

Keeping up with “The Joneses”: reference dependent choice with social comparisons

Alastair Langtry*[†]

June 28, 2022

Abstract

Keeping up with “The Joneses” matters. This paper examines a model of reference dependent choice where reference points are determined by social comparisons. An increase in the strength of social comparisons, even by only a few agents, increases consumption and decreases welfare for everyone. Strikingly, a higher marginal cost of consumption can increase welfare. In a labour market, social comparisons with co-workers create a *big fish in a small pond* effect, inducing incomplete labour market sorting. Further, it is the skilled workers with the weakest social networks who are induced to give up income to become the *big fish*.

JEL: D60, D85, D91, J24

Keywords: social comparisons, reference dependence, networks, labour market sorting.

*University of Cambridge and King’s College, Cambridge CB2 1ST, UK. Email: atl27@cam.ac.uk.

[†]I am grateful to Matthew Elliott, Lukas Freund, Edoardo Gallo, Charles Parry and Gabriela Stockler for insightful comments and discussions, and to seminar participants at the University of Cambridge. This work was supported by the Economic and Social Research Council [award reference ES/P000738/1]. Any remaining errors are the sole responsibility of the author.

“No man is an island, entire of itself”

Meditation XVII, John Donne, 1624

Working life has long provided ample opportunity to compare yourself to your coworkers: from the clothes they wear and the cars they drive to their latest luxury holiday. Trends in clothes, cars and holidays have come and gone, but the ability to compare yourself to your coworkers has not. In contrast, the rapid growth of social media over the past two decades has made it much easier to compare yourself to your friends. Photos and videos of their houses, holidays and parties are far more readily available than ever before. A short scroll through Facebook or Instagram now shows even your loosest of acquaintances’ consumption decisions – people who you would previously have never been able to compare yourself to.

People do not make decisions in a vacuum. Instead, they often evaluate their actions relative to the world – and the people – around them. Taking this casual insight seriously, I build a model of reference dependent choice where agents derive benefits from consumption relative to that of their friends. To do this, I take the benchmark model of Kőszegi and Rabin [2006] and change the determinant of the reference point. In Kőszegi and Rabin [2006], the reference point is an agent’s beliefs about her own future consumption. In contrast, I assume the reference point is a weighted sum of other agents’ consumption, where the weights capture the strength of social comparisons.

To incorporate social comparisons it is necessary to extend Kőszegi and Rabin [2006] to include multiple agents. Agents are embedded in a social network, and their reference point is derived from their neighbours. There is also a constant marginal cost of consumption and a constant marginal of benefit of connections in the social network – to capture the myriad benefits of social connections. Agents make their consumption decisions simultaneously and I study the Nash equilibria of this game.

I first consider a setting where the network is fixed, and agents choose consumption directly. With a fixed network, a unique equilibrium exists when social comparisons are not too strong. Equilibrium consumption (when it exists) relies critically on Bonacich centrality – which captures how well-connected agents are to other well-connected agents.¹ Within the confines of the model, stronger social comparisons, even by only a few agents, increase consumption but *reduce* welfare for all agents. Higher consumption by one agent pushes her neighbours to consume more, through an increase in their reference points. This creates a rush to “keep up with The Joneses”. Everyone spends more on consumption, but nobody gets any further ahead, so welfare falls.

Strikingly, a higher marginal cost of consumption can increase welfare. This arises because it reduces *other people’s* consumption, reducing the need to “keep up with The Joneses” and making

¹Katz [1953] and Bonacich [1987] originally proposed this metric. The link between equilibrium consumption and Bonacich centrality is due to Ballester et al. [2006].

makes an agent better off for a given level of her own consumption. When an agent’s Bonacich centrality is sufficiently high, this effect dominates the standard effects that cause higher costs to reduce welfare. Within my benchmark model, a flat tax would improve then welfare, even if the proceeds are entirely wasted.

Next, I examine the implications of agents choosing their neighbours (i.e. when the network is endogenous). In this case, they separate into a set of “social classes”. This is where agents are only friends with people who have the same level of consumption as they do.

One way of interpreting the cost of consumption is the effort of work needed to acquire it. Such a view implicitly thinks of work as a solitary activity. In reality however, most people work in firms. It is typically possible to increase consumption by moving to a more productive firm, without necessarily working harder. However, moving firms comes with an obvious catch – your coworkers change. Therefore we can also think of the change in coworkers’ consumption as the cost of moving firm to get higher consumption for yourself. This is unavoidable.²

To recognise this reality I then allow agents to choose their firm. This allows us to explore how social comparisons affects labour market sorting and segregation. Sorting in this context is where skilled workers work disproportionately for high productivity firms, and segregation is where skilled workers work disproportionately with each other. I consider a stylised benchmark model with two types of workers and two types of firms. The model is then able to explain two key stylised facts that have emerged from empirical studies of labour markets. First, that sorting is not complete – the most skilled workers are frequently not matching with the most productive firms. Measuring sorting as the correlation between firm and worker types, many authors find a only moderate degree of labour market sorting in practice – a correlation coefficient in the region of 0.4 to 0.6. Bagger and Lentz [2019] find this level of sorting in Danish data, Borovičková and Shimer [2017] in Austrian administrative data, and Torres et al. [2018] and Card et al. [2018] in Portuguese data.

Social comparisons with coworkers can explain the incomplete sorting observed in practice. Given the stark assumptions of the model, they can cause a skilled worker to want to be the *big fish in a small pond*. This means she wants to work with coworkers who earn less than she does, even if her own wages are lower. Consequently, some skilled workers choose to work at low productivity firms in order to work with lower paid coworkers.

The second stylised the fact the model can account for is that sorting has increased over time. Song et al. [2019] and Card et al. [2013] demonstrate that sorting and segregation have risen over time, and both attribute the majority of the increase in income inequality over recent decades to these phenomena. The model predicts increasing sorting when there is technology change that facilitates better comparisons with friends. This is because better comparisons with friends causes the big-fish effect to diminish over time, in turn increasing sorting.

²The idea that coworkers’ income affects agents’ utility has a long history in economics [Marshall, 1890, Veblen, 1899] and is supported by much more recent evidence in a range of settings (see for example Card et al. [2012], Breza et al. [2018]).

I argue that the the rise of social media is precisely such a technology change. Social media has made it far easier to observe the consumption decisions of a wide group of friends and acquaintances. There is also increasing empirical evidence that social media intensifies social comparisons [Burke et al., 2020, Verduyn et al., 2017, 2015, Krasnova et al., 2013], although it is difficult to clearly demonstrate causality. The huge increase in social media use over recent decades would then show up in my model as stronger social comparisons with friends. This makes the *big fish* effect weaker and so drives up sorting in the labour market, as workers find it more important to keep up with their friends rather than to be the *big fish* at work. More recently, the increasing shift to working from home in some industries may make social comparisons with friends more important relative the those with co-workers by causing workers to see their coworkers less frequently. This shift could also exacerbate the effects of social media – driving sorting up further.

Beyond the extent of sorting, the model also predicts the *pattern* of sorting: that is, who works where. Among skilled workers, it is those making the weakest social comparisons with friends who choose the work for low productivity firms and earn higher wages relative to their co-workers, but lower *absolute* wages. This presents a novel mechanism through which social connections increase wages. With weak social connections, the *big fish* effect looms large, and skilled workers choose jobs at low productivity firms. They are not crowded out by those with better social connections, but feel pressured into low productivity firms by a strong *big fish* effect. Hence workers with weaker social networks *choose* lower productivity jobs, rather than having worse job opportunities due to fewer referrals (see, for example, Bolte et al. [2020]). Which of these mechanisms – social comparisons or job referrals – is at work has important policy implications. If social comparisons play the leading role, then Bolte et al.’s [2020] main policy proposal for addressing inequality – one-time affirmative action to provide jobs to disadvantaged groups – will be less effective. In a world where social comparisons are important, social networks create inequality of *aspiration*. Weak social networks cause skilled workers to *prefer* jobs at low productivity firms via the *big fish* effect.

Addressing inequality of opportunity – the goal of Bolte et al.’s [2020] affirmative action – is not sufficient to ensure that skilled workers from disadvantaged groups receive the high productivity jobs. In contrast, strengthening the social networks of skilled workers from disadvantaged groups addresses both issues. It provides them with better job opportunities (via more referrals) and also changes their preferences so that they choose to work in high productivity jobs. While altering social networks is very challenging for a policymaker, it may be necessary if currently disadvantaged groups are to achieve equal labour market outcomes.

Related Literature

This paper continues a strand of the behavioural economics literature that examines the source of reference points.³ Kőszegi and Rabin [2006] provide perhaps the best known contribution to this strand. They suppose that an agent’s reference point depends on her expectations and find that in situations with uncertainty, an agent is willing to pay more for a good when she believes she is more likely to buy it. This paper considers a setting without uncertainty, and so one where Kőszegi and Rabin’s mechanism has no bite.

Beyond my novel motivation for including social comparisons in agents’ preferences, this paper makes two main contributions to the significant literature on social comparisons.⁴ First, is the striking result that a uniform increase in the marginal cost of consumption can penalise agents who are highly central in the social network while benefiting less central agents. There is an obvious and potentially important policy implication. A flat tax on consumption makes the rich better off.

This builds on Frank’s [1985a, 2001] idea that progressive consumption taxes can improve outcomes by “slowing down” the race to keep up with the Joneses. Making status goods more expensive reduces everyone’s consumption of them. But it does so in a way that doesn’t change people’s *position* in terms of consumption levels. So everyone maintains their position, but consumes less of the status goods – making them better off. The core logic of my result mirrors Frank’s argument. Introducing a network highlights that the impacts of government intervention can depend critically on an agent’s position in a social network, and that the Frank’s idea about progressive consumption taxes extend to flat taxes.⁵ It also allows for richer, or at least more explicit, patterns of social comparisons.

Further, it complements the extensive literature on the Easterlin Paradox [Easterlin, 1974].⁶ With the confines of my model, I find that higher income (captured by my model as a lower marginal cost of consumption) can reduce welfare. This is even starker than Easterlin’s finding of little-to-no relationship between income and happiness over time. This is because the only good I have in my model is a status good.

Second, it adds to the branch that focuses on labour market behaviour.⁷ Existing literature

³There is of course a much wider literature concerning reference dependence and loss aversion, with seminal contributions from Tversky and Kahneman [1979], Thaler [1980], Tversky and Kahneman [1991]. See Barberis [2013] for a review of the literature, and DellaVigna [2009], Camerer [1998] for reviews of some of the empirical evidence.

⁴Within economics, this strand dates at least to Veblen [1899]. Other classical contributions include Duesenberry [1949], Layard [1980], Frank [1985a,b], Clark and Oswald [1996], Crès et al. [1997] and Clark and Oswald [1998]. More recent contributions include Hopkins and Kornienko [2004], Luttmer [2005], Frank [2005] and Schelling [2006].

⁵This idea of extending to a flat tax is not entirely novel. Frank [1985a] argues that government regulation can prevent a race to the bottom on workplace safety standards. This is akin to a flat tax in this context. Pressure to “keep up with The Joneses” can induce workers to accept worse conditions in exchange for some additional income. But because everyone faces this pressure, nobody “gets ahead of The Joneses” and are all worse off as a result.

⁶See Clark et al. [2008] for a thorough survey. Some more recent work in this area includes Decancq et al. [2015] and Easterlin and O’Connor [2020]. Easterlin’s original finding is that individual happiness correlates positively with income at a point in time, but aggregate happiness does not rise over time as aggregate incomes increase.

⁷Some important empirical contributions include Pérez-Asenjo [2011], Card et al. [2012], Breza et al. [2018] and

shows that social comparisons impact a variety of outcomes in labour markets – most notably on wage-setting, both in competitive [Frank, 1985b] and imperfectly competitive environments [Goerke and Neugart, 2017]. I explore the impact of social comparisons on labour market sorting.

Close to my paper, Von Siemens [2012] shows how they can induce workers to sort into unemployment. Von Siemens [2012] assumes there are two types of workers, and only one type make social comparisons. He additionally makes a very stark assumption that unemployed workers only compare themselves to other unemployed workers, while the employed compare themselves to both others at their firm and workers at other firms. This asymmetry in the composition of social comparisons then drives some to choose unemployment. Related, Vásquez and Weretka [2021] shows how co-worker altruism might induce wage rigidity and hence involuntary unemployment.

In contrast, all agents in my model make social comparisons and their non-coworkers do not change when they move to a different firm. Additionally, social comparisons in my model can explain why labour market sorting is incomplete and why it has risen over recent decades.

My model is a network game with linear best replies, by now a well-studied class of games.⁸ I make two contributions to this literature. First, I to show a comparative static for rerouting a link (extending well-known results concerning adding or removing a link). Second, I demonstrate one way in which network games can be used to model a well-known behavioural bias. Some of the first results, notably Remark 1 and Proposition 1, are similar in technical detail to results in Ballester et al. [2006] and Ushchev and Zenou [2020] respectively, but their motivation via reference dependent choice is novel.

There is also a growing strand of the networks literature that examines endogenous network formation. Perhaps closest to my paper, Hiller [2017] considers a network game with strategic complementarities where links are two-sided (and so require consent from both agents), finding that equilibrium networks are either complete or empty. A key difference in this paper is that social comparisons are a negative externality (Hiller [2017] assumes positive externalities).⁹ In contrast, I find that agents form groups based on their benefits of forming links, and only link with others in the same group. This creates groups, each with a different level of consumption, who do not interact with one another – very stark social classes emerge.

Related, Ushchev and Zenou [2020] study a game where agents want to be close to the average action in their neighbourhood. They suppose some preference for conformity. In spite of the contrasting motivation and the simplified network component, Ushchev and Zenou [2020, S5.5] also find that the network becomes completely homophilous when agents can choose their links. This means agents only form links with others of the same type (their analysis assumes only two types

Fu et al. [2019]. They demonstrate that social comparisons with coworkers and are important for labour market behaviour. My contribution is theoretical.

⁸A review of this strand of literature is beyond the scope of this paper. See Bramoullé, Galeotti, and Rogers [2016, Chapter 5] and Jackson [2008, Chapter 9] for comprehensive coverage.

