gentzkow & Shapiro (2011) show. And Boxell et al. (2017) show that greater internet use is also not necessarily associated with more political polarization. Specifically, they find that polarization has increased the most among the elderly, who are the least likely to use the Internet and social media, suggesting that the role of these factors is limited.

1.2. The mechanics of echo chambers

An echo chambers is a metaphor based on the acoustic echo chamber, where sounds reverberate in a hollow nclosure. The term has been used to denote the phenomenon of the amplification and reinforcement of beliefs by communication and repetition inside a closed, like minded, community. The Cambridge English Dictionary defines echo chambers as: "An environment in which a person encounters only beliefs or opinions that coincide with their own, so that their existing views are reinforced and alternative ideas are not considered."

To understand echo chambers and their influence our framework decomposes the term into two:

1. Chambers: Individuals segregate with like-minded in terms of preferences, beliefs or attitudes.

2. Echo: Individuals are influenced by the beliefs of those they communicate with in their chamber, in a non-rational manner.

To understand echo chambers we first need to understand why people belong to different chambers. Individuals take many decisions that affect which sources of information or influence they are exposed to. These could be big and important decisions that we take

only infrequently, such as location or career decisions, or smaller and frequent decisions we take, such as when we decide what to read, who to talk to, and what to search for on-line. Sometimes we take these decisions without thinking about how this will affect us later, for example, when moving into a new city mainly because of the salary, we may not anticipate the effect this will have on our future political views. Other times we make a more informed decision, such as when we think about the schools we want to send our children to. In this case we might anticipate the role the school will have on the beliefs and attitudes of our children.

Once people are in their chamber, this could be physical or virtual, the patterns of communication and influence shape and affect attitudes, opinions and even preferences. The term echo reflects the idea that, within a chamber, information might be repeated and exaggerated, but also the fact that you might hear a selection of opinions, those that are close to your initial views. This is especially a problem when people tend to segregate with other like-minded people.

Chambers and echoes are naturally connected and they co-evolve. The choice of chambers affects the types of echo effects we are exposed to. It determines what kind of information will circulate in the chamber and the patterns of repetition and correlation between information sources. In turn, our beliefs, attitudes and preferences influence our choices in terms of future segregation. If you were brought up to fear or dislike other groups in society, chances are you will choose to live in a neighbourhood where these are not represented. Sometimes, segregation and echo effects happen simultaneously. Psychologists have explored a tendency to avoid information or beliefs that do not agree with our own (Bessi 2016). For example, when weeding through the infinite stream of on-line content, individuals' brains simply focus on content which they like and feel close to, given their beliefs or attitudes. Similarly, this happens when people buy and consume news that they know will fit with their biased views (Mullainathan & Shleifer 2005).

Other times these effects happen sequentially. We can shape our children's beliefs by the school choices we make for them, and later, when they make their own decisions, they choose who to segregate with. Sometimes segregation decisions are taken according to aspects which are independent of the unintended consequences of echo effects; one cannot fully control the types of influences one's children will be exposed to.

1.3. Plan of the survey

The survey proceeds as follows. In Section 2 we discuss the literature on segregation. We illustrate the different reasons why individuals will segregate according to their preferences or their beliefs. In Section 3 we survey the relevant cognitive biases of information processing that imply that belief updating will lead to extremism and polarisation. In particular we focus on three prominent cognitive biases, correlation neglect, selection bias and confirmation bias. We discuss why these biases arise naturally in the context of segregation. We also examine the normative implications of these cognitive biases. In Section 4 we illustrate why feedback effects between segregation and cognitive biases are important. We focus on a dynamic model in which segregation and polarised beliefs each fuel the other. In the context of a model of labour market discrimination, we discuss the long term sustainability of segregation, polarisation, and their harmful economic outcomes. In Section 5 we conclude by suggesting avenues for future research, theoretical as well as empirical.

2. The creation of chambers: Segregation

Why do people segregate with those who are like-minded? This phenomenon has long been recognised in the social sciences. Sociologists have observed that in many contexts, people tend to connect with and favour others who are similar.[2] Sociologists find that people exhibit homophily based on demographic or psychological characteristics. Social psychologists have found that this tendency to segregate with similar individuals can be triggered even with minimal distinguishing differences between people.

Within economics, the key explanation for why people segregate is the existence of complementarities. Specifically, people will choose to segregate due to complementarities in preferences, or in beliefs, that enable better economic or political interactions. The models that we discuss below all share this feature: From traditional Tiebout sorting models that focus on complementarities in preferences for public goods, to the more recent literature that considers complementarities through peer effects that enable better cultural transmission of preferences. Similarly, complementarities in beliefs exist when transmission of beliefs is important, when people wish to protect their belief system, or when individuals prefer to interact with those with similar beliefs to facilitate cooperation and communication. As we discuss below, people often misperceive these complementarities and hence segregation can become excessive.

In this Section we first discuss the increase in segregation in recent decades, and how it is linked to political and economic outcomes, such as political polarisation and income inequality. We then put forward models that illustrate why individuals segregate with the like-minded. We first describe segregation according to preferences and then according to beliefs. Both models are important for the purpose of understanding echo chambers. The type of segregation, preference versus belief based, has different implications for dynamic analysis as preferences are typically fixed while beliefs are relatively easy to change and mould.[3]

2.1. The rise of segregation and its consequences

The rise in the use of social media has certainly refuelled the interest of scientists in the causes and consequences of segregation. For example Bakshy et al. (2015) analyse how on-line networks influence exposure to perspectives that cut across ideological lines. They show that individuals' choices played a stronger role in limiting exposure to cross-cutting content. Gilbert & Bergstrom (2009) look at blogs and find that agreement outnumbers disagreement in blog comments by more than 3 to 1.[4] However, off-line segregation is as important, if not more. Gentzkow & Shapiro (2011) analyse the impact of the internet on the segregation of information consumption based on aggregate and individual data. They use an isolation index to define the level of ideological segregation. They find that the internet is more segregated than off-line media, but significantly less segregated than face-to-face interactions. They show that individuals' communication networks are segregated across

[2]For a survey of the research on homophily, see McPherson et al. (2001).

[3]Bohren et al. (2017) provides a theoretical model as well as an empirical illustration of the difference between two types of discrimination, preference-based versus belief-based. They show how one can use the dynamic setting to distinguish between the two.

[4]In the next Section we discuss more empirical work looking at the manifestation of confirmation bias on-line where individuals enhance their beliefs by filtering out news or opinions that do not accord with their views.

work colleagues, friends, family and neighbourhood associations, according to socioeconomic parameters and political preferences.[5]

Indeed, empirical studies suggest that physical, off-line, segregation has increased in Western societies in the last decades. For example, in the US, since 1970, residential segregation has been on the rise.[6] Moreover, this trend in residential choices seems to be correlated with important economic and political variables. Reardon & Bischoff (2011) study the relation between income inequality and segregation according to income in the US. They find that residential segregation and income inequality have been following a remarkably similar trend during 1970-2000. Chetty et al. (2014) look at the relation between segregation in the US and social mobility. They show large gaps between different localities, so that the more segregated areas have much lower social mobility. Alesina & La Ferrara (2005) survey the large literature that studies the relation between segregation and economic outcomes such as growth. Relatedly, a large theoretical and empirical literature has also analysed the effects of segregation according to ethnic groups on different measures of social cohesion, conflict and social attitudes (see Putnam 2007; Field et al. 2008; Sturgis et al. 2011; Uslaner 2012).[7]

The effects of segregation (e.g., income segregation) on political outcomes have also been explored; Bishop (2009) coined the term "The Big Sort", describing the patterns of residential segregation in the US and its effects on polarisation of political beliefs. According to McCarty et al. (2008) there is a close correlation between economic inequality and polarisation in the US; specifically, increased growth in the top of the income distribution leads to higher inequality and demand for conservative legislators. In turn, this increases polarisation and dampens the political response, which further increases inequality. Indeed, Bartels (2008) and Gilens (2012) find that policies more often reflect the preferences of the wealthy compared to those at the bottom.

Segregation can also affect economic opportunities in life through the availability of knowledge and information. Calvó-Armengol & Jackson (2004) show how segregation of individuals into different networks can benefit some and not others as information about job opportunities flows only to selective network members. A similar mechanism underlies Curtis & Warner (1992) who study the benefits of the "old boys network". One way of creating these types of networks is through school choices; school choices tend to be persistent and affected by parental background and beliefs. Evans & Tilley (2012) find that 43% of the privately educated in the UK who have children, have sent them to private schools, nearly five times the rate for parents who went to state schools. Importantly, the different attitudes and beliefs about school graduates have real behavioural implications for labour markets, through occupational choice and employment decisions. For example, in the UK, private school male graduates are up to 10% more likely to be hired to top

[5] About two-thirds (63%) of Americans say family and friends are an important way they get news, whether on-line or off-line; 10% see them as the most important (see Mitchell et al. 2016).

[6] A similar trend is found in the UK. For example, see Cantle & Kauffman (2016).

[7] Some theoretical research points to some potential benefits of segregation. Bala & Goyal (1998) show how in the context of social learning, local learning can prevent herding and thus segregation can be efficient. By incorporating peer effects in the human capital acquisition Chaudhuri & Sethi (2008) highlight another potential efficiency concern with integration; increasing integration increases (decreases) the cost of human capital acquisition for individuals in the high (low) skilled group. They show that increased integration can lead to convergence to an equilibrium that is either skill enhancing or skill reducing. Increasing integration and implementing equal treatment across groups may lead to fundamentally different outcomes.

jobs than state school graduates with the same grades from the same university.[8] Indeed, 50% of private school students believe that people who attend their school will be 'very successful' compared to 9% of state school pupils.[9] Similarly, in the US, private schools lead the tables in terms of placements at top universities, even though students from private schools or selective state schools perform no better than those in standard state schools in achievements tests.[10]

After establishing that segregation is an important aspect in determining political and economic outcomes we now proceed to understand why it happens. Below we discuss the reasons for segregation with like-minded individuals, first according to shared preferences and then according to beliefs/attitudes. Our plan is to illustrate how segregation with like-minded individuals can affect beliefs and create the consequences discussed above.

