

Being up front about Income Inequality*

Brian Hill
CNRS & HEC Paris[†]

October 15, 2020

Abstract

This paper proposes universal provision of information about the income inequality involved in the creation of a good as a means of moderating income inequality. Existing evidence suggests that a section of the population would be willing to pay more for goods whose production involves less excessive income inequality. We show, on a simple model, that supplying inequality information to such a population under competitive markets will in general lead to a reduction in global income inequality. The effect will be stronger the more inequality averse the population. Moreover, the outcome will be socially efficient. Possibilities for (de-centralized) implementation are also discussed.

Keywords: Income inequality, information provision, externalities, inequality aversion.

*Acknowledgements and thanks to be added.

[†]GREGHEC, 1 rue de la Libération, 78351 Jouy-en-Josas, France. E-mail: hill@hec.fr.
URL: www.hec.fr/hill.

1 Introduction

Inequality is an increasingly pertinent, debated and topical issue. Many studies have documented the rise in income and wealth inequality over the past 40 years (e.g. Acemoglu and Autor, 2011; Piketty, 2013). According to some commentators, this rise is not without wider consequences, being relating to a range of recent social and political phenomena.

Income inequality, moreover, is not an issue which leaves the man on the street indifferent. Surveys (e.g. Kohut, 2013; Antunez and Papuchon, 2019) document that many citizens consider it an important issue. Studies suggest that, whilst people agree that some income inequality may be justified by skill, merit, desert or effort, most do not think that current levels of inequality in, say CEO vs. worker pay, can be justified on such grounds.¹ This adds up to a situation in which a significant section of the public seems uncomfortable with current levels of inequality.

The mismatch between public attitudes and current levels of income inequality suggests that it may be an externality.² This chimes well with the platitude about markets not being able to properly incorporate considerations of inequality. Indeed, similar economic tools have been proposed to ‘combat’ undesirable inequalities and ‘correct’ externalities: taxation is a notable example. Considering income inequality as an externality, this paper sets out the theoretical case for a complementary, underused, potentially simpler tool: information. The proposal is to inform potential consumers of every good or service, at the point of purchase, of the income inequality across all those involved in the conception, production, financing, marketing and logistics leading to the existence of that good or service on the market. We also discuss implementation, suggesting that such information need not be mandated by regulation, but can be provided via a mobile phone application.

We theoretically evaluate the potential impact of blanket product-related inequality reporting in a simple competitive-market model. Our first main result indicates the effectiveness of inequality information provision in combatting inequality. It implies that whenever there is a proportion of consumers that are *inequality averse*—they are willing to pay a positive sum to avoid extreme inequality in the creation of the purchased good—inequality information provision will lead to reduction in global income inequality. As we shall discuss, existing experimental and field evidence, including that in a sister paper, points to a significant presence of inequality aversion among consumers. Our second central result shows that, whenever markets are competitive, the market equilibrium under information provision will be socially efficient: the resulting allocation of goods and inequality levels is Pareto optimal with respect to consumers’ preferences. If anything, information provision is efficiency enhancing as an intervention, rather than the contrary.

These properties complement and tie into other qualities of information provision that arise from it being an essentially market-based mechanism. From an economic point of view, this renders it a relatively ‘light’ measure, in several senses. It is non-invasive, unlike taxes, and fully libertarian, insofar as it allows consumers and firms full freedom in their choice of purchase, production

¹See Section 2 for a brief survey of some of this literature.

²The potential externality due to inequality has a long history, dating back at least as far as Thurow (1971).

and investment, and ultimately their contribution to inequality. Aside from the information intervention, it leaves markets unaffected, which is key to obtaining social efficiency. Moreover, it has interesting consequences for the trade off between inequality against productivity, which is at the heart of policy choice regarding income inequality. The public’s opinions on this trade-off are seldom explicitly sought and rarely taken into account; typically, it is either implicitly imposed by fiat (e.g. in a social welfare function) or tacitly assumed to be given ‘clearly enough’ by some widely-shared intuition (as when actors point to the productivity costs of an inequality gain as ‘evidently’ too large or too small, depending on their leanings). Information provision avoids these sticky issues, by directly and explicitly leveraging consumer preferences regarding productivity-inequality trade-offs, and incorporating them in equilibrium by standard market mechanisms. If productivity is lower under information provision (as will typically be the case), it is because sufficiently many consumers consider that the putative productivity gains—and in particular their impact on prices—are outweighed by the accompanying inequality costs.