⁹Related, Staab [2019] studies how agents who make social comparisons form groups when there is a price to group membership. In this setting, he finds that social comparisons can reduce – rather than increase – segregation.

for simplicity).¹⁰

In spirit, this paper is most similar to Ghiglino and Goyal [2010] and Immorlica et al. [2017], which both examine social comparisons taking place on a network.¹¹ Ghiglino and Goyal [2010] consider a two-good general equilibrium exchange economy, where social comparisons apply to one good. In both this paper and in Ghiglino and Goyal [2010] Bonacich centralities are the key determinant of agents' consumption.¹² My focus on a single good and a partial equilibrium approach permits my stronger comparative statics results regarding consumption and welfare. It also facilitates investigation of endogenous network formation and the application to labour markets. Further, it also allows me to drop two assumptions Ghiglino and Goyal [2010] make for tractability: that an agent cares about her consumption relative to the simple average of her friends' consumption, and that the strength of social comparisons only depends on the number friends. Ghiglino and Goyal's 2010 use of Cobb-Douglas utility helps generate linear best replies in a setting with multiple goods. I make alternative, and even starker, simplifications – namely a single good and linear costs

Bramoullé and Ghiglino [2022] subsequently build loss aversion into the framework of Ghiglino and Goyal [2010], and find that loss aversion has a material impact on outcomes. It can generate a continuum of equilibria where all agents consume the same quantity of the status good. This occurs when agents' incomes are sufficiently similar to each other (relative to the strength of loss aversion). In contrast, agents in my model always consume above their reference point (even if only by a small amount), and so loss aversion has little bite. This is because they do not face an immovable budget constraint, but rather a constant marginal cost of consumption.

In its setup, Immorlica et al. [2017] is closer to this paper – it also takes a partial equilibrium approach and focuses on a single good. The key difference is that Immorlica et al. [2017] suppose that agents only make social comparisons with those richer than themselves. This comes at a significant cost to tractability. They find multiple equilibria, and only analyse one – the one with highest consumption by agents. There, they find that agents stratify into a “class structure” according to their position in the network. Consumption then depends on “social class”. In contrast, I assume agents form an overall reference point based on their friends' consumption. This preserves continuity in most functions, provides a unique equilibrium and allows for simpler determination of equilibrium consumption and welfare, and associated comparative statics. I also find an analogy to Immorlica et al.'s (2017) “social classes” when allowing agents to choose their own friends.

The rest of the paper is organised as follows. Section 1 sets out the model. Section 2 presents

¹⁰More recently, Sadler and Golub [2021] have nested a wide variety of endogenous network formation games into a single framework. While my game does not fit because I consider weighted networks, the stark social classes emerging in this paper are very similar to the *ordered overlapping cliques* networks in Sadler and Golub [2021].

¹¹Few treatments of social comparisons prior to Ghiglino and Goyal [2010] include an explicit network, instead focusing on average consumption (e.g. Aronsson and Johansson-Stenman [2015], Hopkins and Kornienko [2004]) or an agent's ordinal rank (e.g. Neumark and Postlewaite [1998], Frank [1985a,b]).

¹²Ghiglino and Goyal [2010] also find that Bonacich centralities determine prices – something outside the scope of this paper.

the main results when the network is fixed. Section 3 considers how agents choose friends. Section 4 considers how agents choose coworkers and jobs. Section 5 concludes.

1 Model

There are n agents. Each agent i simultaneously chooses a level of consumption $x_i \geq 0$. The agents are embedded in a weighted and directed network G . The network G is an $n \times n$ matrix, where the entry G_{ij} denotes the strength of a link from i to j . An agent j is i 's *neighbour* (or *friend*) if and only if $G_{ij} > 0$. Assume i cannot be neighbours with herself (so $G_{ii} = 0$ for all i).

Now define $\alpha_i = \sum_j G_{ij}$ and $g_{ij} = \frac{G_{ij}}{\alpha_i}$ for all i, j . Then $G \equiv \text{diag}(\alpha)g$, where $\alpha = [\alpha_1, \dots, \alpha_n]^T$. Notice that g is row stochastic (i.e. each row sums to one). Call α the *strengths* and g the *structure*. This novel decomposition will help us to better understand the network effects later on.¹³

An agent i 's *reference point* is a weighted sum of her neighbours' consumptions, where the weights are determined by the network G . Mathematically, i 's reference point is $\alpha_i \sum_j g_{ij} x_j$. An agent i cares about her consumption, x_i , only relative to her reference point, not in absolute terms. All agents face a constant marginal cost of consumption, $c > 0$, and receive a marginal benefit $b_i > 0$ of links. So i 's utility function is:

$$u_i = f \left(x_i - \alpha_i \sum_j g_{ij} x_j \right) - cx_i + b_i \alpha_i \sum_j g_{ij} \quad (1)$$

where $f(\cdot)$ is twice continuously differentiable, strictly increasing and concave. To focus on the interesting where agents would choose positive but finite consumption absent any social comparisons, I assume that $f'(0) > c > f'(+\infty)$. Further, assume that for any i , there exists some j such that $b_i = b_j$. That is, every agents shares her value of b_i with at least one other agent.

Discussion. The key feature of reference dependent choice here is that an agent get benefits from consumption *relative to* her reference point, but pays costs for absolute consumption. The key innovation relative to existing work on reference dependent choice is a novel source of the reference point – comparisons with neighbours in a social network. Several of the assumptions I make are for clarity of exposition only. I relax each in the Online Appendix. Doing so does not change the core insights of the model.

- **Linear sub-utility.** In Kőszegi and Rabin [2006], the benefit of consumption relative to the reference point (which they call “gain-loss utility”) is a function of agents’ *utility from*

¹³Mathematically, this is a trivial decomposition. Conceptually however, it allows us to separate out the effects of social comparisons per se from the network effects.

consumption. Formally, it is $f(m(x_i) - m(r_i))$, where r_i is i 's reference point and $m(\cdot)$ is a concave function. I assume this sub-utility function, $m(\cdot)$ is linear.

Relaxing this, and using some strictly concave function $m(\cdot)$ in Equation (1) does not change the insights of the model, but makes them harder to follow. However, it does prevent closed form solutions, and so breaks the tight link between network centrality and outcomes.

- **Single good.** Moving to a K -dimensional consumption bundle with additively separable preferences, as in Kőszegi and Rabin [2006], would have no impact on the results. Here, this means that agents make optimal choices for each good separately. So a K -dimensional bundle collapses down to K separate single-good problems.
- **Homogeneous cost.** My model, embodied by Equation (1), assumes that the marginal cost of consumption is the same for all agents. Introducing arbitrary heterogeneity in the cost parameter, c_i , does not materially change the results, but complicates matters without providing additional insight.
- **No loss aversion.** In Equation (1), the function $f(\cdot)$ is everywhere concave. Capturing loss aversion à la Tversky and Kahneman [1979] involves making $f(\cdot)$ kinked at zero and convex in the negative domain. In my model adding this has no impact, as agents always choose to consume above their reference point (even if only by a very small amount).

Additionally, the linking benefit, b_i , captures the simple idea that social connections bring benefits. These include improved health, job opportunities, access to new information and risk sharing, to name but a few.¹⁴ In the interest of simplicity, I assume that there are no benefits from friends-of-friends. Only direct friends bring benefits. Implicitly this assumes indirect connections are not of first-order importance. In any case, these benefits only play a role when the network is endogenous.

Bonacich centrality. An agent cares about her neighbours' consumption, who in turn care about their neighbours' consumption, and so on. She will end up making comparisons, albeit indirectly, with everyone she is *connected* to in the network. Agent i is connected to j if there exists a *walk* from i to j . A walk is a sequence of links $g_{i,i+1}, \dots, g_{j-1,j}$ such that $g_{k,k+1} > 0$ for all links in the sequence.

The strength with which i compares herself to j depends on the weight of these walks (of any length) from i to j . Bonacich centrality [Bonacich, 1987], counts the total weight of all walks (of any length, from 1 to ∞) from agent i to *all other agents*. It captures the idea that making strong

¹⁴Reynolds and Kaplan [1990] show that social connections increase cancer survival rates, and Newman and Zainal [2020], Gerst-Emerson and Jayawardhana [2015], Kawachi and Berkman [2001] document mental health benefits. Calvo-Armengol and Jackson [2004], Chandrasekhar et al. [2020] show the roles social networks can play in getting a job, and Banerjee et al. [2013] shows improved access to other information. Karlan et al. [2009], Ambrus et al. [2014] demonstrate how networks can facilitate risk sharing.

comparisons with society as a whole involves making strong comparisons with those who themselves make strong comparisons.

Definition 1 (Bonacich Centrality). *The Bonacich centrality of agent i is:*

$$C_i^b = \sum_j \left[\sum_{k=0}^{\infty} G^k \right]_{ij} \quad (2)$$

Bonacich centrality is typically defined as $C^b(G, \beta) = \sum_{k=0}^{\infty} \beta^k G^k$, where β is a decay parameter (see, for example, Jackson [2008, Ch9]). In my model however, the choice of decay parameter is determined by individual preferences as it is implicitly built into the reference strength.

2 Social Comparisons, Consumption and Welfare

The natural starting point is to ask how agents behave in equilibrium. Before turning to the result, it is helpful to see an agent's best response function.

$$BR(x_{-i}) = f'^{-1}(c) + \alpha_i \sum_j g_{ij} x_j \quad (3)$$

Note that best responses are linear in neighbours' consumption. I now show that there is a unique equilibrium set of actions, where consumption is proportional to agents' Bonacich centrality.

Remark 1. *If $\lambda_1 < 1$ then there is a unique Nash Equilibrium with $x_i^* = C_i^b \cdot f'^{-1}(c)$ for all i , where λ_1 is the largest eigenvalue modulus of the matrix G .*

Intuitively, the largest eigenvalue λ_1 captures the extent to which a change is amplified by the network. When $\lambda_1 < 1$, the impact of an agent's consumption on that of her neighbours is (eventually) less than one-for-one. A sufficient condition for this is $\alpha_i < 1$ for all i .¹⁵ In that case, the impact is less than one-for-one effect happens at every step. In contrast, Bonacich centrality is not well defined when $\lambda_1 > 1$, and the solution ceases to have a meaningful economic interpretation. Therefore I will assume that $\lambda_1 < 1$ throughout the remainder of the paper.

Agents with high centrality are more closely connected to other agents with high consumption. This means they have a high reference point, which in turn pushes them to choose high consumption. So consumption exhibits strategic complementarities: higher consumption by one agent raises the marginal benefit of consumption for her neighbours (via the effect on the reference point). This leads to similar outcomes to other network games of strategic complements [Ballester et al., 2006, Calvó-Armengol et al., 2009]. The differences in functional form and in motivation do not matter for this.

¹⁵ $\alpha_i < 1$ for all i implies that G is a substochastic matrix. So $\lambda_1 < 1$ follows from the Perron-Frobenius Theorem (see for example [Seneta, 2006, Ch.1] or Pillai et al. [2005]).

I now turn to the impact of individual components of the model on consumption and welfare. The obvious first component is the strength of social comparisons.

Proposition 1. *If $\lambda_1 < 1$: (i) x_i^* is weakly increasing, and (ii) u_i^* is weakly decreasing, in α_j for all i, j , and strictly so if $i = j$.*

When an agent feels social comparisons more keenly (i.e. her reference strength increases), she chooses higher consumption to better keep up with “The Joneses”. This more intense social competition has a knock-on effect: an increase in i ’s consumption raises her neighbours’ reference points, pushing them to consume more. This effect ripples out through society, increasing *everyone’s* consumption. However this social competition is destructive. Given the assumptions of the model, agents can only ever keep up with “The Joneses” – collectively they get stuck in a rat race and nobody gets ahead. So everyone becomes worse off as their benefits of consumption don’t change, but they have to spend more.

Coupled with the growing evidence that social media intensifies social comparisons (for example Burke et al. [2020], Verduyn et al. [2017, 2015] and Krasnova et al. [2013]), this suggests that social media has been harmful for agents. Such an insight is clearly not new, but this result clarifies the mechanism through which it actually harms its users. The increased importance of online social networks also provides policy-makers with more options, as they are much easier to influence than in-person networks. Online, policy interventions across the entire network are possible. For example, Instagram has trialled removing the number of “likes” attached to a post [Mosseri, 2021, Burke, 2019] in an apparent attempt to limit some of these harmful social comparisons.

A comparative static for the marginal cost of consumption are of clear interest. The first part of the result is obvious and of little interest: when consumption is more expensive, agents do less of it. In contrast, the second part is far less intuitive and has important policy implications.

Proposition 2. *If $\lambda_1 < 1$: (i) x_i^* is strictly decreasing and convex in c for all i , (ii) u_i^* is strictly increasing in c if and only if $\frac{C_i^b - 1}{C_i^b} > \frac{F(c)}{-F'(c) \cdot c}$ for all i , where $F(c) \equiv f'^{-1}(c)$.*

When an agent’s Bonacich centrality is above a threshold, then she becomes *better off* when consumption becomes more costly.¹⁶ The threshold is governed by the concavity of the benefits function, i.e. how quickly marginal utility diminishes when she increases her consumption. Increased cost has three effects. First, it increases the cost of a given amount of consumption, which lowers welfare. Second, it reduces an agent’s consumption, which reduces overall cost. Third, it reduces her neighbours’ consumption, allowing her to reduce her own consumption further, while still keeping up with “The Joneses”. The strength of this final effect depends on how sensitive her consumption is to that of her neighbours. In other words, it depends on her Bonacich centrality. When Bonacich centrality is higher, a given change in marginal cost induces a larger fall

¹⁶Note that $C_i^b \geq 1$, so the left-hand side of part (ii) is positive and increasing in C_i^b . $F(c) > 0$ and $F'(c) < 0$, so the right-hand side is positive.

in consumption, leading the effect of lower overall spending to dominate the higher marginal cost effect.

The condition in part (ii) of Proposition 2 is rather terse, and it is easier to see the intuition for a special case. Suppose that $f(a) = a^\gamma$. Then u_i^* is strictly increasing in c if and only if $C_i^b > \frac{1}{\gamma}$. Here, this link between centrality and the concavity of the benefits function is much more transparent.

Within the confines of the model, this result identifies a condition where taxation will increase welfare, even if the government completely wastes the funds raised. This model focuses on a single good. In reality, there are many goods, and the strength of social comparisons may differ across goods.¹⁷ Goods for which this condition is met are the ideal candidates for raising tax revenue. Note that the threshold depends on an individual agent’s centrality. Therefore, it is possible that some agents (those with high centrality and low utility) benefit from an increase in c , while others (those with low centrality and high utility) lose out. So flat taxes may actually redistribute utility from the well-off to the badly-off (although the well-off agents in this model are those with low consumption). If a policy-maker cares about utility, then flat taxes can be progressive, even if they do not benefit everyone.

The final moving part of the model is the reference structure, g . This is very difficult to examine in general. To help with tractability, I focus on a very simple change to the network structure (while holding the reference strength constant). Call this change a “comparison shift”.