2.2. Preference-based segregation

Many traditional models in economics consider sorting according to preferences. For example, Tiebout models, originating from Tiebout (1956), consider environments in which communities choose the level of local public goods provision and finance them via taxation. In this framework, those who care more about a particular set of goods are better off congregating in their own locality. Thus, preference complementarities can fuel segregation. Relatedly, club good models as in Iannaccone (1992) show how individuals sort themselves into religious groups in order to enjoy complementarities in the production of religious goods such as rituals and communal praying.

Recent contributions in this literature highlight the equilibrium effects which arise when individuals segregate according to taste or income. Bénabou (1996) looks at the effect of segregation on growth. In his paper, agents interact through local public goods, such as school funding, and economy-wide linkages, such as knowledge spillovers. Sorting families into homogeneous communities often minimises the cost of existing heterogeneity, but mixing reduces heterogeneity faster. Integration therefore tends to slow down growth in the short run yet raise it in the long run.[11] Baccara & Yariv (2016) study the formation of peer groups in an environment where each group can produce two distinct public goods that only give utility to people within a group. The agents are free to choose the size of the group and types of group members, along with their choice of public good contribution. When contribution costs are low relative to connection costs, mutually optimal groups are relatively homogeneous.[12] Relatedly, Goyal et al. (2017) look at environments in which individuals prefer to coordinate with others but differ in their preferred action, and examine what groups are formed. The theoretical model predicts different possibilities, some in

[8]These statistics are taken from a recent study by the Social Mobility and Poverty Commission in the UK, which tracked 20,000 students. There is also evidence that state school graduates outperform compared to the expectations about them. These findings are consistent with discrimination due to biased beliefs and the subsequent adjustment of beliefs due to learning. See a 2014 report by HEFCE.

[9]See Nasiroglu (2016).

[10]See the 2007 study of the Centre on Education Policy, using NELS data from 1988-2000, which takes family background into account. See also Abdulkadiroğlu et al. (2014) and Dobbie & Fryer (2014) who find that having peers with high achievements or attending a school with a low racial mix have no effect on pupils' attainment.

[11]See also Durlauf & Sheshadri (2017).

[12]See also Peski (2008).

which more sorting occurs and some in which individuals coordinate on a single action; their experiments show however that agents differentiate themselves more often, inefficiently.

2.2.1. Complementarities in networks: information flows and learning. The literature on social networks has many results pertaining to the study of segregation.[13] One strand of this literature focuses on complementarities in the ability to communicate with others. From the literature about strategic communication (Crawford & Sobel 1982) we know that the level of communication is inversely related to the distance in preferences. As a result we will expect communities that are more homogeneous in terms of preferences to imply higher levels of information sharing. This kind of complementarity is formalised in Galeotti et al. (2013) who study a model of multi-player communication in networks. Privately informed decision makers have different preferences about the actions they take, and communicate to influence each others' actions in their favour. They show that clusters of individuals with similar preferences will facilitate information transmission and will create complementarities in information and hence efficient decision making.

Similarly, Giovanniello (2018) shows how people choose to exchange information with like-minded to the effect that chambers are created. Specifically, she shows that while it is necessarily the case that information will travel along those with similar preferences, this is not sufficient. She considers a model in which voters can be ideologically close but still biased towards different parties. In that case, information e.g., about quality of political candidates, will not be exchanged between such voters. Thus voters have to be both close in their preferences and biased in the same direction.[14]

The above papers, while focusing on preference-based segregation, show that this type of segregation also has implications to what information is shared, and hence on the beliefs of individuals in the network.

2.2.2. Cultural transmission of preferences. Another important channel that encourages individuals to segregate is the case of cultural transmission, analysed first in Bisin & Verdier (2001). Specifically, if parents realise that their offspring's preferences are affected by the community and not just by upbringing, they may choose to live in neighbourhoods where others share the same preferences. Thus, complementarities arise through peer effects on transmission of values. Advani & Reich (2015) show how cultural transmission may hinder economic activity and foster segregation. They assume that individuals face a trade-off between cultural and economic incentives: an individual prefers to maintain her cultural practices, but doing so can inhibit interaction and economic exchange with those who adopt different practices. They find that a small minority group will adopt majority cultural practices and integrate. In contrast, minority groups above a certain critical mass, may retain diverse practices and may also segregate from the majority. They also test their predictions using data on migrants to the United States in the era of mass migration, and find support for the existence of a critical mass of migrants above which the social structure in heterogeneous populations changes discretely towards cultural distinction and segregation. García-Alonso & Wahhajz (2018) analyse the dynamic effects of an increase in cultural diversity within a population, due for example to an immigration wave. They

[13] See Jackson (2011) for a survey.

[14] While in the above papers homophily is endogenously derived, Golub & Jackson (2012) assume homophily and show that this implies a lower speed of convergence to a consensus.

analyse how the pace of change affects the level of segregation.

2.3. Belief-based segregation

We now consider the mechanisms behind why individuals who share similar beliefs might segregate together. These mechanisms and the segregation they create will later feed into our dynamic models of echo effects in the next section.

2.3.1. Cultural transmission of beliefs. We have seen above that when parents consider the values they transmit to their children, then due to peer effects, they may be inclined to segregate with individuals who share similar cultural traits. A similar argument can be applied to segregation according to beliefs rather than according to preferences. To give an example, imagine the thought process of parents who are deciding which school to send their child to. There might be many trade-offs involved in this decision, depending on the characteristics of these schools. But one thing the parents might have in mind is how each school will affect their child's beliefs, through socialisation with friends or through teachers (for example, one school might be secular while the other is a religious school). The parents might be worried that their children's beliefs are amenable to influence. Levy & Razin (2017) incorporate this scenario in a model that studies segregation in schooling (private versus state) and labour market discrimination. They show that the parents' dilemma implies segregation into different schools, according to parents' beliefs about the merits of education in the private versus the state school, and discrimination in the labour market.

2.3.2. Segregation to maintain beliefs: Religious segregation. A related reason for segregation is when individuals or groups seek to actively avoid knowledge or beliefs that are counter to their own. One important environment in which this can arise is the case of religious segregation. Religion plays an important role in the patterns of residential location we observe. Berman (2000) and Razin (2018) document how the ultra-Orthodox Jews in Israel segregate away from the secular Jewish population both physically but also through their lack of participation in the labour market and military service. In current day London, Brimicombe (2007) finds that: "The landscape of religion is found to be more highly segregated in contrast to the landscape of ethnicity." Field et al. (2008) find that more than 70 percent of the population in Ahmedabad in 2002 lived in completely homogeneous neighbourhoods.

While religious individuals may segregate for many reasons, for example, due to complementarities in behaviour, maintaining religious beliefs is another important reason. Levy & Razin (2012) suggest an informational reason for segregation: religious beliefs might be eroded by observing others' behaviours or beliefs. Specifically, Levy & Razin (2012) model a theology of reward and punishment in relation to behaviour in the social sphere. In particular, the theology of the religion makes a connection between the actions of a deity to the behaviour that individuals take in their day to day social interactions. Holding these beliefs allows the religious to sustain cooperative outcomes that may not be available otherwise. But religious beliefs are not static, as they may evolve given the personal experiences of the believer. Religious beliefs must be maintained and protected if they are to be sustained in the long run. To sustain religious beliefs individuals should be guarded from observing behaviours and outcomes that do not agree with their belief system. By segregating in closed communities the religious can sustain their beliefs by not observing the (possibly

good) fortunes of those who "sin".

Attempts to protect communities from information can be seen more generally when organisations may wish to protect a belief system, that may not be immune to updating in the face of real events. Censorship of books, media or internet content is familiar in many authoritarian regimes, creating de-facto segregation according to beliefs.

2.3.3. Segregation due to prejudice about others' behaviour. Another reason for segregation arises when beliefs of individuals are prejudiced against a particular population and individuals segregate to avoid interaction with this population. As a result, people with different beliefs or prejudice segregate as they all share a similar incentive to do so. Bradford & Kelejian (1973) were the first to document what they call the "white flight" from inner city neighbourhoods and towards predominantly white areas; Cantle & Kauffman (2016) document dynamic patterns in the UK from 2001 to 2011 and show strong evidence for this. They observe that: "Between 2001 and 2011 the White British population in England reduced as a percentage of the total population from 86.8% to 79.8% – a decrease of 8%. Although there was a decrease in the proportion of the population who were white in most areas, the decrease was much greater in the areas which had a low proportion of White British in 2001 than in areas which had a high proportion... This does indicate support for 'more mixing and more clustering', but they are not equivalent trends, the clustering is noticeably more marked." Kaufmann & Harris (2015) find that: "For London, between 2001 and 2011, around 620,000 White British people left the city, most of whom moved to whiter areas; whites left London at three times the rate of minorities (about 100,000 of the latter left London)."

While the reasons for this type of segregation could be correlated with income inequality, as possibly richer white individuals move to bigger houses away from inner cities, there is also direct evidence about different views that individuals hold conditional on their location choice. Causality is of course not clear-cut; it may be that once moving, individuals have changed their views to the worse. However, a more direct explanation would be that those that have moved are those with more prejudiced beliefs, and possibly, once segregated, their beliefs had deteriorated even further when they exchanged their beliefs. Indeed, Dustmann & Preston (2007) find strong evidence that racial or cultural prejudice is an important component to attitudes towards immigration in their study using the British Social Attitudes Survey. Similarly, Vertier & Viskanic (2018) show that in areas in France in which refugees were settled (which were randomly assigned), locals had improved their views on foreigners and voted less for the Front National, the extreme right-wing anti-immigration party. This provides evidence for the existence of prejudice.[15]

In Levy & Razin (2018a) we consider an environment in which individuals in the home society are prejudiced against foreign immigrants and are suspicious of their ability to cooperate in economic interactions or of their productivity. As a result, interactions between home society individuals and immigrants are inefficient which makes it worthwhile for prejudiced individuals to segregate away from immigrants. Specifically, it is those with the most prejudiced beliefs against immigrants that will segregate away.