The proposal also has a behavioral angle. On the one hand, reporting production inequalities for all goods and services makes the trade-off between monetary gain (for a cheaper product) and social responsibility (for one that does not exacerbate inequality) clear and salient for consumers. And behavioral economics has taught that attention is an important driver of choice (Kahneman and Tversky, 2000). On the other hand, universal or blanket reporting contrasts with various sorts of ‘labels’—for instance, Corporate Social Responsibility (CSR), Fair Trade or ecological labels—where it is the companies who decide whether to subscribe to them or not. Importantly, unlike optional labelling, blanket reporting identifies and highlights the ‘bad’ cases. Going by the significant psychological evidence that negative cases tend to incite stronger reaction than positive ones (Kahneman et al., 1991; Rozin and Royzman, 2001; Baumeister et al., 2001), one might expect blanket reporting to have a stronger effect on consumption choices.³

Beyond their purely economic advantages, information-based interventions may also have a political dimension. If there are significant inequalities then, simple reasoning suggests, it is because somewhere down the line consumers are buying products or services where the money they spend ends up unequally distributed among those involved in the creation of the product. Informing consumers of the inequality in this distribution can be thought of as a way of empowering them: if they don’t like it, they can immediately and individually do something about it, by altering their consumption choices. Empowerment via purchase decisions is not a new idea; however, once again the proposal of *blanket* reporting suggests that it can be more systematic than, say, one-off boycott campaigns. Of course, with empowerment comes responsibility, and inequality reporting will allow consumers to get a better idea of their responsibilities in the

³Research into existing voluntary labels generally tends to find a willingness to pay among a section of consumers for Fair Trade labeled products, for instance (e.g. Andorfer and Liebe, 2012), as also attested by a market for such products worth around \$8.5 billion in 2017 (<https://www.fairtrade.net/impact/global-sales-overview>). The suggestion here is that when information about the parameter of interest is universally available—saying how bad the inequality (or how unfair the trade) is for all products—then the impact can only be larger. Note that the universal character of the information contrasts with the potential information asymmetry involved in voluntary labelling; see Section 5 for more on the relationship between the current proposal and existing approaches in economics.

current state of inequality, and to do something about it. Finally, survey data reveal that many people significantly underestimate the real level of inequality (see Section 2); blanket reporting of inequality information may thus also help raise awareness of the issue and correct misperceptions, helping citizens develop informed opinions on it.

This voluntarily programmatic paper sets out the case for the theoretical promise and practical feasibility of blanket reporting of the income inequalities involved in product creation. It first quickly reviews existing literature suggesting that people are willing to pay to avoid very large inequalities (Section 2). It then shows, on a simple economic model, how such preferences induce reduced income inequality under universal provision of inequality information. Moreover, the resulting inequality distribution is optimal, by the lights of consumers (Section 3). It discusses possibilities for implementation, drawing some general principles and making some tentative suggestions (Section 4). The paper ends situating the proposal with respect to some parts of the existing economic literature, and considering the potential of similar interventions for issues other than income inequality (Section 5).

2 Attitudes to and opinions about income inequality: current evidence

A range of surveys, across of variety of countries, tells a consistent and by now well-documented story concerning opinions about and attitudes to income inequality (e.g. Osberg and Smeeding, 2006; Kiatpongsan and Norton, 2014; Clark and d'Ambrosio, 2015). Figure 1 presents typical findings concerning the pay ratio between CEOs of large national corporations and unskilled workers—which is a widely used measure in this literature, and serves as a reasonable proxy for income inequality for products produced by large firms. The Figure illustrates three central messages.