Definition 2 (Comparison shift). *A comparison shift is an $n \times n$ matrix D , where $D_{ru} = \phi$, $D_{rd} = -\phi$ for $r, u, d \in N$, and all other elements are equal to zero.*

Mathematically, it is an increase in one element of the matrix G , and a corresponding decrease in another element in the same row. Economically, this is where an agent r makes *stronger* social comparisons with u and correspondingly *weaker* with d .¹⁸ To make sure that solutions always exist, I will work in the setting where $\alpha_i < 1$ for all agents, and only consider perturbations that are *feasible*.¹⁹ To be feasible, we need $\phi \in [0, G_{rd}]$. That is, r cannot reduce the amount she compares herself to d by more than the amount she initially compares herself to d .

With this definition in place, I can now show the impact of this simple change to the network.

Proposition 3. *Consider a comparison shift, D , of magnitude ϕ . Then: (i) x_i^* is strictly increasing, and (ii) u_i^* is strictly decreasing, in ϕ for all i if and only if $C_u^b > C_d^b$.*

The direction of the change in consumption is determined by the relative Bonacich centralities of agents d and u . If r switches to listening to a higher centrality agent (i.e. $C_u^b > C_d^b$) then

¹⁷I extend the analysis to multiple goods in the Online Appendix. But note that this generalisation relies on additive separability of different goods. Further, Bramoullé and Ghiglino [2022] suggests that the generalisation to multiple goods is less straightforward when there is an immovable budget constraint.

¹⁸An increase/decrease in a single element is equivalent to a change in α_i , and so is covered by Proposition 1.

¹⁹The $\lambda_1 < 1$ assumption is not sufficient here. This is because a comparison shift changes the network, and so can change the largest eigenvalue.

everyone’s consumption rises in equilibrium. This is because r has a higher reference point after the change, and so adjusts her consumption upwards. This ripples out through the network as agents adjust to their neighbours’ new, higher, consumption. This increase in consumption then reduces welfare (by the same logic as in Proposition 1).

Characterising the impact of a comparison shift on consumption and welfare relies on a generalisation of the Woodbury matrix identity developed by Chang [2006]. To the best of my knowledge, the use of matrix perturbation results of this type is novel in the economics literature.²⁰

A limitation of Proposition 3 is that it only considers a very simple change to the network structure. This is because the impact of any one comparison shift depends on the whole network immediately before it occurs. So the impact of a set of comparison shifts on a network G is not simply the sum of their individual impacts. Of course, if the total “volume” of the shifts is small, then this naive summing up will be a close approximation, as the interaction effect between the various comparison shifts will be small.²¹ I formalise these insights in the Online Appendix.

Discussion. The first key message for policy-makers is that consumption and welfare can move in opposite directions in this setting. Governments typically use measures of output and consumption as a proxy for welfare. GDP growth is widely sought after, and a failure to achieve it is generally considered a serious problem for an incumbent government. However, this paper suggests that seeking consumption growth – even through innovation, improved productivity, and other measures to reduce the marginal cost of consuming goods – may not succeed in delivering improvements in welfare. In a world with pervasive social comparisons, increasing consumption can be harmful. This means that well-meaning governments could be misguided in their efforts to improve their citizens’ welfare by seeking to increase output and consumption. This key message is not novel. It dates back at least to Abramovitz [1959] and Easterlin [1974] within economics, and is closely related to Easterlin’s Paradox.

Second, flat taxes can aid welfare by reducing how much “The Joneses” consume, reducing an agent’s need to keep up with them. In some cases, this effect can dominate the standard downsides of higher unit costs, and overall welfare can increase. This does not require the government to do anything useful with the revenue raised – taxation for its own sake may be a social good. This builds on Frank’s [1985a, 2001, 2008] idea that progressive consumption taxes improve welfare by “slowing down” the race to keep up with The Joneses. It also complements the standard approach to optimal taxation, which points to concentrating taxes in markets with standard negative externalities, or

²⁰The proof to Proposition 3 uses a simplified version of Chang [2006]. For convenience, I provide a proof of simplified theorem in the Online Appendix.

²¹However, it is easy to analyse the impact of arbitrary changes to the set of Bonacich centralities. For a given value of c , the distributions of consumption and welfare are simply linear transformations of the distribution of Bonacich centralities. The challenge is understanding how centrality responds to changes in the network, not how outcomes respond to changes in centrality. The impact of changes in centrality on aggregate welfare and inequality follow straightforwardly from this insight.

where the distortion due to taxes is low. This paper suggests that taxes should also be targeted towards goods where social comparisons are felt more keenly. In effect, this is because social comparisons act as a (non-standard) negative externality.²²

3 Choosing your friends: Endogenous Network Formation

I now allow agents to *choose* their friends (i.e. form the network endogenously) and examine the network that emerges. This is a natural case to study because in many contexts people are able to choose who they spend time with and pay attention to, and by extension who they compare themselves to.

Assume that links here are *symmetric* (that is $G_{ij} = G_{ji}$), agents can propose any links they like, but links require mutual consent to form, and can be broken unilaterally. So agent i cannot become friends with j unless j is willing to reciprocate. Agents reach an equilibrium when nobody wants to unilaterally cut a link, and no *pair* of agents want to form a link together. This is *pairwise stability* [Jackson and Wolinsky, 1996]. Formally:

Definition 3 (Pairwise stability). *A network G is pairwise stable if:*

- (i) *for all $G_{ij} > 0$: $u_i(G) \geq u_i(G - G_{ij})$ and $u_j(G) \geq u_j(G - G_{ij})$,*
- (ii) *for all $G_{ij} = 0$: if $u_i(G + G_{ij}) > u_i(G)$ then $u_j(G + G_{ij}) < u_j(G)$*

With this definition in place, I can now show that allowing agents to form the network endogenously leads to extreme segregation.

Proposition 4. *In all pairwise stable networks, if $b_i \geq cf'^{-1}(c)$, then $G_{ij} > 0$ only if $b_i = b_j$, and if $b_i < cf'^{-1}(c)$ then $G_{ij} = 0$ for all j .*

The network formation process yields extreme homophily (homophily is the tendency for agents to be disproportionately connected to those similar to themselves). While highly stylised, this is consistent with the high degree of homophily typically found in real social networks [McPherson et al., 2001, Jackson, 2008]. Agents segregate into a well-defined series of “social classes”, where all the agents in a given group have the same value of b_i/c .²³ Any agents whose benefit of linking is not sufficiently high (i.e. $b_i < cF(c)$) will not form any links at all – as the costs of any social comparisons outweigh the benefits of friendships. The condition in Proposition 4 that defines “sufficiently high” linking benefits is terse. To see it more clearly suppose that $f(a) = a^\gamma$. Then an agent i wants to form links if and only if $b_i \geq \gamma^{1/(1-\gamma)}c^{\gamma/(\gamma-1)}$. Even more transparently, when $\gamma = 0.5$, this amounts to $b_i \geq \frac{1}{4c}$.

²²In the Online Appendix, I also show that if there is little correlation between the strength of agents’ social comparisons and the structure of the network, then we can safely ignore the network. This is useful for policy-makers if measuring the network is too expensive or time-consuming.

²³As I have suppressed heterogeneity in c_i for clarity of exposition, it simply amounts to sharing a value of b_i here. The ratio proper has bits when c_i is heterogeneous.

An agent would ideally like to form links with others who have very low consumption, because this minimises the impact on her reference point. Further, she would like to keep forming links until the marginal benefit of doing so equals the marginal cost. However, the second effect precludes the first. An agent i keeps forming links until her consumption reaches her benefit-cost ratio. However, this prevents her from forming links with some agent j who has a lower benefit-cost ratio, because j would not want to form a link with i (as i will have higher consumption than j).

Each “social class” has a different level of consumption. Those with higher marginal benefits of links, relative to the marginal cost of consumption, consume more.

Corollary 1. *In all pairwise stable networks, $x_i^* = \max\{\frac{b_i}{c}, f'^{-1}(c)\}$, and $C_i^b = \max\{\frac{b_i}{c \cdot f'^{-1}(c)}, 1\}$.*

This makes it meaningful to think of “higher” and “lower” social classes. While each group has no distinguishing feature in terms of network structure, they are ordered by the amount that their members consume. In the “higher” social classes, agents make stronger social comparisons within their class, and consequently consume more.

Discussion. When agents can choose their own friends, they end up grouping themselves with like-minded people. This creates a set of “social classes”, where those with the highest benefit from network connections relative to the marginal cost of consumption: (i) form the strongest connections, (ii) make the strongest social comparisons, and (iii) consume the most. While I have suppressed heterogeneity in c for clarity of exposition, this result goes through unchanged with heterogeneous c_i (see Online Appendix). So the driving force is not the benefit of network connections alone. Rather it is the ratio of the benefit of connections and the cost of consumption.²⁴ The result, and its key insight – that social comparisons drive people into distinct social classes – does not depend on which factor dominates.

This is obviously a stylised benchmark – we rarely have complete control over who we spend time with in the real world. Family, coworkers and the neighbours across the street are difficult to avoid. We should not expect to see such a stark result in practice. Rather, Proposition 4 shows that, among the friends agents do freely choose, social comparisons push them to choose others in the same “social class”. Naturally, they will have some links to other social classes, most plausibly through family or coworker ties. It is connections with coworkers that I turn to next.

4 Choosing your coworkers: Labour Market Sorting

One appealing interpretation for the cost parameter c is as a cost of effort of working to acquire consumption. Implicitly, this fits a world where people work on their own and can only increase

²⁴One might argue that the parameter values might be correlated. It is plausible that agents with high values of c_i might also have higher values of b_i , as difficulty in purchasing consumption increase the risk-sharing and insurance benefits of social networks. Any such correlation has no impact on the result.

their consumption by working more or harder. In reality however, most people work for firms. And people can attain higher consumption by moving to more productive firms, without necessarily working harder. However, moving firms comes with an obvious catch – your coworkers change. This is unavoidable. It is this change in the coworkers, and the accompanying changes in social comparisons with them, that constitutes the effective cost of changing firms.²⁵

A significant challenge to modelling these social comparisons when people are changing firms is that the set of coworkers changes entirely. It is not meaningful to separate changes in a person’s firm from their network of coworkers. This is not just a technical challenge, but a conceptual one. When someone changes firms it is unclear who they will form links with in a new firm, and how the structure of these links might change as a result of the move. A natural benchmark is to assume that that people form equally strong links with all of their coworkers. Under this assumption an agent compares herself to average coworker consumption.²⁶

The only difference between friends and coworkers is that an agent stops making social comparisons with her old coworkers when she moves firm, and starts making comparisons with her new ones. In Section Section 2 I ignored the distinction between friends and coworkers because the distinction is mute when agents do not change firms. Note that in reality most of us keep in contact with, and therefore keep making social comparisons with, some people we used to work with. In my model, these people count as friends.²⁷ Here, we can split the network into the friends component and the coworkers component. It is merely conceptual, but nevertheless helpful in thinking about the problem. So i ’s utility function is:

$$u_i = f(x_i - \alpha_{1i} \sum_{j \in \text{friends}} g_{ij} x_j - \alpha_{2i} \bar{x}_m) + b \sum_j \alpha_{1i} g_{ij} \quad (4)$$

where \bar{x}_m is the average consumption of coworkers at i ’s firm. Notice that this is identical to Equation (1) except there is now no direct marginal cost of consumption, and I have separated out friends and coworkers for clarity.

I assume that there are two types of agent, *skilled* (S) and *unskilled* (U), and two types of firm, *high productivity* (H) and *low productivity* (L). Conditional on their type, every agent earns, and therefore consumes, more when she works at a high productivity firm than a low productivity firm. Conditional on the firm’s type, a skilled agent earns more than an unskilled agent. Formally:

²⁵In practice changing firms will often involve changing both coworkers and effort. I abstract from this, and focus on the role of social comparisons.

²⁶Given this assumption, “firm” is best interpreted as a team or department. It is unreasonable to think that an agent compares herself to everyone in her organisation. Far more plausible is the claim that she compares herself to the more modest number of people she works with day-to-day.

²⁷A more nuanced view recognised that when we stop working with someone we often spend less time with them, and so will make weaker social comparisons. My model copes with this easily. Such a person should be considered both a friend and a coworker. The social comparisons that fall away when you stop working with them are the coworker component. The comparisons that remain are the friends component.

$x_{iH} > x_{iL}$, for all i , and $x_{Sm} > x_{Um}$ for $m \in \{H, L\}$.²⁸

Further, each firm has a fixed number of job openings. This is stark, and corresponds to a short-run world where firms face some friction that prevents them from either adjusting wages or creating/destroying jobs. I also assume that all firms prefer to hire skilled over unskilled workers. This is a consequence of pushing the wage-setting process into the background. For simplicity, I also assume that firms are large – in the sense that they employ lots of workers. This is not an important assumption, but makes the the average coworker consumption, \bar{x}_m , much simpler.²⁹

I also assume that the total strength of social comparisons, $\alpha_{1i} + \alpha_{2i} \equiv \alpha_i$, is fixed for each agent and (to guarantee existence) that $\alpha_i < 1$. This builds in a trade-off between social comparisons with friends and with coworkers. One motivation for this is that an agent’s propensity to compare herself with others is fixed outside of this model, by a set of underlying traits like her personality. Given that, the relative comparisons with different people depend on the time she spends with them. Spending more time comparing herself with friends necessarily entails less time comparing herself with coworkers. This assumption is a particularly parsimonious way of introducing such a trade-off.

We can now examine how social comparisons determines agents’ *sorting* into firms. But first, we need a formal definition of sorting. With only two types of agents and two types of firms, there is a very simple definition.

Definition 4 (Sorting). *Sorting is equal to the fraction of skilled workers that work for high productivity firms, c .*

Sorting captures the extent to which workers and firms match along lines of skill and productivity. Absolute deviations of c from $1/2$ capture *segregation* – the extent to which workers are concentrated in firms with other workers of the same type. If the consumption premium from working at a high productivity firm (compared to a low productivity firm) is larger for skilled workers than for unskilled ones, then skilled workers will never disproportionately concentrate in low productivity firms in equilibrium. So sorting and segregation coincide, and it suffices to simply consider c . This is because the only incentive for a skilled worker to choose a low productivity firm is to reduce her reference point by having lower-earning co-workers. If skilled workers worked disproportionately for low productivity firms, then this incentive would not be present.

The key insight in this section is that social comparisons with coworkers *induce* less labour market sorting.

Proposition 5. *If the strength of social comparisons with coworkers weakly increases for all workers, then labour market sorting weakly decreases.*

²⁸I leave the process of wage-setting in the background as it is not the focus of this paper. Note however that I implicitly assume that an agent’s wage does not depend on who else works at the firm. I am abstracting from notions of team production.

²⁹This allows me to ignore the impact a single worker has on the social comparisons others at the firm.