We have discussed in this Section how individuals are motivated to segregate with like-

[15] Alesina et al. (2018) show evidence for the prevalence of misperceptions of natives about immigrants in their own country.

minded. While the traditional literature has looked at segregation according to preferences, segregation according to beliefs or attitudes has been only recently explored. Within these chambers of like-minded individuals, echoes can easily be created, which is the topic of our next Section.

3. The creation of the echo: Behavioural biases in belief updating

A large body of literature shows that segregation affects beliefs. In social psychology, contact theory posits that beliefs are affected by segregation through the different interactions between people from different groups (see Allport 1954; Hewstone & Brown 1986; Pettigrew & Tropp 2006; Hewstone 2009; Lowe 2018). Boisjoly et al. (2006), Algan et al. (2015), Burns et al. (2016) and Vertier & Viskanic (2018) show how interacting with different individuals affects attitudes towards one another. Kaufmann & Harris (2015) find significant effects of segregation on attitudes about the benefits of immigration.

How beliefs are affected by others' beliefs about us or by observing other pieces of information is of course a more general problem and not specific to segregation. Throughout each day we are exposed to large amounts of information, some of which we seek actively, some we consume more passively. How good are we at aggregating all these pieces of information? In economics, the traditional assumption of rationality implies that individuals are efficient in gleaning information from their surroundings. However, both political scientists and psychologists have typically taken a more pessimistic view of our ability to process information. In political science for example, a large literature documents the incompetency of voters in collecting and processing information. Voters have been shown to be poorly informed about what they vote on (Campbell et al. 1960; Kinder & Sears 1985; Bartels 1996; Delli Carpini & Keeter 1996) and to use the information they do have incorrectly (Lau & Redlawsk 2001; Achen & Bartels 2004; Wolfers 2007; Leigh 2009; Healy et al. 2010; Huber et al. 2012). As Bartels (1996) writes:

> One of the most striking contributions to the political science of half a century of survey research has been to document how poorly ordinary citizens approximate a classical ideal of informed democratic citizenship.

Psychologists have also taken a grim view of individuals' ability to make sense of the information presented to them. A good example for this is the strong response to the rationality assumption in economics, in a series of papers by Daniel Kahneman and Amos Tversky. These papers have revealed different biases that inflict individuals who are exposed to different pieces of information (see Rabin 1998). More recently, these results have spurred new research in economics, in behavioural economics and bounded rationality, that incorporates some of these biases into economic models.

In this section we survey a few of these cognitive biases as they relate to the creation of echo chambers. A good starting point to think about this would be to imagine yourself in your daily interactions with the people around you. You spend your day reading newspapers and on-line news content, talking to friends and family, talking to colleagues at work and might spend some time on social networks. Daily, this might amount to large quantities of information that you may want to sit back and process before you go to bed.

If individuals are rational and have correct beliefs about the nature of interactions in their network, no echo effects will exist. On average people will hold correct beliefs and there will be a limit to how polarised or extreme beliefs might be. In this survey, instead,

we are interested in the types of cognitive biases that might arise when people interact and glean information within their chambers.

A few aspects of your above interactions imply that it is not easy to aggregate all this information properly. For one thing, the network of interactions in your social milieu might imply that you cannot really follow where the information a friend is telling you came from. For example, a friend, Amir, might be telling you something. But Amir might have also talked with Neeve and you have told Neeve something similar the previous day. How then should you weigh what Amir tells you? Often, in these situations we might err by simply treating what Amir tells us as an independent piece of information. This is what we term correlation neglect below.

But there is another problem that could make your inference complicated which is related to the composition of your social network. In particular, as we saw above, one reason you like talking to Amir or Neeve is that they are very much like you. Therefore, Amir and Neeve will most likely say things that agree with your own views. In these cases, some individuals might err by over-weighing what Amir and Neeve say, due to a selection bias.

Finally, every now and then, perhaps at your workplace you encounter other individuals, such as Francesco, that have very different things to say than Amir and Neeve. In these situations do you fully take into consideration what Francesco says? "Confirmation Bias" is a bias that psychologists have documented in which not only will you put too low a weight on what Francesco says, but you might also become stronger in your opposing conviction after the encounter.[16]

The key mechanism we explore in this paper is how segregation and cognitive biases work together to create the effects of echo chambers. For this reason, below we focus on the above three cognitive biases which are tightly related to the features of segregation. We will now go over the mechanisms through which cognitive biases will exacerbate polarisation in the presence of segregation.

3.1. A basic model to introduce biases

The simplest way to think of the cognitive biases we consider is to assume that individuals do observe the information of others directly, but have difficulties interpreting this information. Consider the following model. Individuals try to learn about the state of the world ω, which could be high or low, $\omega \in \{l, h\}$. They all have a common prior that the states are equally likely.

For example, the state could correspond to the fate of the UK after Brexit, where a low state implies low growth and a high state high growth. Knowing the state informs the group about its policies. In the Brexit example above, information about the state will inform voters how to vote in a referendum about Brexit.

Individuals start with some beliefs about the states. Let q^i denote the belief of individual i that the state is high, with $1 - q^i$ denoting the belief of that individual that the state is low. The individual's belief could have been generated by receiving a signal $s \in \{l, h\}$, with an accuracy $\Pr(s = h|\omega = h) = \Pr(s = l|\omega = l) = q \geq \frac{1}{2}$. In this case, Bayes rule implies

[16]Note that there are other mechanisms that imply that information is not aggregated properly in a network or group. One such example is social learning, when individuals do not observe the information of others, but a coarse action representing this information, which is not a sufficient statistic of this information (Banerjee 1992; Bikhchandani et al. 1992).

that receiving a signal h will yield the (high) belief $q = \Pr(\omega = h|s = h)$, and receiving a signal l will yield the (low) belief that $\Pr(\omega = h|s = l) = 1 - q$. For example, this signal could be generated by reading an informative newspaper article about the effects of Brexit on the UK labour market.

When individuals interact in their social network they share their opinions with each other. To focus attention on cognitive biases, rather than any strategic considerations, let us assume that individuals share their true beliefs with each other. When exposed to these different opinions, how do individuals update their beliefs? This is what we consider in the sections below. Note that while we focus on three biases that relate to the creation of echo chambers, correlation neglect, selection bias and confirmation bias, this is in no way an exhaustive list of relevant biases. Alternatives include for example models in which individuals manipulate their own beliefs as in the motivated beliefs literature(Bénabou & Tirole 2011; Bénabou 2013; Bénabou & Tirole 2016; Le Yaouanq 2018), which we discuss in Section 3.8.

3.2. Correlation neglect

As discussed above, there are many reasons to believe that in social networks individuals' sources of information are correlated in complex ways. Correlation neglect is a cognitive bias where individuals simply ignore such correlation structures. Therefore, individuals with correlation neglect treat all sources of information as conditionally independent. This is a simple way to combine information sources into a unique prediction.

A recent empirical and experimental literature has shown that in complex environments, decision makers indeed ignore correlations to some degree. For example, Ortoleva & Snowberg (2015) use data from the 2010 Cooperative Congressional Election Study (CCES) to show how correlation neglect shapes political views. Eyster & Weizsäcker (2011), Kallir & Sonsino (2009), Bai et al. (2015) and Enke & Zimmerman (2017) provide experimental evidence for correlation neglect. Jiao et al. (2016) provide evidence for correlation neglect in stock prices that are discussed in on-line discussion groups.[17]

We now proceed to define correlation neglect formally, in the context of our model. To define correlation neglect, assume that N individuals exchange their beliefs q^i. In reality these beliefs might stem from a complex web of correlation relations. However if individuals neglect this correlation, then their new correlation neglect (CN) belief, q^{CN}, will be uniquely determined as:

$$q^{CN} = \frac{\prod_{i=1}^{n} q^i}{\prod_{i=1}^{n} q^i + \prod_{i=1}^{n} (1 - q^i)} \qquad 1.$$

Thus, for example, if a share α of N individuals had received the h signal and have belief $q > \frac{1}{2}$, and a share $1 - \alpha$ had received the l signal and have belief $1 - q < \frac{1}{2}$, then if

[17] Neglecting correlation is not necessarily a bias of naïve individuals; scientists and data analysts have long treated forecasts as independent. The Naïve-Bayes classifier, a method to analyse data by assuming different aspects of it are independent, is one of the work horses of operations research and machine learning. This method has had surprising success and is extensively used. Querubin & Dell (2017) document how this approach was employed by the US military in the Vietnam war to assess which hamlets should be bombed based on multidimensional data collected from each hamlet. For more on the Naïve-Bayes approach see Russell & Norivg (2003) and Domingos & Pazzani (1996).

all exchange their beliefs, we have that

$$q^{CN} = \frac{q^{\alpha N}(1-q)^{(1-\alpha)N}}{q^{\alpha N}(1-q)^{(1-\alpha)N} + (1-q)^{\alpha N}(q)^{(1-\alpha)N}},$$

with $q^{CN} \to 1$ for a large N and $\alpha > \frac{1}{2}$, and $q^{CN} \to 0$ for a large N and $\alpha < \frac{1}{2}$. If for example the true information structure that had generated these initial beliefs involves correlation, so that all those that received the same signal had the same information source, then post-communication beliefs would become excessively extreme.

More generally, it is easy to see from (1) that the belief updating function satisfies the following properties. First, confident individuals are very persuasive. For example, if $q^i = 1$ (or alternatively $q^i = 0$) for some i, so that some individual has extreme beliefs, then she fully convinces all others. Second, beliefs are monotone: they increase in peers' beliefs. Finally, belief updating can also exhibit extremism and polarisation: For a set of beliefs where all are higher (lower) than a half, then updated beliefs would be higher (lower) than the maximum (minimum) belief in the set. For example if $q^i > \frac{1}{2}$ for all i, then the correlation neglect belief q^{CN} will satisfy $q^{CN} > \max_i q^i$. If $q^i < \frac{1}{2}$ for all i, then $q^{CN} < \min_i q^i$. Thus observing a selection of similar beliefs will induce extreme beliefs. Moreover, if society segregates into two groups, one with people who have high beliefs and one with those with low beliefs, polarisation will arise.

The prevalence of full correlation neglect in complex environments: The above definition captures individuals who fully neglect the possibility of correlation. Some individuals might be concerned about the "correlation neglect" that is implicit in this Naïve-Bayes approach or simply have misspecified models of the correlation. Ellis & Piccione (2017) provide an axiomatic characterisation of individuals that cannot account for correlation (or complexity in their terminology). Levy & Razin (2018c) propose a model in which individuals neglect correlation to some degree. Specifically, if we think of a modified correlation-neglect belief that allows for some correlation, then Levy & Razin (2018c) show that it can be written as

$$q^{CN}_{\boldsymbol{\lambda}} = \frac{\lambda_h \prod_{i=1}^{n} q^i}{\lambda_h \prod_{i=1}^{n} q^i + \lambda_l \prod_{i=1}^{n} (1-q^i)} = \frac{\lambda_h q^{CN}}{\lambda_h q^{CN} + \lambda_l (1-q^{CN})} \qquad 2.$$

where λ_h, λ_l are parameters that capture degrees of correlation across the information in the different states, high or low. However, when $n \to \infty$, in many environments we would have $q^{CN} \to 1$ or $q^{CN} \to 0$. If in addition λ_h, λ_l are bounded, i.e., $\lambda_h, \lambda_l > 0$, and $\lambda_h, \lambda_l < \infty$, we will have $q^{CN}_{\boldsymbol{\lambda}} \to q^{CN}$. In other words, correlation neglect can arise when we face big data, whose naive interpretation according to the q^{CN} is sufficiently precise (but not necessarily correct). Thus, in complex environments full correlation neglect is likely to arise. Even if individuals consider some degree of correlation, a large data set will overwhelm this and they will behave as if they have full correlation neglect.[18]

[18] Another general model of correlation neglect is provided by Spiegler (2018) using the tool of Bayesian networks. A simple network is a relation between random variables x_1, x_2 and x_3. For example, it may be that x_1 induces x_2 as well as x_3 but x_2 and x_3 are independent. A decision maker can neglect correlation so that she believes that $\Pr(x_1, x_2, x_3) = \Pr(x_1)\Pr(x_2)\Pr(x_3)$ (for applications of this approach, see Eliaz et al. (2018) and Schumacher & Thysen (2018)).

3.3. Selection bias

When the information you are exposed to is not randomly assigned, selection bias might arise. For example, suppose everyone in your vicinity has $q^i > \frac{1}{2}$. On the one hand, you might take this observation to mean that there is large evidence that the state is probably one. But in reality your observation might be a result of the fact that you and the people around you all chose to interact with one another. If the latter is the case, you should decrease the weight you put on the opinions of those close to you. Failure to do this to the right degree is termed selection bias which we now define in the context of our model.

Below we introduce a formulation of selection bias used in Levy & Razin (2017) to model socialisation in schools. In the context of our simple information model, suppose that all those with the high signal h and hence the high beliefs q, a share α among the N individuals, communicate only between themselves, and similarly, all with the low signal l and hence low beliefs $1-q$, a share $1-\alpha$, communicate only with each other. Selection bias arises as individuals do not take this segregation into account. Rather, individuals assume that the opinions they see were generated uniformly from opinions in the population. In these two groups, beliefs involving selection bias (SB), will differ and will depend on the signal l or h:

$$q^{SB}(h) = \frac{q^{\alpha N}}{q^{\alpha N} + (1-q)^{\alpha N}} > q^{SB}(l) = \frac{(1-q)^{(1-\alpha)N}}{(1-q)^{(1-\alpha)N} + (q)^{(1-\alpha)N}} \qquad 3.$$

Frick et al. (2018) use a similar notion of selection bias that they term assortativity neglect, and provide a theoretical foundation for it as a model of misperception in a segregated society.

Note that selection bias is related to correlation neglect and so will give rise to similar dynamics. In particular, within each group we will have a process of extremism in which beliefs become more extreme when individuals exchange information. A result of the different composition of groups will result in polarisation of opinions across groups. To understand exactly the patterns of extremism and polarisation one would have to combine the analysis of endogenous segregation discussed in the last section with the evolution of beliefs modelled here. This will be the topic of Section 4.

3.4. Confirmation bias

While selection bias arises because of our choices of whom to interact with, confirmation bias is a bias that arises from the way we interpret what we see. Confirmation bias refers to the propensity to ignore or misinterpret information that runs counter to one's own belief. One of the first experiments that is associated with confirmation bias is the one by Lord et al. (1979). They show that individuals exposed to the same information can polarise their beliefs in different directions. Thus, information must be interpreted differently.[19]

Using our model we can represent confirmation bias in the following way. In a sense, for confirmation bias to arise, we do not need segregation per se, as the segregation arises cognitively, through the misinterpretation of certain pieces of information. Suppose again that a share α observed the signal h and have high beliefs $q > \frac{1}{2}$ and a share $1-\alpha$ observed

[19]Rabin & Schrag (1999) provide a model of confirmatory bias according to which contrary pieces of evidence are simply viewed as confirming one's beliefs, which may lead to very strong –and wrong– beliefs.

the low signal l and have low beliefs $1-q<\frac{1}{2}$. Suppose the individual who has posterior q, so she had observed the signal h, as in Rabin & Schrag (1999), misperceives low posteriors as high ones with probability $\beta>0$. Thus, given such confirmation bias (CB) and her signal h, she ends up with the following belief:

$$q^{CB}(h)=\frac{(q)^{(\alpha+(1-\alpha)\beta)N}(1-q)^{(1-\alpha)(1-\beta)N}}{(q)^{(\alpha+(1-\alpha)\beta)N}(1-q)^{(1-\alpha)(1-\beta)N}+(1-q)^{(\alpha+(1-\alpha)\beta)N}(q)^{(1-\alpha)(1-\beta)N}}$$

While an individual with posterior $1-q$ who had observed the signal l will interpret high posteriors as low ones with probability β, and hence will end up with beliefs:

$$q^{CB}(l)=\frac{(q)^{\alpha(1-\beta)N}(1-q)^{((1-\alpha)+\alpha\beta)N}}{(q)^{\alpha(1-\beta)N}(1-q)^{((1-\alpha)+\alpha\beta)N}+(1-q)^{\alpha(1-\beta)N}(q)^{((1-\alpha)+\alpha\beta)N}}$$

Again, when N is large enough, we have that $q^{CB}(h)\to 1$ and $q^{CB}(l)\to 0$ when β large enough (compared with α). Thus, confirmation bias is sufficient to create two chambers with polarisation.

Here segregation is not physical but rather created by selective interpretation of information: if when browsing on-line individuals interpret the content in their own way, by way of confirmation bias, then *de facto* they are segregating away from others who interpret information differently. This makes individuals become more convinced in their views and hence creates polarisation.

3.5. Environments that facilitate biases

A question one may want to ask is when should we expect the above biases to arise. For example, when an individual reads news from different outlets, on-line or off-line, if these are truly independent pieces of information, then she still behaves optimally even if she suffers from correlation neglect. Therefore, to understand the relevance of correlation neglect, we need to understand the sources of correlation in our environment. We now discuss several environments in which we expect the above biases to be more prominent.

On-line replication of news: One avenue for correlation neglect to arise is the replication of on-line (as well as off-line) news content. There is a good reason to think that consumers of news media are likely to suffer from correlation neglect to some extent. For one thing, news items are constantly copied and repackaged across outlets. Cagé et al. (2017) study copyright in news media, following pieces of news as they trickle through different outlets including social media. They document how pieces of news are often copied multiple times and across different outlets. In addition they find that only 32% of on-line content is original. Still, despite the prevalence of copying, media outlets hardly name the sources they copy. Thus readers are exposed to repeated news, potentially without being aware of it.

News aggregation websites are another example of how media is copied and the sources of information are made harder to trace. These sites publish their own news as well as links to similar news in other sites, which therefore expose individuals to repetition of news.

Exposure to multiple sources of information: It is also clear that people read multiple sources of information. Individual-level survey data on 18 countries from Reuters Institute for the Study of Journalism shows that the average news consumer uses about five news sources per week. More generally Kennedy & Prat (2017) and Prat (2017) document the consumption patterns of news consumers showing that individuals use multiple outlets

to learn about news.[20] Communication among individuals also implies that, indirectly, they are exposed to even more sources.

Segregation and complexity of communication in networks: As we discussed in the Section above, segregation, be it physical or on-line, is an inherent trait of society. Segregation patterns are very complex and imply that individuals might have multiple social networks they belong to. This complexity implies that it is hard to follow both the selection that is involved in what you are exposed to as well as the correlation structures between the pieces of information you consume.

Repeated communication in groups and more generally in networks is often considered to impose large informational requirements on individuals. Individuals may be unaware of the structure of the network, so that while they know who they communicate with, they might not know their neighbours' neighbours. This implies that it may be very difficult to trace the path that a piece of information takes in an environment with repeated communication.