Social comparisons with coworkers can cause a skilled worker to choose to work for low productivity firms because they want to be a “*big fish in a small pond*”. This is because average wages are lower at low productivity firms, and having lower coworker wages reduces her reference point. When comparisons with coworkers are strong enough, the benefit of improving her wage *relative to her co-workers* outweighs the cost of a lower *absolute* wage. If social comparisons with coworkers become stronger, then the “big fish” effect becomes more prominent, so skilled workers shift to low productivity firms.

Labour market sorting in turn impacts a range of other outcomes. Therefore we can easily link social comparisons and these other outcomes. For example, if output and wages (which are implicitly functions of agent and firm types) are supermodular, then an increase in the strength of social comparisons with coworkers weakly reduces output and reduces the variance in wages.³⁰ If the functions are submodular, the results are flipped.

Notice that the structure of the social network (of friends) does not play a direct role here. This is because it, and friends’ wages, do not change in response to a worker’s choice of firm. In contrast, the *strength* of social comparisons with friends does matter, because of the assumption that total weight a workers places of social comparisons is fixed. This means that stronger comparisons with friends causes weaker comparisons with coworkers. Consequently, outcomes for skilled workers – where they work and how much they earn – are intimately tied to their social networks.

Proposition 6. *There exists a threshold value α_1^{crit} such that a skilled worker works at a high productivity firm if and only if $\alpha_{1i} \geq \alpha_1^{crit}$.*

This means that all skilled workers at high productivity firms make stronger social comparisons with friends than all skilled workers at low productivity firms. Those with the strongest social comparisons with friends experience the weakest desire to be a “big fish in a small pond”, and so choose to work for high productivity firms.

The result gives a complete characterisation of which type of firm every skilled worker chooses. A fraction c of skilled workers work at high productivity firms. These are precisely the fraction c of skilled workers with the strongest comparisons with friends. The result does not extend to unskilled workers because unskilled workers in my model are forced to take whatever vacancies are left over, and firms are indifferent between all unskilled workers.

4.1 Discussion

Conceptually, this section merely divides an agent’s neighbours into coworkers and friends (that is, people she stops comparing herself with when she changes firm, and those who remain constant) and considers changes in coworkers as the cost of moving firms. However, this decomposition assumes

³⁰The variance result does not extend to standard measure of inequality, such as the Gini coefficient, because the strength of social comparisons affect the average wage.

that agents do actually make social comparisons with coworkers and are willing to change jobs as a result. There is ample empirical evidence to support this assumption. In field experiments both Cohn et al. [2014] and Breza et al. [2018] find that creating social comparisons within teams reduces the effort exerted by the lower-paid team members. Obviously my model abstracts from effort, but these studies show that comparisons with coworkers are present and do affect behaviour. Closer to my model, in another field experiment Card et al. [2012] find that disclosing coworkers' pay makes relatively lower paid less satisfied with their job and more likely to search for a new one.

The model can then explain two key stylised facts in labour market behaviour: that sorting and segregation in labour markets are incomplete and that they have risen significantly over recent decades [Song et al., 2019, Card et al., 2013].³¹ These analyses raise two questions. Why didn't skilled workers sort into high productivity firms and only work with similarly-skilled peers in the first place? And what caused sorting and segregation to increase? In my model, the answer to the first question is straightforward – social comparisons. Comparisons with coworkers make agents to want to be a “big fish in a small pond”, causing incomplete sorting.

An increase in sorting follows from agents making (relatively) weaker social comparisons with coworkers. I contend that this has been driven by changes in technology – most importantly the rapid growth in social media use over recent decades.³² This works through two mechanisms. First, it has made it easier to observe looser friends' and acquaintances' consumption [Liu et al., 2016, Ellison et al., 2007], and compare yourself to them [Verduyn et al., 2017]. Second it has made it easier to maintain contact with *ex*-coworkers. In my model this transforms coworkers in friends. Remember that the definition of friends is simply people who you continue comparing yourself to when you change firm.³³ In contrast, there has been no change of comparable magnitude to the ways in which we can compare ourselves to our coworkers.

There are of course a number of existing explanations for rising sorting. Some popular arguments include: rising return to skills (e.g. Acemoglu and Autor [2011]), increasing use of outsourcing, and increasing worker-firm complementarities. Rising returns to skills would mechanically increase sorting and segregation, but Song et al. [2019, p38] show that this can only explain a (small) minority of the changes. Increases in outsourcing are similarly an important component. Goldschmidt and Schmieder [2017] show it explains 9% of the increase in German wage inequality since the 1980's. Again, this can only be one part of a bigger picture.

An important question is whether my model has different implications to existing explanations. In a sense, I simply provide another mechanism linking changes in technology to increases in

³¹Song et al. [2019] find the majority of the increase in inequality in the US over the past 40 years is due to this increased sorting and segregation. So it is obviously an important question on its own terms. Song et al. [2019] use a confidential database from the U.S. Social Security Administration, which contains earnings records for every individual ever issued a U.S. Social Security number. Earnings data are uncapped and contain all sources of income deemed as remuneration by the U.S. Internal Revenue Service. They use the full dataset from 1981 to 2013.

³²In 2021, 72% of Americans use at least one social media platform, up from only 5% in 2005 [Pew Research, 2021].

³³There is absolutely no requirement that you actually like them.

labour market sorting. The precise technology change – the rise of social media – is different to other explanations, but this has happened at the same time as many other changes that stem from increasing internet use. So how can we distinguish my mechanism from existing one?

My model provides clear predictions – albeit under stark assumptions – about *who works where*, in addition to to just the *extent* of sorting. Skilled workers choose high wage jobs because the pressure from their social network to earn more outweighs their desire to be the “big fish” in the workplace. So skilled workers with weaker social networks are the ones who choose low productivity firms. This is because they feel the “big fish” pressure (relatively) more strongly. It is the particular pattern of sorting – the link between an individual’s strength of social comparisons and where she chooses to work – that separates my model from existing explanations.

There is anecdotal evidence that suggests people with weaker social networks – in the sense of their friends having lower consumption – end up with lower paying jobs themselves. See for example Jackson [2019] and Putnam [2016]. This is in line with the predictions of the my model.³⁴ In my set-up, it is something akin to a poverty of aspiration that pushes some skilled workers into lower productivity, lower paying, firms. They are not held back by a lack of ability or insider access. Rather, it is their preferences that drive them to earning lower wages.³⁵

This link between social networks, preferences, and in turn, labour market decisions, is in line with existing work. Existing empirical work shows that an agent’s social connections can affect participation [Nicoletti et al., 2018]), productivity [Mas and Moretti, 2009], and enrolment in retirement plans [Duflo and Saez, 2003], among others. My result suggests a new margin – namely the choice of where to work.

5 Conclusion

This paper has taken seriously the simple idea that keeping up with “The Joneses” matters. Doing so has a number of striking implications for policymakers. Stronger social comparisons increase consumption but reduce welfare. The usual close relationship between consumption and welfare need not exist. This suggests standard consumption-based proxies for welfare, such as G.D.P., may be a poor guide when targeting welfare. It complements work related to the Easterlin puzzle [Easterlin, 1995, Frey and Stutzer, 2002], which asks why happiness correlates with incomes at a point in time, but doesn’t continually increase as incomes rise. If we believe that social comparisons are a significant part of agents’ preferences, then we should be much more careful in advocating for continued economic growth (Easterlin [2013] makes a similar argument using macro-level data).

Further, increasing the cost of consumption can benefit everybody – taxation can be welfare

³⁴Note however, that it is also consistent with a complementary explanation – that social networks affect agents’ *ability* to get high paying jobs.

³⁵The implications for policy-makers are mixed. On the one hand, giving (currently) poorer skilled agents stronger social networks can increase their earnings. However, this requires changing these agents’ preferences, and leaves them with lower utility.

improving, even if the government wastes the revenue. So goods where social comparisons are important are ideal candidates for taxation when a government does need to raise revenue. This complements the standard advice to target goods with negative externalities and markets where taxes do not distort behaviour much. When the network is endogenous – when agents can choose links for themselves – a clear set of “social classes” emerge. Agents only form friendships with people like them. This happens because everyone wants to escape the pressures of making social comparisons with those richer than themselves, and as a consequence they can only become friends with those at the same level of consumption.

Extending the model to examine a simple labour market helps explain two stylised facts from empirical studies of labour markets. First, labour market sorting is incomplete – some skilled workers continue to work for low productivity firms. Second, the degree of sorting has been rising over time. Social comparisons with co-workers induce incomplete sorting because they make agents want to be a “big fish a small pond”. At a certain point they are willing to move to a low productivity firm, losing income in absolute terms, in exchange for becoming better off relative to their (new) co-workers. I contend that the rise of social media has made it much easier to compare yourself to your friends, and so has made co-workers a relatively less important source of social comparisons. This trend would drive the increase in labour market sorting found in the data [Song et al., 2019].

Social comparisons can also determine which workers work for which firms, at least among skilled workers. Skilled workers with weaker social networks will choose to work for low productivity firms. Without the social pressures from their friends to earn high incomes, the “big fish in a small pond” effect looms relatively large for them. This suggests that social networks affect labour market outcomes through impacts on agents’ *preferences*, rather than just opportunities.

References

- M. Abramovitz. The welfare interpretation of secular trends in national income and product. In *The allocation of economic resources: Essays in honor of Bernard Francis Haley*, pages 1–22. Stanford University Press, 1959.
- D. Acemoglu and D. Autor. Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of Labor Economics*, volume 4, pages 1043–1171. Elsevier, 2011.
- A. Ambrus, M. Mobius, and A. Szeidl. Consumption risk-sharing in social networks. *American Economic Review*, 104(1):149–82, 2014.
- T. Aronsson and O. Johansson-Stenman. Keeping up with the joneses, the smiths and the tanakas: On international tax coordination and social comparisons. *Journal of Public Economics*, 131: 71–86, 2015.
- J. Bagger and R. Lentz. An empirical model of wage dispersion with sorting. *The Review of Economic Studies*, 86(1):153–190, 2019.
- C. Ballester, A. Calvó-Armengol, and Y. Zenou. Who’s who in networks. wanted: The key player. *Econometrica*, 74(5):1403–1417, 2006.
- A. Banerjee, A. G. Chandrasekhar, E. Duflo, and M. O. Jackson. The diffusion of microfinance. *Science*, 341(6144), 2013.
- N. C. Barberis. Thirty years of prospect theory in economics: A review and assessment. *Journal of Economic Perspectives*, 27(1):173–96, 2013.
- L. Bolte, N. Immorlica, and M. O. Jackson. The role of referrals in immobility, inequality, and inefficiency in labor markets. 2020. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3512293.
- P. Bonacich. Power and centrality: A family of measures. *American Journal of Sociology*, 92(5): 1170–1182, 1987.
- K. C. Border. *Fixed Point Theorems with Applications to Economics and Game Theory*. Cambridge University Press, 1985.
- K. Borovičková and R. Shimer. High wage workers work for high wage firms. Technical report, National Bureau of Economic Research, 2017.
- Y. Bramoullé and C. Ghiglino. Loss aversion and conspicuous consumption in networks. 2022.
- Y. Bramoullé, A. Galeotti, and B. Rogers. *The Oxford Handbook of the Economics of Networks*. Oxford University Press, 2016.
- E. Breza, S. Kaur, and Y. Shamdasani. The morale effects of pay inequality. *The Quarterly Journal of Economics*, 133(2):611–663, 2018.

- C. Burke. This is what removing instagram “likes” could do for your mental health, November 2019. URL <https://www.bustle.com/p/how-instagram-removing-likes-could-affect-your-mental-health-19342825>. Accessed on 23 May, 2021.
- M. Burke, J. Cheng, and B. de Gant. Social comparison and facebook: Feedback, positivity, and opportunities for comparison. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–13, 2020.
- A. Calvo-Armengol and M. O. Jackson. The effects of social networks on employment and inequality. *American Economic Review*, 94(3):426–454, 2004.
- A. Calvó-Armengol, E. Patacchini, and Y. Zenou. Peer effects and social networks in education. *The Review of Economic Studies*, 76(4):1239–1267, 2009.
- C. F. Camerer. Prospect theory in the wild: Evidence from the field. In D. Kahneman and A. Tversky, editors, *Choices, Values and Frames*. Cambridge: Russell Sage Foundation, 1998.
- D. Card, A. Mas, E. Moretti, and E. Saez. Inequality at work: The effect of peer salaries on job satisfaction. *American Economic Review*, 102(6):2981–3003, 2012.
- D. Card, J. Heining, and P. Kline. Workplace heterogeneity and the rise of west german wage inequality. *The Quarterly Journal of Economics*, 128(3):967–1015, 2013.
- D. Card, A. R. Cardoso, J. Heining, and P. Kline. Firms and labor market inequality: Evidence and some theory. *Journal of Labor Economics*, 36(1):13–70, 2018.
- A. G. Chandrasekhar, M. Morten, and A. Peter. Network-based hiring: Local benefits; global costs. Technical report, National Bureau of Economic Research, 2020.
- F. C. Chang. Inversion of a perturbed matrix. *Applied Mathematics Letters*, 19(2):169–173, 2006.
- A. E. Clark and A. J. Oswald. Satisfaction and comparison income. *Journal of Public Economics*, 61(3):359–381, 1996.
- A. E. Clark and A. J. Oswald. Comparison-concave utility and following behaviour in social and economic settings. *Journal of Public Economics*, 70(1):133–155, 1998.
- A. E. Clark, P. Frijters, and M. A. Shields. Relative income, happiness, and utility: An explanation for the easterlin paradox and other puzzles. *Journal of Economic literature*, 46(1):95–144, 2008.
- A. Cohn, E. Fehr, B. Herrmann, and F. Schneider. Social comparison and effort provision: Evidence from a field experiment. *Journal of the European Economic Association*, 12(4):877–898, 2014.
- H. Crès, C. Ghiglino, and M. Tvede. Externalities, internalization and fluctuations. *International Economic Review*, pages 465–477, 1997.
- K. Decancq, M. Fleurbaey, and E. Schokkaert. Happiness, equivalent incomes and respect for individual preferences. *Economica*, 82:1082–1106, 2015.
- S. DellaVigna. Psychology and economics: Evidence from the field. *Journal of Economic Literature*, 47(2):315–72, 2009.