The network literature has typically taken one of two avenues. One avenue is the fully rational approach whereby individuals are fully aware of the network and the equilibrium and update using Bayes rule (see Acemoglu et al. 2014). The second avenue is to assume that individuals follow a particular heuristic when updating. A leading example is the DeGroot heuristic, where individuals average their and others' beliefs, as in Golub & Jackson (2010) and De Marzo et al. (2003). De Marzo et al. (2003) analyse a model of multiple rounds of communication (in a network) when players have correlation neglect. They show that this implies that views will become concentrated on a one-dimensional conflict. See Jackson (2011) for a survey of social networks and information diffusion in networks.

These are two polar ways to model information diffusion, one based on full rationality and the other based on an ad hoc heuristic. A third avenue, which we explore in Levy & Razin (2018b), is to account for correlation neglect. Note that the DeGroot heuristic does not lead to polarisation of beliefs, as beliefs are averaged; however, using q^{CN} as above leads to polarisation and extreme beliefs.[21]

In the social learning in networks literature, some have identified correlation neglect with a redundancy bias (Gagnon-Bartsch & Rabin 2016), whereas Eyster & Rabin (2010) use a form of neglect of one's action from the information of others in their naïve herding model. Bohren (2016) also considers model misspecification in the context of herding.[22]

Machine learning and AI: On-line browsing has become more and more complex throughout the years. Today, the algorithms used by search engines and other stakeholders on the internet have implications for the creation of on-line echo chambers. Filter bubbles is the term used to describe the propensity of search engines to match individuals with content that would appeal to them. New machine learning and AI algorithms have been shown to amplify existing biases in our society. Recent examples are the experience of Microsoft with Tay, a twitter chat box and the experiment of MIT researchers with an AI algorithm called Norman, showing how very different outcomes result from feeding the application with different information.[23] Moreover, algorithms of news and content aggregation, which

[20]See also "In Changing News Landscape, Even Television is Vulnerable: Trends in News Consumption: 1991-2012" a report by Pew Research.

[21]See also Dandekar et al. (2013) who show that the DeGroot model does not yield polarisation, and Molavi et al. (2018) who model other forms of non-Bayesian social learning in networks.

[22]See also Guarino & Jehiel (2013) and Mueller–Frank & Neri (2013).

[23]See O'Brien (2018).

are based on complex, non-transparent algorithms, muddy the waters in terms of our understanding of the correlation structures behind the multitude of pieces of information we are exposed to.

Bakshy et al. (2015) analyse how on-line networks influence exposure to perspectives that cut across ideological lines. They examined how 10.1 million U.S. Facebook users interact with socially shared news. They directly measured ideological homophily in friend networks and examined the extent to which heterogeneous friends could potentially expose individuals to cross-cutting content. They then quantified the extent to which individuals encounter comparatively more or less diverse content while interacting via Facebook's algorithmically ranked News Feed and further studied users' choices to click through to ideologically discordant content. They show that both the algorithmic ranking and to a larger degree individuals' choices played a role in limiting exposure to cross-cutting content.

Concentration of ownership implies correlation: The intervention of owners in the editorial decisions of their news outlets has always been an important issue in the debate about the regulation of the media industry.[24] It is one of the reasons behind a common call to have independent editorial boards. For example, in the UK, in June 2017, the culture secretary decided to refer 21st Century Fox's £11.7bn bid to seize full control of satellite broadcaster Sky to the Competition and Markets Authority, for a fuller, "phase two" investigation. The FT reports that behind this decision was the fact that "While Fox and News Corp are separate companies, the Murdoch Family Trust has material influence across both companies." To secure the deal, 21st Century Fox has to take some measures that "[...] include setting up a separate editorial board with a majority of independent members to oversee Sky News and a commitment to maintain Sky-branded news for five years at current funding levels."[25]

3.6. Strategic manipulation of cognitive biases

The existence of the cognitive biases we surveyed above opens the door for interested parties to take advantage of consumers or voters. A recent example is the use of Facebook by Cambridge Anlaytica and partner organisations to affect different political campaigns across the world. One way in which the Facebook data was used was to have targeted messages, tailored to the characteristics of users. In addition, Cambridge analytica allegedly shared its data with other organisations working on the same campaigns to create repeated messaging to the same individuals under different frames. Another example is the use of real-time information about which messages were resonating to shape Donald Trump's travel schedule during the 2016 election campaign. So, if there was a spike in clicks on an article about immigration in a county in Pennsylvania or Wisconsin, then Trump would visit the place and deliver an immigration-focused speech.[26]

[24] Some insight into the nature of such interventions can be gleaned from the evidence given to the Leveson Inquiry in the UK. From Hickman (2012): "Andrew Neil, who edited The Sunday Times between 1983 and 1994, recalled in Full Disclosure that although the proprietor did not expect to see his views repeated immediately in the next paper "he had a quiet, remorseless, sometimes threatening way of laying down the parameters within which you were expected to operate... stray too far too often from his general outlook and you will be looking for a new job." The former Times and Sunday Times editor Harold Evans said that Murdoch broke all of his promises of editorial independence after taking over titles."

[25] See Bond (2017).

[26] See Illing (2018).

A recent literature in economics and political science sheds light on strategic influence in the presence of cognitive biases. Levy et al. (2018a,b) analyse how interested parties can influence an individual who has correlation neglect and apply their results to the media market. Giovanniello (2018) analyses a model of informative campaign advertising, and shows how the ability of voters to strategically communicate with each other shapes the advertising strategies of two competing parties. Mullainathan & Shleifer (2005) analyse the news market when individuals like to read news that agree with their views, or confirm their bias. They show how this leads firms to slant their news reports in the direction of such bias. Prat (2017) develops a measure of media power which is based on fully impressionable readers with correlation neglect.

3.7. Are cognitive biases (and polarisation) necessarily harmful?

The above discussion has shown evidence for the existence and prevalence of cognitive biases in acquiring information. These biases will lead to individuals holding wrong and biased beliefs. But what are the costs of having such wrong, and sometimes polarised, beliefs?

While intuitively we might think that cognitive biases are bad for voters, a recent literature in behavioural political economy shows that sometimes these biases might also have some positive impact on aggregate welfare. Levy & Razin (2015a) analyse a voting model with heterogeneous voters and a common value shock. All voters prefer the policy on the right in a right-wing state of the world, and the policy on the left in a left-wing state of the world, albeit with different intensities. Each voter receives signals about the state of the world and makes voting decisions given this information and her preferences. Signals are correlated but "behavioural" voters neglect the correlation in these sources, while rational voters do not. The key result in this paper is that correlation neglect can be – and is, in many standard environments – beneficial for information aggregation: Even if each behavioural voter does not vote optimally from her own point of view (compared to a rational voter), the whole electorate may reach better, more informed, outcomes (compared to a rational electorate). Intuitively, correlation neglect magnifies the effect of information on individuals' behaviour. Individuals who might otherwise stick with the policy that accords with the direction of their political preferences may be swayed to change their vote if they believe that their information is sufficiently strong in the opposite direction. This implies that individuals base their vote more on their information rather than on their preferences. Thus, while correlation neglect is harmful for individuals, it may be better for society on average. Levy & Razin (2015b) show, in the context of political polarisation, that polarisation in voter opinions that is due to correlation neglect does not necessarily translate to polarisation in political platforms of parties.

Lockwood (2018) shows the implications of confirmation bias, in a political agency setting. In his paper, as opposed to the rest of the literature that focuses on behavioural voters, either voters or politicians can have this bias. In the baseline case where voters have this bias and where only the politician's actions are observable before the election, confirmation bias decreases pandering by the incumbent, and can raise voter welfare as a consequence. Similarly, Ashworth & Bueno de Mesquita (2014) show that voter incompetence, modelled as the voter's lack of ability to be properly informed, can sometimes improve politicians incentives to choose the right policies (due to a reduced signalling motivation).

The key idea in the literature surveyed above is that the political system, even without taking into consideration cognitive biases, is already flawed. It sometimes blocks information

from being aggregated efficiently, or its electoral incentives induce politicians to behave not in line with voters' preferences. Starting from an inefficient benchmark, it is sometimes useful for voters to be overconfident or for voters to ignore in some way their information and therefore induce less distortive behaviour by politicians.

But different cognitive biases might imply very different normative results. Levy & Razin (2015a) show that when voters have confirmation bias the election aggregates less information than when voters have correlation neglect. Therefore, it is important to understand empirically what is the underlying cognitive bias that voters have.

3.8. Other biases

We have considered models in which individuals are restricted from updating information properly which implies that echo chambers can arise. This failure of belief updating arises for example when the environment is too complex to understand (e.g., networks of communication), or, when individuals face some cognitive constraints. Other models in the literature have instead analyzed how individuals may be compelled to manipulate their own beliefs in order to affect their behaviour. For example, if individuals believe that hard work induces high rewards, they know they will work harder; in turn they may be motivated to influence their beliefs in this direction. This "motivated beliefs" incentive, explored for example in Bénabou & Tirole (2006) and Bénabou & Tirole (2011), can also then create clusters of individuals with similar beliefs. For example, Bénabou (2013) considers how complementarities in group activities compel individuals to manipulate beliefs in the same way.

4. The dynamics of echo chambers

The above two sections have surveyed the literatures on segregation (chambers) and cognitive biases (echoes). In this section we analyse the feedback effects between segregation decisions and the effects of segregation on beliefs. Intuitively, our perceptions about the world are shaped, in part, by where we live and who we interact with. But our decisions about where to live and who to talk to are also shaped by our beliefs. Therefore, to fully understand the implications of echo chambers one has to understand how they evolve. This feedback effect is also important for empirical work; if we fail to take it into account we might make wrong inferences about causality. For example, Dustmann & Preston (2001) analyse how segregation in neighbourhoods affects attitudes towards minorities. They show that earlier studies that have only looked at one direction of causality, i.e., how segregation and social exclusion affect beliefs and attitudes towards minorities, have biased results due to neglecting location choices which depend on these beliefs.