- J. S. Duesenberry. *Income, saving, and the theory of consumer behavior*. Harvard University Press, 1949.
- E. Dufflo and E. Saez. The role of information and social interactions in retirement plan decisions: Evidence from a randomized experiment. *The Quarterly journal of economics*, 118(3):815–842, 2003.
- R. A. Easterlin. Does economic growth improve the human lot? some empirical evidence. In *Nations and households in economic growth*, pages 89–125. Elsevier, 1974.
- R. A. Easterlin. Will raising the incomes of all increase the happiness of all? *Journal of Economic Behavior & Organization*, 27(1):35–47, 1995.
- R. A. Easterlin. Happiness, growth, and public policy. *Economic Inquiry*, 51(1):1–15, 2013.
- R. A. Easterlin and K. O’Connor. The easterlin paradox. *Available at SSRN 3743147*, 2020.
- N. B. Ellison, C. Steinfield, and C. Lampe. The Benefits of Facebook “Friends”: Social Capital and College Students’ Use of Online Social Network Sites. *Journal of Computer-Mediated Communication*, 12(4):1143–1168, 2007.
- R. H. Frank. *Choosing the right pond: Human behavior and the quest for status*. Oxford University Press, 1985a.
- R. H. Frank. The demand for unobservable and other nonpositional goods. *The American Economic Review*, 75(1):101–116, 1985b.
- R. H. Frank. *Luxury fever: Why money fails to satisfy in an era of excess*. Simon and Schuster, 2001.
- R. H. Frank. Positional externalities cause large and preventable welfare losses. *American economic review*, 95(2):137–141, 2005.
- R. H. Frank. Should public policy respond to positional externalities? *Journal of Public Economics*, 92(8-9):1777–1786, 2008.
- B. S. Frey and A. Stutzer. What can economists learn from happiness research? *Journal of Economic Literature*, 40(2):402–435, 2002.
- J. Fu, M. Sefton, and R. Upward. Social comparisons in job search. *Journal of Economic Behavior & Organization*, 168:338–361, 2019.
- K. Gerst-Emerson and J. Jayawardhana. Loneliness as a public health issue: the impact of loneliness on health care utilization among older adults. *American Journal of Public Health*, 105(5):1013–1019, 2015.
- C. Ghiglino and S. Goyal. Keeping up with the neighbors: social interaction in a market economy. *Journal of the European Economic Association*, 8(1):90–119, 2010.
- L. Goerke and M. Neugart. Social comparisons in oligopsony. *Journal of Economic Behavior & Organization*, 141:196–209, 2017.

- D. Goldschmidt and J. F. Schmieder. The rise of domestic outsourcing and the evolution of the german wage structure. *The Quarterly Journal of Economics*, 132(3):1165–1217, 2017.
- J. Hadar and W. R. Russell. Rules for ordering uncertain prospects. *The American Economic Review*, 59(1):25–34, 1969.
- T. Hiller. Peer effects in endogenous networks. *Games and Economic Behavior*, 105:349–367, 2017.
- E. Hopkins and T. Kornienko. Running to keep in the same place: Consumer choice as a game of status. *American Economic Review*, 94(4):1085–1107, 2004.
- R. A. Horn and C. R. Johnson. *Matrix analysis*. Cambridge University Press, 2012.
- N. Immorlica, R. Kranton, M. Manea, and G. Stoddard. Social status in networks. *American Economic Journal: Microeconomics*, 9(1):1–30, 2017.
- M. O. Jackson. *Social and economic networks*. Princeton University Press, 2008.
- M. O. Jackson. *The human network: How your social position determines your power, beliefs, and behaviors*. Vintage, 2019.
- M. O. Jackson and A. Wolinsky. A strategic model of social and economic networks. *Journal of Economic Theory*, 71(1):44–74, 1996.
- D. Karlan, M. Mobius, T. Rosenblat, and A. Szeidl. Trust and social collateral. *The Quarterly Journal of Economics*, 124(3):1307–1361, 2009.
- L. Katz. A new status index derived from sociometric analysis. *Psychometrika*, 18(1):39–43, 1953.
- I. Kawachi and L. F. Berkman. Social ties and mental health. *Journal of Urban Health*, 78(3):458–467, 2001.
- B. Köszegi and M. Rabin. A model of reference-dependent preferences. *The Quarterly Journal of Economics*, 121(4):1133–1165, 2006.
- H. Krasnova, H. Wenninger, T. Widjaja, and P. Buxmann. Envy on facebook: a hidden threat to users’ life satisfaction? *Proceedings of the 11th International Conference on Wirtschaftsinformatik*, 2013.
- R. Layard. Human satisfactions and public policy. *The Economic Journal*, 90(360):737–750, 1980.
- D. Liu, S. E. Ainsworth, and R. F. Baumeister. A meta-analysis of social networking online and social capital. *Review of General Psychology*, 20(4):369–391, 2016.
- E. F. Luttmer. Neighbors as negatives: Relative earnings and well-being. *The Quarterly journal of economics*, 120(3):963–1002, 2005.
- A. Marshall. *Principles of Economics*. London: Macmillan, 1890.
- A. Mas and E. Moretti. Peers at work. *American Economic Review*, 99(1):112–45, 2009.
- A. Mas-Colell, M. D. Whinston, J. R. Green, et al. *Microeconomic Theory*. Oxford University Press, 1995.

- M. McPherson, L. Smith-Lovin, and J. M. Cook. Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27(1):415–444, 2001.
- A. Mosseri, April 2021. URL <https://twitter.com/mosseri>. Accessed on 23 May, 2021.
- D. Neumark and A. Postlewaite. Relative income concerns and the rise in married women’s employment. *Journal of Public Economics*, 70(1):157–183, 1998.
- M. G. Newman and N. H. Zainal. The value of maintaining social connections for mental health in older people. *The Lancet Public Health*, 5(1):12–13, 2020.
- C. Nicoletti, K. G. Salvanes, and E. Tominey. The family peer effect on mothers’ labor supply. *American Economic Journal: Applied Economics*, 10(3):206–34, 2018.
- E. Pérez-Asenjo. If happiness is relative, against whom do we compare ourselves? implications for labour supply. *Journal of Population Economics*, 24(4):1411–1442, 2011.
- Pew Research Center. Social media factsheet. Technical report, 2021. URL <https://www.pewresearch.org/internet/fact-sheet/social-media/>.
- S. U. Pillai, T. Suel, and S. Cha. The perron-frobenius theorem: some of its applications. *IEEE Signal Processing Magazine*, 22(2):62–75, 2005.
- R. D. Putnam. *Our kids: The American dream in crisis*. Simon and Schuster, 2016.
- P. Reynolds and G. A. Kaplan. Social connections and risk for cancer: prospective evidence from the alameda county study. *Behavioral Medicine*, 16(3):101–110, 1990.
- M. Rothschild and J. E. Stiglitz. Increasing risk: I. A definition. *Journal of Economic Theory*, 2(3):225–243, 1970.
- E. Sadler and B. Golub. Games on endogenous networks. *arXiv preprint arXiv:2102.01587*, 2021.
- T. C. Schelling. *Micromotives and macrobehavior*. WW Norton & Company, 2006.
- E. Seneta. *Non-negative Matrices and Markov Chains*. Springer Science & Business Media, 2006.
- J. Song, D. J. Price, F. Guvenen, N. Bloom, and T. Von Wachter. Firming up inequality. *The Quarterly Journal of Economics*, 134(1):1–50, 2019.
- M. Staab. The formation of social groups under status concern. *Working Paper*, 2019.
- R. Thaler. Toward a positive theory of consumer choice. *Journal of Economic Behavior & Organization*, 1(1):39–60, 1980.
- S. Torres, P. Portugal, J. T. Addison, and P. Guimaraes. The sources of wage variation and the direction of assortative matching: Evidence from a three-way high-dimensional fixed effects regression model. *Labour Economics*, 54:47–60, 2018.
- A. Tversky and D. Kahneman. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–291, 1979.

- A. Tversky and D. Kahneman. Loss aversion in riskless choice: A reference-dependent model. *The Quarterly Journal of Economics*, 106(4):1039–1061, 1991.
- P. Ushchev and Y. Zenou. Social norms in networks. *Journal of Economic Theory*, 185:104969, 2020.
- J. Vásquez and M. Weretka. Co-worker altruism and unemployment. *Games and Economic Behavior*, 130:224–239, 2021.
- T. Veblen. *The Theory of the Leisure Class: An Economic Study of Institutions*. New York: Macmillan, 1899.
- P. Verduyn, D. S. Lee, J. Park, H. Shaback, A. Orvell, J. Bayer, O. Ybarra, J. Jonides, and E. Kross. Passive facebook usage undermines affective well-being: Experimental and longitudinal evidence. *Journal of Experimental Psychology: General*, 144(2):480, 2015.
- P. Verduyn, O. Ybarra, M. Résibois, J. Jonides, and E. Kross. Do social network sites enhance or undermine subjective well-being? a critical review. *Social Issues and Policy Review*, 11(1): 274–302, 2017.
- F. A. Von Siemens. Social preferences, sorting, and competition. *The Scandinavian Journal of Economics*, 114(3):780–807, 2012.

A Proofs

Before presenting the proofs, I make two simplifications to notation. (1) Let $B = (I - G)^{-1}$, and $B_{ij} \equiv (I - G)_{ij}^{-1}$, (2) $F(\cdot) \equiv f'^{-1}(\cdot)$. We only use these where they aid exposition.

Proof of Remark 1. In matrix notation, the First Order Condition is: $x^*(I - G) = f'^{-1}(c)$. $\lambda_1 < 1$ implies that $(I - G)^{-1}$ exists, is unique, and can be represented by the Neumann series $(I - G)^{-1} = \sum_{k=0}^{\infty} G^k$. \square

At this point, it is helpful to state a relationship that holds in equilibrium. It is obvious (so a proof is omitted), but I will use it in various proofs.

Lemma A.1. *In equilibrium, $f'(x_i^* - \alpha_i \sum_j g_{ij} x_j^*) = c$ and $x_i^* - \alpha_i \sum_j g_{ij} x_j^* = f'^{-1}(c)$*

Proof of Proposition 1. (i) $(I - \text{diag}(\alpha)g)^{-1} = \sum_{k=0}^{\infty} (\text{diag}(\alpha)g)^k$. This is the a standard Neumann Series (see for example Horn and Johnson [2012]). If a given α_j increases, then all elements of $(\text{diag}(\alpha)g)^k$ weakly increase for any k . Since $x_i^* \propto \sum_j (I - \text{diag}(\alpha)g)_{ij}^{-1}$, then x_i^* is weakly increasing for all i . This relationship must be strict when $i = j$ since $(\text{diag}(\alpha)g)_{ii}$ is strictly increasing in α_i by construction.

(ii) Take equilibrium utility, $u_i^* = f(x_i^* - \alpha_i \sum_j g_{ij} x_j^*) - cx_i^*$, and differentiate with respect to α_k .

$$\frac{du_i^*}{d\alpha_k} = \left(\frac{dx_i^*}{d\alpha_k} - \alpha_i \sum_j g_{ij} \frac{dx_j^*}{d\alpha_k} \right) f'(x_i^* - \alpha_i \sum_j g_{ij} x_j^*) - c \frac{dx_i^*}{d\alpha_k}$$

Applying Lemma A.1 an simplifying yields: $\frac{du_i^*}{d\alpha_k} = -\alpha_i c \sum_j g_{ij} \frac{dx_j^*}{d\alpha_k}$ Finally, noticing that x_i^* is weakly decreasing in α_k for all i, k (and that $c > 0, \alpha \geq 0, g_{ij} \geq 0$) yields the result. \square

Proof of Proposition 2. (i) Follows from the equilibrium found in Proposition 1 and the fact that $f'^{-1}(\cdot) \equiv F(c)$ is strictly decreasing and convex. **(ii)** First, recall that $x_i^* = C_i^b F(c)$, and that C_i^b depends only on the network, and not on c . Therefore $\frac{dx_i^*}{dc} = F'(c)C_i^b$ Now take equilibrium utility, $u_i^* = f(x_i^* - \alpha_i \sum_j g_{ij} x_j^*) - cx_i^*$, and differentiate with respect to c .

$$\frac{du_i^*}{dc} = F'(c)(C_i^b - \alpha_i \sum_j g_{ij} C_j^b) f'(x_i^* - \alpha_i \sum_j g_{ij} x_j^*) - cF'(c)C_i^b - F(c)C_i^b$$

Notice that $F(c)(C_i^b - \alpha_i \sum_j g_{ij} C_j^b) = x_i^* - \alpha_i \sum_j g_{ij} x_j^*$, which is equal to $F(c)$ by Lemma A.1. Therefore $(C_i^b - \alpha_i \sum_j g_{ij} C_j^b) = 1$. Applying Lemma A.1 again and simplifying yields:

$$\frac{du_i^*}{dc} = F'(c)c - cF'(c)C_i^b - F(c)C_i^b$$

By rearranging, we can see that $\frac{du_i^*}{dc} > 0$ if and only if: $\frac{C_i^b - 1}{C_i^b} > \frac{F(c)}{-F'(c) \cdot c}$ (during this rearranging the inequality sign will flip when we divide through by $(1 - C_i^b)$ because this term is negative). \square

The following Lemma restates the results of Chang [2006]. It underpins all results relating to comparison shifts of the network. To my knowledge, this result is not well known in economics, so I provide a proof in Appendix G.

Lemma A.2 (A perturbation theorem. Chang [2006]). *If D is a comparison shift, then $(I - [G + D])^{-1} - (I - G)^{-1} = H$. Where:*

$$H = \frac{\phi}{1 - \phi(B_{ur} - B_{dr})} \begin{pmatrix} B_{1r}(B_{u1} - B_{d1}) & \cdots & B_{1r}(B_{un} - B_{dn}) \\ \vdots & \ddots & \vdots \\ B_{nr}(B_{u1} - B_{d1}) & \cdots & B_{nr}(B_{un} - B_{dn}) \end{pmatrix}$$

Lemma A.3. *The change in agent i 's consumption due a comparison shift D , Δx_i^* , is equal to:*

$$\frac{\phi B_{ir}(C_u^b - C_d^b)}{1 - \phi(B_{ur} - B_{dr})} \cdot f'^{-1}(c)$$

Proof of Lemma A.3. From Proposition 1: $x_i^* = f'^{-1}(c) \sum_j B_{ij}$ (by the definition of B). From Lemma A.2, the optimal action after a comparison shift D is: $x_i'^* = f'^{-1}(c) \sum_j (B_{ij} + H_{ij})$. Denote the change in optimal action $x_i'^* - x_i^* \equiv \Delta x_i^*$. Then using the definition of H : $\Delta x_i^* = f'^{-1}(c) \sum_j \frac{\phi B_{ir}(B_{uj} - B_{dj})}{1 - \phi(B_{ur} - B_{dr})}$. Finally, using the definition of Bonacich centrality (Definition 1) yields the result. \square

Proof of Proposition 3. (i) Arbitrarily small changes in ϕ are possible, so Δx_i^* is continuously differentiable in ϕ . For ease of exposition, let: $B_{ir}(C_u^b - C_d^b)f'^{-1}(c) \equiv y$ and $(B_{ur} - B_{dr}) \equiv z$. Let $\hat{x}_i^* \equiv x_i^* + \Delta x_i^*$ be the new equilibrium value of consumption. Then $\hat{x}_i^* = x_i^* + y\phi(1 - z\phi)^{-1}$. By standard rules of differentiation (and some rearranging): $\frac{d\hat{x}_i^*}{d\phi} = y \cdot (1 - z\phi)^{-2}$. $y > 0$ as $B_{ir} > 0$, $f'^{-1}(c) > 0$ (by construction), and $C_u^b - C_d^b > 0$ (by assumption). Therefore $\frac{d\hat{x}_i^*}{d\phi} > 0$. (ii) taking the second expression from Lemma A.1, and substituting it into equilibrium utility, we have that $\hat{u}_i^* = f(F(c)) - c\hat{x}_i^*$. Since \hat{x}_i^* is increasing in ϕ , then \hat{u}_i^* is clearly decreasing in ϕ . \square

It is now helpful to state an implication of Lemma A.1.