To illustrate the feedback effect between segregation and beliefs we focus on an example of schooling. Levy & Razin (2017) analyse how echo chambers in schools can sustain polarised beliefs that imply labour market discrimination. The model describes a society with non-overlapping generations, infinite periods, and three stages in each period. In the peer influence (Echo) stage, segregation affects beliefs. In this stage, individuals' beliefs about schools are shaped by their parents' beliefs and by their school peers, where they ignore selection bias. In the labour market stage, discrimination may arise based on such beliefs. Employers decide whether to hire an employee, based on the school she graduated from and their own beliefs about the schools' effect on productivities. Labour market

experience also entails learning about true productivities. In the school choice (Chambers) stage, beliefs and labour market discrimination affect segregation choices. In this stage, parents choose which school –state or private– to send their offspring to. Thus, the model puts forward explicitly the feedback effect between echoes and chambers.

The model uses "imperfect empathy" in parental school choice, as in Bisin & Verdier (2001). Parents base their decisions on their expectation about how their children will fare in the labour market. But their child's labour market experience will be shaped both by what others will think of her as well as her own beliefs. Therefore, parents have to form expectations about how the school will affect their child's future beliefs and behaviour. The "imperfect empathy" assumption means that parents evaluate their child's welfare using their own belief, not the one that their child will end up holding. This creates homophily, that is, parents would rather their children segregate with like-minded others so that their child's belief does not stray too far from their own. This endogenous homophily, along with selection bias, will imply that beliefs can become polarised.

Levy & Razin (2017) find a simple necessary and sufficient condition that characterises when segregation, polarised beliefs and discrimination persist in the long run. When the condition is satisfied, in all equilibria, there are polarised beliefs about the productivity of graduates from the different schools (over and above actual productivity differences). Parents who send their children to a private school believe that the difference between the schools is greater than it really is. Parents who send their children to a state school realise that there is discrimination, believe it is not justified, and are priced out of private school. Finally, those who went to private (state) school will also send their children to a private (state) school. Thus, the "old boys" network is endogenously formed.

The analysis centres on the race between echo chamber effects and true learning.[27] First, history matters; to create long-run segregation and polarised beliefs, those in the private school have to start from a relatively low opinion of state school graduates. Second, the higher the intensity of socialisation in schools is, the easier it is to create segregation and polarisation. Finally, polarised beliefs are easier to sustain the less individuals learn about others from their labour market experience. Importantly, the cycle of segregation and polarised beliefs can also be broken down. This arises in the model when those who segregate into the private school have sufficiently mixed beliefs so that belief polarisation cannot arise.

In the dynamic model above, the school choices of parents affect the beliefs of their children, and these in turn affect their schooling choices when they are parents themselves. A few papers have taken an alternative approach to model this feedback effect in a static model. Frick et al. (2018) analyse a model in which individuals segregate into different interaction groups but could hold misperceived beliefs about what happens in other groups. Their equilibrium notion, termed local perception equilibrium, has an observational consistency requirement so that their perception about those they interact with must be correct. They then move on to show that misperceptions similar to our notion of selection bias above have the property that they are part of an equilibrium no matter the environment. Similarly, Windsteiger (2018), who analyses segregation in a political economy model, suggests a notion of equilibrium that also demands that beliefs about one's interaction group are

[27]Wrong beliefs arise in this model not because individuals stop "experimenting", as in Piketty (1995) or Fudenberg & Levine (1993), but because their peers' beliefs are pessimistic enough and thus the echo chamber effect outweighs any positive learning.

always correct. Moreover, Windsteiger (2018) adds an additional requirement about the misperceived beliefs about other groups. She assumes that beliefs must be consistent with those in neighbouring groups not wanting to switch groups. She shows that this additional restriction refines the set of equilibria in a useful way.

5. Future research

In this Section we conclude our survey by pointing at the potential avenues for future research stemming from the discussion above. We consider relevant issues for empirical as well as experimental and theoretical work.

A central empirical challenge at the heart of studying echo chambers is causality. As we saw above, there are feedback effects between the formation of chambers and the kind of beliefs they instil in their occupiers. How can we disentangle whether individuals in segregated neighbourhoods have polarised beliefs due to self-selection or due to a different process of belief formation once segregation arises?

Even when we focus on analysing how beliefs evolve in a chamber following segregation, empirical challenges remain. Specifically, consider the case of prejudice against immigrants or foreigners. Contact theory focuses on interactions between individuals as the vehicle by which stereotyping and prejudice can be reduced (Allport 1954; Hewstone & Brown 1986). According to this theory, individuals who interact with other groups start using information gleaned from personal experiences rather than stereotypes. Pettigrew & Tropp (2006) show how interaction between different groups can substantially reduce attitudinal and behavioural measures of negative evaluation.[28]

Whether contact is helpful however, may depend on the specific interactions between individuals. One needs more data about the nature of interactions between the groups (see Cantle 2001). For example, data about residential segregation might not be enough. We might want to gather data about the distribution of interactions between different groups. Is group A interacting with group B mainly as employers versus employees or are they engaged in more cooperative interactions? A recent study by Lowe (2018), shows how different types of integration, collaborative and adversarial, may have different effects. Lowe recruited 1,261 participants and randomly assigned these young Indian men from different castes to participate in month-long cricket leagues; he shows that collaborative contact reduces discrimination, leading to more cross-caste friendships and 33% less own-caste favouritism while adversarial contact generally has no, or even harmful, effects.

For policy making, it is important to understand the sources and mechanisms that drive echo chambers. For example, are on-line echo chambers supply or demand driven? Papers such as Bakshy et al. (2015) point to a demand driven effects, showing that individuals' choices played an important role in limiting exposure to cross-cutting content. As we saw above, understanding the types of biases that drive these effects is also important as potential remedies depend on the particular biases. More research along these lines is needed to inform our strategies for tackling echo chambers and their effects. To this end, experiments can provide a valuable way to understand under what environments behavioural biases of information processing can be mitigated. A recent paper by Enke (2017) provides

[28] A recent theoretical contribution to contact theory is by Desmet et al. (2018), who develop a measure of antagonism that relies on a "local learning multiplier", a measure of how local interaction affects antagonism towards other groups in society at large.

results about how selection biases can be mitigated. Similarly, Laudenbach et al. (2017) conduct experiments showing how correlation neglect can be overcome depending on the context, framing effects, and the complexity of the problem.

From a theoretical point of view, there are a few methodological issues to consider as well as policy implications. Methodologically, equilibrium analysis when individuals have wrong beliefs has to be adapted. Recent advances provide equilibrium notions for misperception or misspecified models. These include behavioural equilibrium notions such as Cursed Equilibrium (Eyster & Rabin 2005) and Analogy-Based Equilibrium (Jehiel 2005). Esponda & Pouzo (2016) provide a solution concept for games with players who have misspecified models of the world (Berk-Nash equilibrium).

An important extension of current research is the role for government intervention. There are several ways in which governments can intervene to improve outcomes. Some government interventions can be targeted at preventing echoes, and some at preventing segregation. To prevent echoes, one option is public campaigns to inform and correct individuals' wrong beliefs as well as reducing polarisation. A second, less direct, role for governments in the context of echo chambers is regulation of media markets. Concentration of media ownership can allow for strategic manipulation of correlation neglect, as Levy et al. (2018a) show, which provides another reason for the break-up of media conglomerates. To prevent on-line segregation, targeted algorithms facilitate provision of information that is already aligned with individuals' views rather than unbiased information, and these may be regulated. For example, ensuring that algorithms are transparent is somewhat in line with recent EU regulation and will help clarify how information is generated and targeted. Finally, to prevent physical segregation, government may target the role that the private market plays at creating opportunities for segregation, such as segregated schools or neighbourhoods.

DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 681579. For excellent research assis-tance and helpful comments we thank Clement Minaudier and Heidi Thysen. We also thank Lisa Windsteiger for helpful comments.

LITERATURE CITED

Abdulkadiroğlu A, Angrist J, Pathak P. 2014. The elite illusion: achievement effects at Boston and New York exam schools. *Econometrica* 82:137–96

Acemoglu D, Bimpikis K, Ozdaglar, A. 2014. Dynamics of information exchange in endogenous social networks. *Theor. Econ.* 9(1):41–97

Achen C, Bartels L. 2004. *Why shark attacks are bad for democracy.* Unpublished manuscript, Vanderbilt Univ.

Advani A, Reich B. 2015. Melting pot or salad bowl: The formation of heterogeneous communities. Work. Pap. 15/30, IFS

Akerlof GA, Shiller R. 2009. *Animal spirits: How human psychology drives the economy, and why it matters for global capitalism.* Princeton: Princeton University Press

Alesina A, La Ferrara E. 2005. Preferences for redistribution in the land of opportunities. *J. Public Econ.* 89(5-6):897–931

Alesina A, Miano A, Stantcheva S. 2018. Immigration and Redistribution. Work. Pap. 24733, NBER

Algan Y, Do Q-A, Dalvit N, Le Chapelain A, Zenou Y. 2015. *How social networks shape our beliefs: a natural experiment among future French politicians.* Unpublished manuscript, Sciences Po

Allport G. 1954. *The Nature of Prejudice.* Cambridge, MA: Addison-Wesley

Ashworth S, De Mesquita EB. 2014. Is voter competence good for voters?: Information, rationality, and democratic performance. *Am. Polit. Sci. Rev.* 108(3):565–87

Baccara M, Yariv L. 2016. Choosing peers: homophily and polarization in groups. *J. Econ. Theory* 165:152–78

Bai J, Golosov M, Qian N, Kai Y. 2015. *Understanding the influence of government controlled media: evidence from air pollution in China.* Unpublished manuscript, Kellogg School of Management

Bakshy E, Messing S, Adamic LA. 2015. Exposure to ideologically diverse news and opinion on Facebook *Science* 348(6239):1130-32.