Lemma A.4. *For a fixed network G , $\frac{du_i^*}{dG_{ij}} = -x_j^*c_i + b_i$*

Proof of Lemma A.4. Take eq. (1), substitute in equilibrium values of consumption, x_i^* , and utility, u_i^* , and then differentiate with respect to G_{ij} . This yields:

$$\frac{du_i^*}{dG_{ij}} = \left(\frac{dx_i^*}{dG_{ij}} - x_j^* \right) f' \left(x_i^* - \sum_j G_{ij}x_j^* \right) - c_i \frac{dx_i^*}{dG_{ij}} + b_i$$

From Lemma A.1 we know that $f'(x_i^* - \sum_j G_{ij}x_j^*) = c_i$. Substituting this in and simplifying yields the result. \square

It is also clearer to prove Corollary 1 and Proposition 4 together.

Proof of Proposition 4 and Corollary 1. (i) First, notice that when i has no links, $x_i^* = f'^{-1}(c)$. Therefore, $b_i < cf'^{-1}(c)$ then i will never want to form a link (as no agent can have consumption below $f'^{-1}(c)$). **(ii)** Now consider cases where $b_i \geq cf'^{-1}(c)$. Consider any pair of agents i, j such that $b_i = b_j$. Recall that in section 1 we *assumed* that for any i , there exists some j such that $b_i = b_j$. Since $\frac{du_i^*}{dG_{ij}} = -x_j^*c_i + b_i$ for all agents, then i and j want to form (or strengthen) a link if $x_i^* < \frac{b_i}{c}$ and $x_j^* < \frac{b_j}{c}$. Therefore, $x_i^* \geq \frac{b_i}{c}$ for all i . **(iii)** Now notice that i does not want to form a link with any agent k where $b_k > b_i$ (or wants to cut a link if one exists). This is because $x_k^* \geq \frac{b_k}{c}$ which implies that $-x_k^*c + b_i < 0$. Since links require mutual consent, there can be no links between two agents i, j where $b_i \neq b_j$. This is because either $b_i > b_j$ or $b_j > b_i$, so one of the agents would want to cut the link. **(iv)** Since agents can only form links with others who share a value of b , then $x_i^* \leq \frac{b_i}{c}$. Otherwise, $-x_i^*c + b_j < 0$ (since $b_i = b_j$) and so j would cut the link. Therefore $x_i = \frac{b_i}{c}$ for all i where $b_i < cf'^{-1}(c)$. **(v)** We know from Remark 1 that for any network G , $C_i^b = \frac{x_i^*}{f'^{-1}(c)}$. The Bonacich centrality result in corollary 1 follows trivially from substituting the value x_i^* found above into this equation. \square

Proof of Proposition 5. Firm choice. Firms always choose skilled workers over unskilled workers (by assumption), so it suffices to consider where the skilled workers choose to work. Unskilled workers fill up all remaining vacancies (since we have assumed that there are as many vacancies as workers). A skilled worker i will apply for a job at a high productivity firm only if and only if:

$$f(x_{SH} - \alpha_{1i} \sum_l g_{lj}\bar{x}_j - \alpha_{2i}\bar{x}_H) \geq f(x_{SL} - \alpha_{1i} \sum_l g_{lj}x_j - \alpha_{2i}\bar{x}_L) \quad (\text{A.1})$$

Simplifying and rearranging yields $x_{SH} - \alpha_{2i}\bar{x}_H \geq x_{SL} - \alpha_{2i}\bar{x}_L$. After further rearranging: $1/\alpha_{2i} \geq (\bar{x}_H - \bar{x}_L)/(x_{SH} - x_{SL})$. Let c be the fraction of skilled workers working for at high productivity firms. Then $\bar{x}_H = cx_{SH} + (1-c)x_{UH}$. Using this and the corresponding expression for \bar{x}_L yields:

$$\frac{1}{\alpha_{2i}} \geq \frac{[(x_{SH} - x_{UH}) + (x_{SL} - x_{UL})] \cdot c + (x_{UH} - x_{UL})}{(x_{SH} - x_{SL})}. \quad (\text{A.2})$$

$(x_{SH} - x_{UH})$ [resp. $(x_{SL} - x_{UL})$] is the skill premium at the high [resp. low] productivity firm. $(x_{SH} - x_{SL})$ [resp. $(x_{UH} - x_{UL})$] is the ‘‘firm premium’’ for a skilled [resp. unskilled] worker. By assumption, all these terms are strictly positive, so RHS of Equation (A.2) is strictly increasing in c . Clearly the LHS of Equation (A.2) is strictly decreasing in α_{2i} . For any c there exists a critical value of α_{2i} such that equation A.2 holds with equality.³⁶ Denote this value α_{2i}^{crit} . So α_{2i}^{crit} is a

³⁶i.e. skilled worker i is indifferent between working at a high productivity and low productivity firm

strictly decreasing function of c . Denote this function $\alpha_{2i}^{crit} = h(c)$, with $h'(c) < 0$.

Existence. Let $F(\cdot)$ denote the C.D.F. of α_{2i} (i.e. the fraction of skilled workers with an α_{2i} value of less than or equal to α_2). Therefore $F(h(c))$ is the fraction of skilled workers who want to work at high productivity firms, given that a fraction c of skilled workers actually work at high productivity firms. An equilibrium is then a fixed point of this equation: $F(h(c)) = c$. RHS of this is strictly increasing in c , and LHS is weakly decreasing in c , so there exists a fixed point, c^* .

Monotonicity. Suppose that α_{2i} weakly increases for all i . Then $F(a)$ weakly decreases for all a (this follows from well-known properties of C.D.F.s). This in turn reduces $F(h(c))$ for any $h(c)$, which weakly reduces c^* . Notice that $F(h(c)) = 1$ for all values of c where $h(c) \geq \max_i \alpha_{2i}$. \square

Proof of Proposition 6. From Proposition 5, for any c , there exists an agent-specific critical value α_{2i}^{crit} such that: **(a)** when $\alpha_{2i} = \alpha_{2i}^{crit}$, i is *indifferent* between high productivity and low productivity firms, **(b)** when $\alpha_{2i} < \alpha_{2i}^{crit}$, i strictly prefers a *high* productivity firm, **(c)** when $\alpha_{2i} > \alpha_{2i}^{crit}$, i strictly prefers a *low* productivity firm. Proposition 5 also proved that there exists some equilibrium level of labour market sorting, c^* .

This yields an equilibrium critical value for α_2 . Denote it $\alpha_{2,eq}^{crit}$. Therefore, all skilled workers with $\alpha_{2i} \leq \alpha_{2,eq}^{crit}$ work at high productivity firms. All skilled workers with $\alpha_{2i} \geq \alpha_{2,eq}^{crit}$ work at low productivity firms. Finally, notice that $\alpha = \alpha_{1i} + \alpha_{2i}$ (by assumption), so a threshold in α_2 is easily converted to a threshold in α_1 . \square

End of main text.
This page is intentionally left blank.

For Online Publication

B A Comment for Policy-makers

Consider a special case where agents’ reference strengths are uncorrelated with the structure of the social network. Its key feature is that the structure of the social network will not have any impact on outcomes.

Assumption 1 (Uncorrelated case). *The reference structure, g , is uncorrelated with the reference strength: $\text{corr}(g_{ij}, \alpha_j) = 0$ for all agents i, j .*

While Assumption 1 is appears rather abstract, it covers the natural case where all agents have the same reference strength ($\alpha_i = \alpha$ for all i), so it is not vacuous. Under this assumption, the network effects drop out.

Corollary B.1. *Under Assumption 1: x_i^* and u_i^* do not depend on the reference structure, g .*

Assumption 1 requires that i ’s neighbours are, on average, “Mr and Mrs Average” in terms of how strongly they make social comparisons. In turn, their neighbours are also, on average, “Mr and Mrs Average”, and so on throughout the whole social network. There is no way for a global network effect to build – it will always get averaged out. The clearest case where this assumption holds is where all agents have the same reference strength. This shows how homogeneity of the local social comparisons shuts down the effect of heterogeneity in global network structure.

This simplification also allows us to make stronger claims about the impact of reference strengths. First, an equilibrium exists if and only if the average reference strength is less than one, and Bonacich centralities collapse to a simple function of the reference strengths (in particular, $C_i^b = 1 + \frac{\alpha_i}{1-\alpha}$ for all i).

The tight link between an agent’s reference strength and their optimal consumption makes it straightforward to see how changes in the distribution of these reference strengths feeds through into the distribution of consumption. First [resp. second] order stochastically dominant shifts in the distribution of the the references strengths causes a first [resp. second] order stochastically dominant shift in the distribution of equilibrium consumption.³⁷ Inequality in social comparisons drives consumption inequality, even in a setting where agents are otherwise identical.

Implications for policy-making. Measuring networks in practice – especially the weighted networks present in my model – is typically difficult and resource intensive. In a perfect world with unlimited resources and attention, we should of course examine whether this assumption holds in a

³⁷First Order Stochastic Dominance and Second Order Stochastic Dominance formalise the notions of (unambiguously) “bigger” and “more spread out” respectively. They are first due to Hadar and Russell [1969] and Rothschild and Stiglitz [1970]. More modern coverage can be found in [Mas-Colell et al., 1995, Ch 6.D]. I prove the claim regarding equilibrium outcomes below. The additional claims follow straightforwardly from this.

network, and proceed accordingly. In reality, however, corollary B.1 presents a pragmatic approach for policy-makers and governments short on time and money, so long as the assumption is not violated too egregiously.

Proofs

Proof of Corollary B.1. First, let $a \equiv \text{diag}(\alpha)$ for convenience. Using a Neumann Series representation, we can express $(I-ag)^{-1} = \sum_{k=1}^{\infty} (ag)^k$. Therefore, $C_i^b = \sum_j (I-ag)_j^{-1} = \sum_j \sum_{k=1}^{\infty} [(ag)^k]_{ij}$. We deal with the first two terms individually, and then all further terms by induction. Clearly $\sum_j \mathbf{1}\{i=j\} = 1$, and $\sum_j (ag)_{ij} = \alpha_i \sum_j g_{ij} = \alpha_i$ by construction. Induction: for $k=2$;

$$\sum_j (ag)_{ij}^2 = \sum_j \sum_s (ag)_{is} (ag)_{sj} = \sum_j \sum_s \alpha_i g_{is} \alpha_s g_{sj} = \alpha_i \sum_s g_{is} \alpha_s \sum_j g_{sj}$$

By Assumption 1: $\sum_s g_{is} \alpha_s = \bar{\alpha} \sum_s g_{is}$. Then notice that $\sum_j g_{sj} = 1$ for any s . This yields: $\sum_j (ag)_{ij}^2 = \alpha_i \cdot \bar{\alpha}$. Now assume that for $k=t > 2$, $\sum_j (ag)_{ij}^t = \alpha_i \cdot \bar{\alpha}^{t-1}$. Then show for $k=t+1$. By definition $\sum_j G_{ij}^{t+1} = \sum_j [(ag) \cdot (ag)^t]_{ij}$. For clarity of exposition, we let $G_{ij}^t = \mathcal{A}_{ij}$.

$$\sum_j (ag)_{ij}^{t+1} = \sum_j \sum_s \alpha_i g_{is} \mathcal{A}_{sj} = \alpha_i \sum_s g_{is} \sum_j \mathcal{A}_{sj}$$

Using the assumption for $k=t$, $\sum_j \mathcal{A}_{sj} = \alpha_s \bar{\alpha}^{t-1}$. Therefore, $\sum_j (ag)_{ij}^{t+1} = \alpha_i \cdot \bar{\alpha}^{t-1} \sum_s g_{is} \alpha_s$. Using Assumption 1 and the fact that rows of g sum to 1 as before: $\sum_j (ag)_{ij}^{t+1} = \alpha_i \cdot \bar{\alpha}^t$. Therefore, $C_i^b = \sum_j \sum_{k=0}^{\infty} [(ag)^k]_{ij} = 1 + \alpha_i (1 + \bar{\alpha} + \bar{\alpha}^2 + \bar{\alpha}^3 + \dots)$. Clearly this does not depend on the network. Finally, it follows from Remark 1 that x_i^* and u_i^* do not depend on the network structure, g . \square

Remark B.1. Under Assumption 1: (i) an equilibrium exists if and only if $\bar{\alpha} < 1$, and (ii) when $\bar{\alpha} < 1$: $C_i^b = 1 + \frac{\alpha_i}{1-\bar{\alpha}}$ for all i , where $\bar{\alpha} = \frac{1}{n} \sum_i \alpha_i$

Proof of Remark B.1. From the proof to Corollary B.1, we have $C_i^b = 1 + \alpha_i / (1 - \bar{\alpha})$. An equilibrium exists if and only if Bonacich centrality, C_i^b is well defined. Clearly this happens if and only if $\bar{\alpha} < 1$. \square

C Non-linear sub-utility

The model in Section 1 assumes that the sub-utility is linear. This was a simplification compared to the Kőszegi and Rabin [2006] benchmark. However, this was only for clarity of exposition. It is easier to understand how the model works when more things are linear. It also allows the proofs to use simpler machinery that readers, especially those somewhat familiar with networks, ought to be more comfortable with.

Here, I introduce non-linear sub-utility. The model is exactly the same as in Section 1, except that i 's utility function is now:

$$u_i = f \left(m(x_i) - \alpha_i \sum_j g_{ij} m(x_j) \right) - cx_i + b_i \sum_j \alpha_i g_{ij} \quad (1^*)$$

where $m(\cdot)$ is twice continuously differentiable, strictly increasing, strictly concave, and WLOG $m(0) = 0$. I will not provide intuition alongside the results because while the maths becomes more complex with this additional non-linear function, there is no new economic insight.

One disadvantage of introducing this non-linear sub-utility is that a closed form characterisation of equilibrium behaviour is no longer possible. This makes the network effects far harder to see. Nevertheless, existence and uniqueness are unaffected.

Remark C.1. *Suppose $\alpha_i < 1$ for all i . There is a unique Nash Equilibrium.*

Proposition 1 then goes through unaffected.

Proposition C.1. *If $\alpha_i < 1$ for all i : (i) x_i^* is weakly increasing, and (ii) u_i^* is weakly decreasing, in α_j for all i, j , and strictly so if $i = j$.*

Proposition 2 is subject to some minor modification because the threshold in part (ii) is difficult to characterise with non-linear $m(\cdot)$

Proposition C.2. *If $\alpha_i < 1$ for all i : (i) x_i^* is strictly decreasing in c , (ii) u_i^* is strictly increasing in c if x_i^* is sufficiently sensitive to c (specifically, if $\frac{dx_i^*}{dc} < \min\left\{\frac{-x_i^*}{c}, \frac{m'(x_i^*)}{c \cdot m''(x_i^*)}\right\}$).*

Without a closed form solution for x_i^* , it is not possible to show the tight link between changes to the network structure and changes in equilibrium consumption (i.e. an analogue to Proposition 3). Nevertheless, simply looking at the First Order Condition is suggestive.

$$m(x_i) - \alpha_i \sum_j g_{ij} m(x_j) = F \left(\frac{c}{m'(x_i)} \right). \quad (C.1)$$

We can see that agents' actions are complements. If one of i 's neighbours consumes more, this pushes up i 's reference point, which in turn pushes her to consume more in an effort to keep up.

Intuitively, if i shifts from comparing herself with low consumption friends to high consumption neighbours, then her reference point increases, and so she chooses to consume more. The more “The Joneses” consume, the more you need to do to keep up with them.

Proofs

Proof of Remark C.1. Consider the First Order Condition [FOC] (with some minimal rearranging) – Equation (C.1). LHS is strictly increasing in x_i . $m'(x_i)$ is strictly decreasing in x_i , and $F(\cdot)$ is strictly decreasing in its argument. So RHS is strictly decreasing in x_i . Hence there is a unique value of x_i that solves the First Order Condition. I use x_i^* to denote *equilibrium* consumption (i.e. when *all* FOCs hold simultaneously). For clarity, let \hat{x}_i denote the value that solves just agent i 's FOC, for given values of x_j , $j \neq i$.

Existence. Note that \hat{x}_i is increasing in x_j . So if no agent wants to choose more than A when all other agents are choosing A , then there cannot exist an equilibrium where any agent chooses A or more. Suppose all other agents choose $x_j = A$. Then i wants to choose less than A if and only if: $m(A)[1 - \alpha_i] > F(c/m'(A))$.³⁸ Since LHS is increasing in A , and RHS is decreasing in A , there must exist some finite value of A such that this inequality is true for all i . Denote this \hat{A} . Given this, we only need to consider the *compact* space $[0, \hat{A}]^n$. The functions are continuous by assumption. Therefore Brouwer's Fixed Point Theorem (e.g. Border [1985]) guarantees existence.

Uniqueness. Proof by contradiction. Denote the equilibrium with the smallest action x^* . Suppose there is another equilibrium, $x^{*'} \equiv x^* + D$ (it is convenient and WLOG to write the second equilibrium as the initial one, plus some change). Consider the agent i whose action increases the most when moving from x^* to $x^{*'}$. That is, the agent i such that $D_i \geq D_j$ for all $j \neq i$. Solve the first order condition for this agent i , assuming that all other agents are playing the equilibrium $x^{*'}$. This is, we solve:

$$x_i - \alpha_i \sum_j g_{ij}[x_j^* + D_j] = F\left(\frac{c}{m'(x_i)}\right)$$

WLOG let the solution takes the form $x_i^* + z_i$.

$$x_i^* + z_i - \alpha_i \sum_j g_{ij}x_j^* - \alpha_i \sum_j g_{ij}D_j = F\left(\frac{c}{m'(x_i^* + z_i)}\right)$$

³⁸Note that after we pulled $m(A)$ out of the summation, $\sum_j g_{ij} = 1$, and so disappears.

Substituting in the solution to the FOC from the initial equilibrium x^* (and rearranging):

$$z_i - \alpha_i \sum_j g_{ij} D_j = F\left(\frac{c}{m'(x_i^* + z_i)}\right) - F\left(\frac{c}{m'(x_i^*)}\right)$$

Since $F(\cdot)$ is decreasing in its argument.

$$z_i - \alpha_i \sum_j g_{ij} D_j < 0$$

This actually assumes that $z_i > 0$. But if $z_i \leq 0$ then we have already reached a contradiction, since we require the increase in i 's action to be larger than the increase in all other agents' actions (when moving from x^* to $x^{*'}).$ Then applying the requirement that $D_i \geq D_j$ for all $j \neq i$:

$$z_i < \alpha_i \sum_j g_{ij} D_j \leq \alpha_i \sum_j g_{ij} D_i = \alpha_i D_i < D_i$$

Contradiction. □

Proof of Proposition C.1. For some fixed parameters g, α, c, b and functions $f(\cdot), m(\cdot)$ there is a unique equilibrium x^* . **(i)** Suppose that $\alpha_j \uparrow$ for some j . Then $\hat{x}_j \uparrow$ (i.e. j 's optimal action rises *conditional on* all other agents' actions). In turn $\hat{x}_j \uparrow \implies \hat{x}_k \uparrow$ for all k s.t. $g_{kj} > 0$. In turn this increases x_i for all i (but only weakly so, since there is no guarantee that there exists a directly path from j to i).

Remark C.1 guarantees a unique equilibrium. So this process must eventually converge to an equilibrium. But clearly all equilibrium actions weakly rises, and strictly so for the "initial" agent (who experienced the increase in α_j). **(ii)** In equilibrium we have:

$$m(x_i^*) - \alpha_i \sum_j g_{ij} m(x_j^*) = F\left(\frac{c}{m'(x_i^*)}\right)$$

We can substitute this back into the utility function to find equilibrium utility.

$$u_i^* = f\left(F\left[\frac{c}{m'(x_i^*)}\right]\right) - cx_i^*$$

We know from (i) that $\alpha_j \uparrow$ weakly increases x_i^* . Finally, notice that u_i^* is strictly decreasing in x_i^* . To see this clearly, note that $x_i^* \uparrow \implies m'(x_i^*) \downarrow \implies F(c/m'(x_i^*)) \downarrow \implies f(F(c/m'(x_i^*))) \downarrow$. □

Proof of Proposition C.2. For some fixed parameters g, α, c, b and functions $f(\cdot), m(\cdot)$ there is a unique equilibrium x^* . **(i)** Suppose that $c \uparrow$. This decreases RHS of the FOC (Equation (C.1)). To restore equality, it must be that $\hat{x}_i \downarrow$ (as this decreases LHS and increases RHS). In turn $\hat{x}_i \downarrow \implies \hat{x}_j \downarrow$. Remark C.1 guarantees a unique equilibrium. So this process must eventually

converge to an equilibrium. But all equilibrium actions strictly decrease. **(ii)** Equilibrium utility is:

$$u_i^* = f\left(F\left[\frac{c}{m'(x_i^*)}\right]\right) - cx_i^*.$$

This was derived in the proof to Proposition C.1. It is clear that a sufficient condition for u_i^* to be strictly increasing in c is that: (a) cx_i^* and (b) $c/m'(x_i^*)$ are strictly decreasing in c (recall that $f(F(\cdot))$ is decreasing in its argument)³⁹ Consider each condition separately. First,

$$\frac{d(cx_i^*)}{dc} = c \cdot \frac{dx_i^*}{dc} + x_i^*,$$

which is less than zero if and only if $\frac{dx_i^*}{dc} < -x_i^*/c$. Second,

$$\frac{d(c/m'(x_i^*))}{dc} = [m'(x_i^*)]^{-1} - cm''(x_i^*)[m'(x_i^*)]^{-2} \frac{dx_i^*}{dc}$$

which is less than zero if and only if $\frac{dx_i^*}{dc} < -m'(x_i^*)/[c \cdot m''(x_i^*)]$. □

An equivalent sufficient condition for u_i^* to be strictly decreasing in c would be easy. The method for proving it would be similar to above. The condition would simply be $\frac{dx_i^*}{dc} > \max\left\{\frac{-x_i^*}{c}, \frac{m'(x_i^*)}{c \cdot m''(x_i^*)}\right\}$. Due to the inability to find closed form solutions for x_i^* when $m(\cdot)$ is nonlinear, I am not able to provide a tight link between network centrality and how welfare is affected by cost changes. Nevertheless, the flavour of Proposition 2 goes through – agent's whose consumption is highly sensitive to costs end up experiencing welfare gains when the marginal cost of consumption rises. This is again because the benefits (in terms of a lower reference point) of lower consumption by neighbours more than offsets the directly higher cost of consumption.

³⁹So a decrease in $c/m'(x_i^*)$ increases $f(F(c/m'(x_i^*)))$.

D Loss Aversion

In section 1 we assume that $u_i = f(x_i - \alpha_i \sum_j g_{ij} x_j)$, where $f(\cdot)$ is strictly increasing and concave. This captures references dependence and diminishing sensitivity, but *not* loss aversion. To capture loss aversion, we now assume that $f(\cdot)$ is strictly concave [resp. convex] in the positive [resp. negative], and kinked at zero. This is in line with the Kahneman and Tversky's canonical setup [Tversky and Kahneman, 1979, 1991]. Formally, we can write these properties as the following restrictions on the function $f(\cdot)$: (i) $f : a \rightarrow \mathbb{R}$ for $a \in \mathbb{R}$, (ii) $f'(a) > 0 \forall a$, (iii) $f''(a) < 0 \forall a > 0$, $f''(a) > 0 \forall a < 0$, (iv) $\lim_{a \rightarrow 0^-} |f'(a)| > \lim_{a \rightarrow 0^+} |f'(a)|$, (v) $f''(a) \in \mathbb{R} \forall a \neq 0$, $f(0) = 0$.

Adding the functional form requirements for loss aversion does not affect Remark 1. This is because agents are always able to, and always choose to, consume above their reference point in this model. Therefore the convexity in the negative domain that is the core addition of loss averse preferences (over and above reference dependent preferences) has no bite.

E Multiple Goods

In section 1 we assume that there is only 1 good. Here, consider the model with K goods (i.e. a K -dimensional consumption bundle), as in Kőszegi and Rabin [2006]. This generalisation has no effect.

There are K goods, x_1, \dots, x_K . Each agent i simultaneously chooses a consumption bundle $(x_{i1}, x_{i2}, \dots, x_{iK}) \in \mathbb{R}_+^n$. Following Kőszegi and Rabin [2006], I assume that preferences are additively separable over goods. All other elements of the model are the same as in section 1. Therefore i 's utility function is:

$$u_i = \sum_{k=1}^K \left[f \left(x_{ik} - \sum_j \alpha_i g_{ij} x_{jk} \right) - c_k x_{ik} \right] + b_i \sum_j \alpha_i g_{ij} \quad (\text{E.1})$$

where $f(\cdot)$ has the same properties as in section 1. Now consider agents' First Order Conditions.

$$\frac{du_i}{dx_{ik}} = f' \left(x_{ik} - \sum_j \alpha_i g_{ij} x_{jk} \right) - c_k = 0 \quad \text{for all } i \quad \text{and for all } k \quad (\text{E.2})$$

These First Order Conditions are clearly identical to those in the 1-good case. Obviously there are now K First Order Conditions for each agent, but x_{ik} only appears in the First Order Conditions relating to good k , and never in any relating to $k' \neq k$. Therefore the solution for each good k is the same as it would be if k were the only good.

F Network Structure

This section presents technical results regarding the effects of comparison shifts, and formalises claims made in section 2. The natural starting point is to characterise the exact impact of a single comparison shift.

Proposition F.1. *The change in agent i 's consumption due a comparison shift D , Δx_i^* , is equal to:*

$$\frac{\phi B_{ir}(C_u^b - C_d^b)}{1 - \phi(B_{ur} - B_{dr})} \cdot f'^{-1}(c)$$

Proof of Proposition F.1. This is a restatement of Lemma A.3 (which was needed to prove other results from the main text). \square

This exact characterisation is not very user-friendly and is largely technical. However, it is useful because it forms the basis for a number of further results, most importantly Proposition 3. Its implications for the effects of a comparison shift on different agents are immediate.

Corollary F.1. *Given a comparison shift D , the change in optimal actions is:*

- (i) *in the same direction for all agents,*
- (ii) *proportional to the amount that agent i compares herself to r (the subject of the comparison shift) prior to the shift.*

Proof of Corollary F.1. Proposition F.1 characterises the change in agent i 's optimal action, Δx_i^* . The only term in the expression for Δx_i^* that depends on i is B_{ir} . **(i)** Since $B_{ij} \geq 0$ for all i, j , Δx_i^* has the same sign for all i . **(ii)** Δx_i^* is equal to B_{ir} multiplied by some terms that do not depend on i . Therefore Δx_i^* is proportional to B_{ir} (i.e. the amount that i compares herself to r). \square

It also allows us to examine the nature of the returns to the magnitude of a comparison shift. When an agent r moves from direct comparison with a lower centrality agent to direct comparison with a higher centrality agent (i.e. $C_u^b > C_d^b$), then there are increasing returns to the magnitude of a comparison shift if and only if agent r influences agent u more than she influences agent d , but the gap is not too large.

Corollary F.2. *Given a comparison shift D , if $C_u^b > C_d^b$, then for all i , x_i^* is: convex in ϕ if and only if $B_{ur} - B_{dr} \in (0, \frac{1}{\phi})$, and concave in ϕ otherwise.*

Proof of Corollary F.2. For ease of exposition, let: $B_{ir}(C_u^b - C_d^b)f'^{-1}(c) \equiv y$ and $(B_{ur} - B_{dr}) \equiv z$. Let $\hat{x}_i^* \equiv x_i^* + \Delta x_i^*$ be the new equilibrium value of consumption. So $\hat{x}_i^* = x_i^* + y\phi(1 - z\phi)^{-1}$. Now consider the second derivative. $\frac{d^2 \hat{x}_i^*}{d\phi^2} = 2yz(1 - z\phi)^{-3}$. We can partition values of z into three cases: (i) $z < 0$, (ii) $0 \leq z \leq \frac{1}{\phi}$, (iii) $z > \frac{1}{\phi}$.