Bala V, Goyal S. 1998. Learning from neighbours. *Rev. Econ. Stud.* 65:595–621

Banerjee A. 1992. A simple model of herd behavior. *Q. J. Econ.* 107(3):797–817

Barber M, McCarty N. 2015. Causes and consequences of polarization. In *Solutions to Political Polarization in America* ed. N. Persily, pp. 15–58. Cambridge: Cambridge University Press

Bartels L. 1996. Uninformed votes: information effects in presidential elections. *Am. J. Polit. Sci.* 40(1):194–230

Bartels L. 2008. *Unequal Democracy: The Political Economy of the New Gilded Age.* Princeton: Princeton University Press

Bénabou R. 1996. Heterogeneity, stratification, and growth: macroeconomic implications of community structure and school finance *Am. Econ. Rev.* 86(3):584–609

Bénabou R, Tirole J. 2006. Belief in a just world and redistributive politics. *Q. J. Econ.* 121(2):699–746

Bénabou R, Tirole J. 2011. Identity, morals, and taboos: beliefs as assets. *Q. J. Econ.* 126(2):805–55

Bénabou R. 2013. Groupthink: collective delusions in organizations and markets. *Rev. Econ. Stud.* 80(2):429–62

Bénabou R, Tirole J. 2016. Mindful economics: the production, consumption, and value of beliefs. *J. Econ. Perspect.* 30(3):141–64

Berman E. 2000. Sect, Subsidy, and Sacrifice: An economist's view of ultra-Orthodox Jews. *Q. J. Econ.* 115(3):905—53

Bessi A. 2016. Personality traits and echo chambers on Facebook. *Comput. Hum. Behav.* 65:319–24

Bikhchandani S, Hirshleifer D, Welch I. 1992. A Theory of fads, fashion, custom, and cultural change as informational cascades. *J. Polit. Econ.* 100(5):992–1026

Bishop B. 2009. *The Big Sort: Why the Clustering of America is Tearing Us Apart.* NY: Mariner Books

Bisin A, Verdier T. 2001. The economics of cultural transmission and the dynamics of preferences. *J. Econ. Theory* 97(2):298–319

Bohren JA. 2016. Informational herding with model misspecification. *J. Econ. Theory* 163:222–47

Bohren JA, Imas A, Rosenberg M. 2017. *The dynamics of discrimination: theory and evidence.* Unpublished manuscript, University of Pennsylvania

Boisjoly J, Duncan G, Kremer M, Levy D, Eccles J. 2006. Empathy or antipathy? The impact of diversity. *Am. Econ. Rev.* 96(5):1890–1905

Bond D. 2017. How Fox's bid for Sky came under scrutiny. *The Financial Times*, Jun. 29, https://www.ft.com/content/6dd32342-5ccd-11e7-b553-e2df1b0c3220?mhq5j=e3

Boxell L, Gentzkow M, Shapiro JM. 2017. Greater Internet use is not associated with faster growth

in political polarization among US demographic groups. *P. Natl. Acad. Sci. USA* 114(40):10612–17

Bradford D, Kelejian H. 1973. An econometric model of the flight to the suburbs. *J. Polit. Econ.* 81(3):566–89

Brimicombe AJ. 2007. Ethnicity, Religion, and Residential Segregation in London: Evidence from a Computational Typology of Minority Communities. *Environ. Plann. B* 34(5):884–904

Burns J, Corno L, La Ferrara E. 2016. *Interaction, stereotypes and performance. Evidence from South Africa.* Unpublished manuscript, University of Cape Town

Cagé J, Hervé N, Viaud ML. 2017. The production of information in an online world: Is copy right? Work. Pap. 12066, CEPR

Calvó-Armengol A, Jackson M. 2004. The effects of social networks on employment and inequality. *Am. Econ. Rev.* 94(3):426–54

Campbell A, Converse PE, Miller WE, Stokes DE. 1960. *The American Voter.* New York: Wiley & Sons

Cantle T. 2001. *Community Cohesion: A Report of the Independent Review Team.* London: Home Office.

Cantle T, Kauffman E. 2016. Is segregation increasing in the UK? *Open Democracy.* https://www.opendemocracy.net/wfd/ted-cantle-and-eric-kaufmann/is-segregation-on-increase-in-uk

Chaudhuri S, Sethi R. 2008. Statistical discrimination with peer effects: can integration eliminate negative stereotypes? *Rev. Econ. Stud.* 75(2):579–96

Chetty R, Hendren N, Kline P, Saez E. 2014. Where is the land of opportunity? The geography of intergenerational mobility in the United States. *Q. J. Econ.* 129(4):1553–1623

Crawford VP, Sobel J. 1982. Strategic information transmission. *Econometrica* 50(6):1431–51

Curtis S, Warner J. 1992. Matchmaker, matchmaker: The effect of old boy networks on job match quality, earnings, and tenure. *J. Labor Econ.* 10:306–30

Dandekar P, Goel A, Lee D. 2013. Biased assimilation, homophily, and the dynamics of polarization. *P. Natl. Acad. Sci. USA* 110:5791–96

De Marzo PM, Vayanos D, Zwiebel J. 2003. Persuasion bias, social influence and unidimensional opinions. *Q. J. Econ.* 118:909–68

Del Vicario M, Bessi A, Zollo F, Petroni F, Scala A, Caldarelli G, Stanley HE, Quattrociocchi W. 2016. The spreading of misinformation online. *P. Natl. Acad. Sci. USA* 113(3):554–59.

Delli Carpini MX, Keeter S. 1996. *What Americans Know about Politics and Why It Matters.* New Haven; London: Yale University Press

Desmet K, Gomes J, Ortuño–Ortín I. 2018. The geography of linguistic diversity and the provision of public goods. Work. Pap. 24694, NBER

Dubois E, Blank G. 2018. The echo chamber is overstated: the moderating effect of political interest and diverse media. *Communication & Society* 21(5):729–45

Durlauf SN, Seshadri A. 2017. Understanding the Great Gatsby curve. In *NBER Macroeconomics Annual 2017*, Vol. 32, ed. M Eichenbaum, JA Parker, pp. 333–93. Chicago: University of Chicago Press

Dustmann C, Preston I. 2007. Racial and economic factors in attitudes to immigration. *BE J. Econ Anal. Poli.* 7(1)

Dustmann C, Preston I. 2001. Attitudes to ethnic minorities, ethnic context and location decisions. *Econ. J.* 111:353–73

Dobbie W, Fryer RG. 2014. The impact of attending a school with high-achieving peers: evidence from the New York City exam schools. *Am. Econ. J.-Appl. Econ.* 6(3):58–75

Domingos P, Pazzani MJ. 1996. Beyond independence: conditions for the optimality of the simple Bayesian classifier. *Machine Learning: Proceedings of the Thirteenth International Conference* 105–12

Eliaz K, Spiegler R, Thysen H. 2018. *Strategic interpretations.* Unpublished manuscript, Tel Aviv

Univ.
Ellis A, Piccione M. 2017. Correlation misperception in choice. *Am. Econ. Rev.* 107(4):1264–92
Enke B. 2017. *What you see is all there is.* Unpublished manuscript, Harvard Univ.
Enke B, Zimmerman F. 2017. Correlation neglect in belief formation. *Rev. Econ. Stud.*. In print.
Esponda I, Pouzo D. 2016. Berk—Nash equilibrium: A framework for modeling agents with misspecified models. *Econometrica* 84(3):1093–1130
Evans G, Tilley, J. 2012. How parties shape class politics: explaining the decline of the class basis of party support. *Brit. J. Polit. Sci.* 42(1):137–61
Eyster E, Rabin M. 2005. Cursed equilibrium. *Econometrica* 73(5):1623–72
Eyster E, Rabin M. 2010. Naïve herding in rich-information settings. *Am. Econ. J.-Microecon.* 2(4):221–43
Eyster E, Weizsäcker G. 2011. Correlation neglect in financial decision-making. Work. Pap. 1104, DIW Berlin
Field E, Levinson M, Pande R, Visaria S. 2008. Segregation, rent control, and riots: The economics of religious conflict in an Indian city. *Am. Econ. Rev.* 98(2):505–10
Frick M, Iijima R, Ishii Y. 2018. *Dispersed behavior and perceptions in assortative societies.* Work. Pap. 2128, Cowles Foundation
Fudenberg D, Levine D. 1993. Self confirming equilibrium. *Econometrica* 61(3):523–45
Gagnon-Bartsch T, Rabin M. 2016. *Naive social learning, mislearning, and unlearning.* Unpublished manuscript, Harvard Univ.
Galeotti A, Ghiglino C, Squintani F. 2013. Strategic information transmission networks. *J. Econ. Theory* 148(5):1751–69
García-Alonso MDC, Wahhajz Z. 2018 *Social diversity and bridging identity.* Unpublished manuscript, Kent Univ.
Gentzkow M, Shapiro JM. 2011. Ideological segregation online and offline. *Q. J. Econ.* 126:1799–1839
Gentzkow M, Shapiro JM, Sinkinson M. 2014. Competition and ideological diversity: historical evidence from US newspapers *Am. Econ. Rev.* 104(10):3073–3114
Gilbert E, Bergstrom T. 2009. Blogs are echo chambers. In *42nd Hawaii International Conference on System Sciences*, pp. 1–10.
Gilens M. 2012. *Affluence and Influence: Economic Inequality and Political Power in America.* Princeton: Princeton University Press
Giovanniello MA. 2018. *Echo chambers: voter-to-voter communication and political competition.* Unpublished manuscript, University of Verona
Golub B, Jackson MO. 2010. Naïve learning in social networks and the wisdom of crowds. *Am. Econ. J.-Microecon.* 2(1):112–49
Golub B, Jackson MO. 2012. How homophily affects the speed of learning and best-response dynamics. *Q. J. Econ.* 127(3):1287–1338
Goyal S, Hernández P, Martínez–Cánovas G, Moisan F, Muñoz–Herrera M, Sánchez A. 2017. *Integration and segregation.* Unpublished manuscript, Cambridge Univ.
Guarino A, Jehiel P. 2013. Social learning with coarse inference. *Am. Econ. J.-Microecon.* 5(1):147–74
Healy AJ, Malhotra N, Mo CH. 2010. Irrelevant events affect voters' evaluations of government performance. *P. Natl. Acad. Sci. USA* 107(29):12804–09
Hewstone M. 2009. Living apart, living together? The role of intergroup contact in social integration. *P. Brit. Acad.* 162:243–300.
Hewstone M, Brown R. 1986. Contact is not enough: an intergroup perspective on the 'Contact Hypothesis'. In *Contact and Conflict in Intergroup Encounters* ed. M Hewstone, R Brown, pp. 1–44. Cambridge, MA: Basil Blackwell
Hickman M. 2012. The world according to Rupert Murdoch. *The Independent*, Apr. 26, http://www.independent.co.uk/news/uk/crime/the-world-according-to-rupert-murdoch-