In case (ii), $\frac{d^2 \hat{x}_i^*}{d\phi^2} \geq 0$. In cases (i) and (iii), $\frac{d^2 \hat{x}_i^*}{d\phi^2} < 0$. These observations follow straightforwardly from the fact that $y > 0$ and $\phi > 0$. The second derivative determines convexity/concavity. Note that if $(C_u^b - C_d^b) < 0$ then $y < 0$ and so all results flip. \square

Now consider a *composite comparison shift*, $\hat{D} = D_1 + \dots + D_Z$, which is constructed by summing up $n \geq 2$ comparison shifts. The impact of a composite comparison shift is not equal to the sum of the effects of each comparison shift that form part of it. This is because the impact of any one comparison shift depends on the *whole* network immediately before it occurs.

Corollary F.3. *Consider a composite comparison shift \hat{D} . The impact of $\hat{D} = D_1 + \dots + D_Z$ is not the same as the sum of the impacts of D_1, \dots, D_Z .*

Proof of Corollary F.3. W.L.O.G any composite comparison shift \hat{D} can be expressed as the sum of comparison shifts D_1, \dots, D_Z : $\hat{D} \equiv \sum_{i=1}^Z D_i$. It follows from Lemma A.2 that H is a function of both the comparison shifts D_i and $(I - G)^{-1}$.

Trivially $(I - G - X)^{-1} \neq (I - G)^{-1}$ for any $X \neq 0$. Therefore the effect of a comparison shift D_i depends on the network immediately prior to the comparison shift: $H(D_i, (I - G - \sum_{j=1}^{i-1} D_j)^{-1}) \neq H(D_i, (I - G)^{-1})$. The effect of a composite comparison shift \hat{D} is therefore not equal to the sum of the effects of individual comparison shifts: $\sum_{i=1}^Z H(D_i, (I - G - \sum_{j=1}^{i-1} D_j)^{-1}) \neq \sum_{i=1}^Z H(D_i, (I - G)^{-1})$. \square

For a given network, it would be straightforward to calculate the change in actions due to a more complex change in the reference structure using an application of the formula provided by Chang [2006]. However, we cannot obtain analytic results due to the interactions between the effects of each comparison shift. Even with a few comparison shifts, the outcome would be too complicated to yield any insight. However, if the comparison shifts are all of an equal size, for example Δ , then interactions between the are of the order Δ^2 . Therefore, when considering small changes to the network (i.e. small $Z \cdot \Delta$) we can reasonably disregard the interactions – a naive summation is a close approximation for the actual aggregate effect.

Corollary F.4. *If the total change to the network, $Z \cdot \Delta$, is small, then the impact of a composite comparison shift \hat{D} approximately equal to the sum of the impacts of each comparison shift that makes up \hat{D} .*

Proof of Remark F.4. Let the composite comparison shift \hat{D} be constructed from a series of comparison shifts (D_1, \dots, D_Z) , each of a size Δ . $H(D_i, (I - G - X)^{-1})$ is the change in the optimal actions following a comparison shift D_i , when the starting network is $(G + X)$. From Lemma A.2: $H(D_i, (I - G)^{-1}) = \Delta \cdot fn((I - G)^{-1}) = \mathcal{O}(\Delta)$. Then $H(D_i, (I - G - D_j)^{-1}) = \Delta \cdot fn((I - G - D_j)^{-1}) = \Delta \cdot fn((I - G)^{-1} + \mathcal{O}(\Delta)) = H(D_i, (I - G)^{-1}) + \mathcal{O}(\Delta^2)$. By a simple induction argument we can see that

$H(D_i, (I - G - \sum_{j=1}^J D_j)^{-1}) = H(D_i, (I - G)^{-1}) + \mathcal{O}(\Delta^2)$. Therefore the effect of the earlier comparison shifts D_1, \dots, D_{i-1} on the change in optimal actions induced by D_i is on the order Δ^2 . If $Z \cdot \Delta$ is sufficiently small, then we can ignore these interaction effects, which are collectively of the order $Z \cdot \Delta^2$. \square

This approach works because we are able to take a linear approximation (i.e. ignore any terms that are $\mathcal{O}(\delta^2)$). The total change to agents' Bonacich centrality will be well approximated by simply adding up the effects of each comparison shift. However, we should be very wary of extrapolating from small changes to large ones. It is difficult to characterise the interactions and small changes may not be indicative of large ones. Past observations from small or localised changes may cease to be a useful guide in the face of large-scale social change happens.

Finally, we prove that assuming $\alpha_i < 1$ for all i is sufficient to guarantee existence, even after a comparison shift (something we claimed, but did not prove, in section 2).

Remark F.1. *If $\alpha_i < 1 \forall i \in N$ then $(I - (ag + D))^{-1}$ exists for any network ag , and any comparison shift D .*

Proof of Remark F.1. This follows trivially from Proposition 1, which proves that $\alpha_i < 1$ for all i is a sufficient condition to ensure equilibrium existence for any network G . Since a comparison shift leaves α_i unchanged for all i (by definition), then Proposition 1 continues to apply. \square

However, if $\alpha_i > 1$ for some i , then it is possible that a solution exists for some, *but not all*, reference structures g . In this instance it is necessary to check that a solution exists both before and after the shift. That is: $\lambda_1(G) < 1$ and $\lambda_1(G + D) < 1$. While it is possible for only one of these conditions to be met, the results cannot apply unless both hold.

G Chang (2006): A perturbation theorem

I restate Lemma A.2 for clarity and then present a proof. This is a special case of Chang [2006], which simplifies the proof. I have also aligned the notation to match my paper.

Lemma (A perturbation theorem. Chang [2006]). *If D is a comparison shift, then $(I - [G + D])^{-1} - (I - G)^{-1} = H$. Where:*

$$H = \frac{\phi}{1 - \phi(B_{ur} - B_{dr})} \begin{pmatrix} B_{1r}(B_{u1} - B_{d1}) & \cdots & B_{1r}(B_{un} - B_{dn}) \\ \vdots & \ddots & \vdots \\ B_{nr}(B_{u1} - B_{d1}) & \cdots & B_{nr}(B_{un} - B_{dn}) \end{pmatrix}$$

Proof of Lemma A.2. This proof is a simplified version of Chang [2006]. By the Woodbury Identity Matrix: if A and D are $n \times n$ matrices, and A is non-singular, then $(A - D)^{-1} = E + E(I - DE)^{-1}DE$, where $E \equiv A^{-1}$ [Woodbury (1950) and Sherman and Morrison (1949)]. Now partition the matrices D and E :

$$D = \begin{bmatrix} \overline{D} & 0 \\ 0 & 0 \end{bmatrix} \quad E = \begin{bmatrix} \overline{E} & E_2 \\ E_1 & E_3 \end{bmatrix} \quad \overline{E} = \begin{bmatrix} \overline{E} \\ E_1 \end{bmatrix} \quad \underline{E} = \begin{bmatrix} \overline{E} & B_E \end{bmatrix}$$

where \overline{D} is the smallest matrix that contains non-zero elements of D , and \overline{E} contains the transpose of the elements in \underline{E} .⁴⁰ Simple matrix algebra yields: $(A - D)^{-1} = E + \overline{E}(I - \overline{D} \overline{E})^{-1} \overline{D} \underline{E}$.

Now recall that D is a ‘comparison shift’ (as per Definition 2). Therefore: $\overline{D} = \begin{bmatrix} D_{ru} & D_{rd} \end{bmatrix} = \begin{bmatrix} \phi & -\phi \end{bmatrix}$, and so:

$$\overline{E} = \begin{bmatrix} E_{ur} \\ E_{dr} \end{bmatrix}, \quad \overline{E} = \begin{bmatrix} E_{1r} \\ \vdots \\ E_{nr} \end{bmatrix} \quad \text{and} \quad \underline{E} = \begin{bmatrix} E_{u1} & \cdots & E_{un} \\ E_{d1} & \cdots & E_{dn} \end{bmatrix}$$

This yields;

$$(A - D)^{-1} = E + \begin{bmatrix} E_{1r} \\ \vdots \\ E_{nr} \end{bmatrix} \cdot \left(I - \begin{bmatrix} D_{ru} & D_{rd} \end{bmatrix} \begin{bmatrix} E_{ur} \\ E_{dr} \end{bmatrix} \right)^{-1} \cdot \begin{bmatrix} D_{ru} & D_{rd} \end{bmatrix} \cdot \begin{bmatrix} E_{u1} & \cdots & E_{un} \\ E_{d1} & \cdots & E_{dn} \end{bmatrix}$$

Now multiply out, substitute in $D_{ru} = \phi$ and $D_{rd} = -\phi$, noticing that the matrix inverse $(I - \overline{D} \overline{E})^{-1}$ is a scalar, and rearrange.

$$(A - D)^{-1} - A^{-1} = \frac{\phi}{1 - \phi(B_{ur} - B_{dr})} \begin{pmatrix} B_{1r}(B_{u1} - B_{d1}) & \cdots & B_{1r}(B_{un} - B_{dn}) \\ \vdots & \ddots & \vdots \\ B_{nr}(B_{u1} - B_{d1}) & \cdots & B_{nr}(B_{un} - B_{dn}) \end{pmatrix}$$

This result holds for any non-singular $n \times n$ matrix A . Letting $A = (I - G)$ yields the result. \square

⁴⁰So if \overline{D} consists of elements \overline{D}_{ij} for $i \in I$, $j \in J$, then \overline{E} consists of elements \overline{D}_{ji} for $i \in I$, $j \in J$.

H Heterogeneous Costs

The cost parameter c_i reflects agents' underlying propensity/ability to take the action, x_i . In this section, we examine the implications of introducing cost heterogeneity. The key takeaway is that allowing for heterogeneous costs has relatively little impact on the insights of the main model.

With this generalisation, equilibrium play depends on the individual entries in the full matrix B , rather than only on Bonacich centralities (the row sums of B). Nevertheless, the existence condition is unaffected, and the solution takes a similar form. At this point, it is helpful to interpret individual elements of the matrix B and to introduce a notion of “generalised Bonacich centrality”.

Definition H.1 (Comparisons). *The comparison matrix is $B \equiv (I - G)^{-1}$, where B_{ij} captures how much i compares herself to j .*

Bonacich centrality captures an agent's connectedness to the network as a whole. The comparison matrix breaks this down to the individual agent level. An element B_{ij} measures the total weight of walks from i to j , and captures the extent to which i compares herself to j . This is a dis-aggregation of Bonacich centrality. We can then weight these individual level comparisons with a function of the cost parameters to obtain a generalised notion Bonacich centrality.⁴¹

Definition H.2 (Generalised Bonacich Centrality). *The vector of Bonacich centralities for a network $G \equiv ag$ is $C^{gen} = B \cdot f'^{-1}(c)$. The Bonacich centrality of agent i is $C_i^{gen} = \sum_j B_{ij} f'^{-1}(c_j)$*

With this definition we can restate Remark 1, accounting for heterogeneous costs. The condition for existence depends only on the network and so is unaffected by cost heterogeneity. However, the optimal actions are now proportional to our new notion of generalised centrality, rather than the usual Bonacich centralities.

Proposition H.1 (Existence and Solution). *With heterogeneous costs, an equilibrium exists if and only if $\lambda_1 < 1$. If this condition is met, then there is a unique Nash Equilibrium: $x_i^* = C_i^{gen}$*

Proof. This follows from the proofs to Proposition 1, with the only change that $f'^{-1}(c_j)$ now depends on j , and so cannot be pulled out through the summation sign. \square

The other results from Section 2 also extend to this heterogeneous cost setting. The proofs here only provide the required extension from their homogeneous cost analogues.

Proposition H.2 (Reference strength). *If $\lambda_1 < 1$: (i) x_i^* is weakly increasing, and (ii) u_i^* is weakly decreasing, in α_j for all i, j , and strictly so if $i = j$.*

Proof. **(i)** Proposition 1 shows that if a given α_j increases, then all elements of $(ag)^k$ weakly increase for any k . Consequently, all elements $(I - ag)_{ij}^{-1}$ increase. **(ii)** having

⁴¹It is clear that when $c_j = c$ for all j this collapses back to the original Bonacich centrality.

found that x_i^* is increasing in α_k , the second part of the proof to Proposition 1 goes through unchanged. \square

Proposition H.3 (Cost). *If $\lambda_1 < 1$: (i) x_i^* is strictly decreasing and convex in c_j for all j , (ii) u_i^* is strictly increasing in c_j for all $j \neq i$.*

Proof. (i) follows straightforwardly from Proposition H.1 and the fact that $F(\cdot)$ is strictly decreasing and convex. (ii) equilibrium utility is $u_i^* = f(x_i^* - \sum_{k \neq i} G_{ik} x_k^*) - c_i x_i^*$. Differentiate with respect to c_j :

$$\frac{du_i^*}{dc_j} = \left(\frac{dx_i^*}{dc_j} - \sum_{k \neq i} G_{ik} \frac{dx_k^*}{dc_j} \right) f'(\cdot) - c_i \frac{dx_i^*}{dc_j}$$

Now recall that $f'(x_i^* - \sum_j x_j^*) - c_i = 0$ in equilibrium, and that $x_i^* = C_i^{gen}$ for all i , $\frac{dx_k^*}{dc_j} = \frac{dC_k^{gen}}{dc_j} = B_{kj} F'(c_j) < 0$. Substituting these in yields:

$$\frac{du_i^*}{dc_j} = -f'(\cdot) \sum_{k \neq i} G_{ik} B_{kj} F'(c_j) < 0$$

\square

This result is somewhat different to the homogeneous cost version. Because the cost parameter is now agent-specific, the outcome is much simpler. An increase in an agent j 's cost pushes down her consumption, relaxing the need for others to keep up with "The Joneses" (in this case, agent j). This increases welfare for all $i \neq j$. Since i has not experienced an increase in her own costs, there is no off-setting effect. The impact of someone else's cost parameter on your welfare is unambiguous.

Proposition H.4. *The change in agent i 's action due a comparison shift D is equal to;*

$$\frac{\phi B_{ir} (C_u^{gen} - C_d^{gen})}{1 - \phi (B_{ur} - B_{dr})}$$

Proof. This follows from Proposition F.1 but replacing $C_i^b \cdot f'^{-1}(c)$ with C_i^{gen} . \square

All other results concerning comparison shifts also follow as a result of this, including an analogue to Proposition 3. This is because they are also based on Proposition F.1 in the homogeneous cost case.

Proposition H.5 (Endogenous network). *In all pairwise stable networks, if $b_i \geq c_i f'^{-1}(c_i)$, then $G_{ij} > 0$ only if $\frac{b_i}{c_i} = \frac{b_j}{c_j}$.*

Proof. This follows straightforwardly from the proof to Proposition 4, replacing c with the agent-specific version as appropriate. \square