7679254.html
Huber GA, Hill SJ, Lenz GS. 2012. Sources of bias in retrospective decision-making: experimental evidence on voters limitations in controlling incumbents. *Am. Polit. Sci. Rev.* 106(4):720–41
Iannaccone LR. 1992. Religious markets and the economics of religion. *Soc. Compass* 39(1):123–31
Illing S. 2018. Cambridge Analytica, the shady data firm that might be a key Trump-Russia link, explained. *Vox*, Apr. 4, https://www.vox.com/policy-and-politics/2017/10/16/15657512/cambridge-analytica-facebook-alexander-nix-christopher-wylie
Jackson MO. 2011. An overview of social networks and economic applications. In *The Handbook of Social Economics*, ed. J Benhabib, A Bisin, MO Jackson pp. 511-85. Amsterdam: North Holland Press
Jehiel P. 2005. Analogy-based expectation equilibrium. *J. Econ. theory* 123(2):81–104
Jiao P, Veiga A, Walther A. 2016. *Social media, news media and the stock market.* Unpublished manuscript, Oxford Univ.
Kallir I, Sonsino D. 2009. The neglect of correlation in allocation decisions. *South. Econ. J.* 75(4):1045–66
Kaufmann E, Harris G. 2015. White flight or positive contact?: local diversity and attitudes to immigration in Britain. *Comp. Polit. Stud.* 48(12):1563–90
Kennedy P, Prat A. 2017. *Where do people get their news?* Unpublished manuscript, Columbia Univ.
Kinder DR, Sears DO. 1985. Public opinion and political action. In *The Handbook of Social Psychology*, 3rd edition, ed. G Lindzey, E Aronson, pp. 659-742. New York: Random House
Lau R, Redlawsk DP. 2001. Advantages and disadvantages of cognitive heuristics in political decision making. *Am. J. Polit. Sci.* 45(4):951–71
Laudenbach C, Ungeheuer M, Weber M. 2017. *How to overcome correlation neglect?* Unpublished manuscript, Goethe University Frankfurt
Leigh A. 2009. Does the world economy swing national elections? *Oxford B. Econ. Stat.* 71(2):163–81
Levy G, Razin R. 2012. Religious beliefs, religious participation and cooperation. *Am. Econ. J.-Microecon.* 4(3):121–51
Levy G, Razin R. 2015a. Correlation neglect, voting behaviour and information aggregation. *Am. Econ. Rev.* 105(4):1634–45
Levy G, Razin R. 2015b. Does polarisation of opinions lead to polarisation of platforms? The case of correlation neglect. *Q. J. Polit. Sci.* 10(3):321-55
Levy G, Razin R. 2017. The coevolution of segregation, polarised beliefs and discrimination: the case of private vs. state education. *Am. Econ. J.-Microecon.* 9(4):141–70
Levy G, Razin R. 2018a. *Immigration into prejudiced societies: the effects of segregation and echo chambers.* Unpublished manuscript, LSE
Levy G, Razin R. 2018b. Information diffusion in networks with the Bayesian Peer Influence Heuristic. *Game. Econ. Behav.* 109:262–70
Levy G, Razin R. 2018c. *Combining forecasts in the presence of ambiguity over correlation structures.* Unpublished manuscript, LSE
Levy G, Moreno De Barreda I, Razin R. 2018a. *Media power via the correlation of news content.* Unpublished manuscript, LSE
Levy G, Moreno De Barreda I, Razin R. 2018b. *Persuasion with correlation neglect.* Unpublished manuscript, LSE
Le Yaouanq Y. 2018. *A model of ideological thinking.* Unpublished manuscript, LMU
Lockwood B. 2018. Confirmation bias and electoral accountability. *Q. J. Polit. Sci.*. In press.
Lord CG, Ross L, Lepper MR. 1979. Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *J. Pers. Soc. Psychol.* 37(11):2098–2109.
Lowe M. 2018. *Types of contact: a field experiment on collaborative and adversarial caste integration.* Unpublished manuscript, MIT

McCarty N, Poole KT, Rosenthal H. 2008. *Polarized America: The Dance of Ideology and Unequal Riches.* Cambridge, MA: MIT Press

McPherson M, Smith-Lovin L, Cook JM. 2001. Birds of a feather: homophily in social networks. *Annu. Rev. Sociol.* 27(1):415–44

Mitchell A, Gottfried J, Barthel M, Shearer E. 2016. *The Modern News Consumer.* Pew Research Center, http://www.journalism.org/2016/07/07/pathways-to-news/

Molavi P, Tahbaz-Salehi A, Jadbabaie A. 2018. A theory of non-Bayesian social learning. *Econometrica* 86(2):445–90

Mueller–Frank M, Neri C. 2013. *Social learning in networks: theory and experiments.* Unpublished manuscript, University of Navarra

Mullainathan S, Shleifer A. 2005. The market for news. *Am. Econo. Rev.* 95(4):1031–53

Nasiroglu L. 2016. State school children believe privately educated pupils are more successful. *YouGov*, Feb. 2, https://yougov.co.uk/news/2016/02/02/state-school-children-believe-privately-educated-p/

O'Brien SA. 2018. MIT scientists created an AI-powered 'psychopath' named Norman. *CNN*, Jun. 7, https://money.cnn.com/2018/06/07/technology/mit-media-lab-normal-ai/index.html

Ortoleva P, Snowberg E. 2015. Overconfidence in political behavior. *Am. Econ. Rev.* 105(2):504–35

Peski M. 2008. *Complementarities, group formation and preferences for similarity.* Unpublished manuscript, University of Toronto

Pettigrew T, Tropp L. 2006. A meta-analytic test of intergroup contact theory. *J. Pers. Soc. Psychol.* 90(5):751–83

Piketty T. 1995. Social mobility and redistributive politics. *Q. J. Econ.* 110(3):551–83

Prat A. 2017. Media power. *J. Polit. Econ.*. In press.

Putnam RD. 2007. E pluribus unum: diversity and community in the twenty-first century. *Scand. Polit. Stud.* 30(2):137–74

Quattrociocchi W, Scala A, Sunstein CR. 2016. *Echo chambers on Facebook* Unpublished manuscript, IMT Lucca

Querubin P, Dell M. 2017. Nation building through foreign intervention: evidence from discontinuities in military strategies. *Q. J. Econ.* 133(2):701—64

Rabin M. 1998. Psychology and economics. *J. Econ. Lit.* 36(1):11–46

Rabin M, Schrag J. 1999. First impressions matter: a model of confirmatory bias. *Q. J. Econ.* 114(1):37–82.

Razin A. 2018. *Israel and the World Economy: the Power of Globalization* Cambridge, MA: MIT Press

Reardon S, Bischoff K. 2011. *More Unequal and More Separate: Growth in the Residential Segregation of Families by Income, 1970-2009.* Report prepared for the Russell Sage US2010 project

Russell S, Norvig P. 2003. *Artificial Intelligence: A Modern Approach* Upper Saddle River, NJ: Prentice Hall. 2nd ed.

Schumacher H, Thysen H. 2018. *Equilibrium contracts and boundedly rational expectations.* Unpublished manuscript, KU Leuven

Spiegler R. 2018. *Can agents with causal misperceptions be systematically fooled?* Unpublished manuscript, Tel Aviv Univ.

Sturgis P, Brunton-Smith I, Read S, Allum N. 2011. Does ethnic diversity erode trust? Putnam's 'Hunkering-Down' thesis reconsidered. *Brit. J. Polit. Sci.* 41(1):57–82

Sunstein CR. 2001. *Echo chambers: Bush vs. Gore, Impeachment, and Beyond* Princeton: Princeton University Press

Tiebout CM. 1956. A pure theory of local expenditures. *J. Polit. Econ.* 64(5):416–24

Uslaner E. 2012. *Segregation and Mistrust: Diversity, Isolation, and Social Cohesion.* Cambridge: Cambridge University Press

Vertier P, Viskanic M. 2018. *Dismantling the 'Jungle': migrant relocation and extreme voting in France.* Unpublished manuscript, Sciences Po

Windsteiger L. 2018. *The redistributive consequences of segregation.* Work. Pap. 2017-12, Max Planck Institute for Tax Law and Public Finance

Wolfers J. 2007. *Are voters rational? Evidence from gubernatorial elections.* Unpublished manuscript, University of Pennsylvania