Firstly, people have definite views on the ideal pay ratio, and these are not reductive. In particular, most people think that some inequality is justified.⁴ This is consistent with them considering certain inequality levels to be acceptable on the grounds of fairness, merit or desert, as well as other factors (Fehr and Schmidt, 2003; Almås et al., 2010). Nevertheless, many state moderate levels of ideal inequality (the median ideal ratio in the global sample graphed in Figure 1a is 4.6 : 1, Kiatpongsan and Norton 2014; in the French sample in Figure 1b, it is 7.7 : 1).

Secondly, they generally think, correctly, that the actual pay ratio is higher than their ideal value: in the French data graphed in Figure 1b, 92% of subjects estimated the pay ratio as higher than their ideal value.

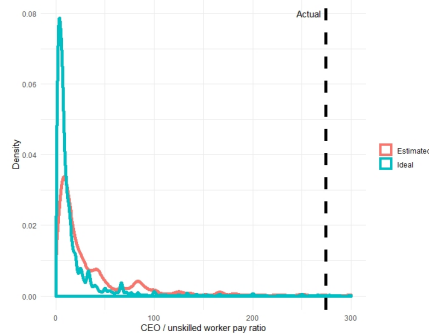
Thirdly, people grossly underestimate the extent of pay inequality, generally by a factor of ten or more. For instance, in the French data, the median estimated pay ratio was 20:1, whereas the actual pay ratio was closer to 270:1. 93% of participants underestimated the pay ratio.

These facts suggest the potential effectiveness of an information-based intervention focussed on income inequality. The first two reveal that people have

⁴For instance, in the French data graphed in Figure 1b, only 94 out of 7762 subjects stated that a CEO should ideally earn the same or less than an unskilled worker.



(a) CEO to unskilled worker pay ratio across 16 countries (2009). Figure drawn from Kiatpongsan and Norton (2014). The centre of the diagram represents a pay ratio of 1:1, the outermost ring represents a pay ratio of 351:1. Readers wishing a zoom in on the estimated and ideal ratios (the red area and the barely visible blue area in the middle) are referred to Kiatpongsan and Norton (2014).



(b) CEO to unskilled worker pay ratio in France (2018).

The Figure plots the distribution of estimated and ideal CEO-to-unskilled worker pay ratios across participants in the French DREES 2018 survey (DREES, 2018; DREES is the French Centre for Research, Study, Evaluation and Statistics, under the auspices of the Health and Social Affairs Ministry). The actual value plotted is calculated from the average net salary of CEOs among CAC40 companies (CAC40 is the benchmark French stock market index), as reported in Lacombe (2019), and the 2015 average net salary of worker in France, according to the latest (2015) INSEE data (Tavernier, 2019), corrected for inflation.

Figure 1: Actual, estimated and ideal pay ratios of CEOs to unskilled workers

views on acceptable levels of inequality, which could inform their purchasing decisions. Moreover, they are aware that these views may not correspond to reality. The final fact—that people are badly misinformed about current levels of inequality—suggests that the information involved in the intervention is not already possessed by many people, hence leaving space for significant effects of information provision.

These suggestions receive further support from experimental studies. Mohan et al. (2018) provide evidence from an incentivised field experiment suggesting a significant proportion of subjects opt to purchase from retailers with low CEO-to-average employee salary ratio (around 5 : 1) over retailers with a significantly higher ratio (of the order of 700 : 1). Mohan et al. (2015) run a series of non-inventive-compatible studies suggesting that both the willingness to buy from a firm and the amount subjects are willing to pay for a good are impacted by inequality information, with subjects willing to pay of the order of \$2 to \$5 more on a \$25 good for a reduction in inequality of a factor of 200. Hill and Lloyd (2020), in an incentivised, behavioural choice experiment run on a representative sample of the English population, find a significant impact of the inequality involved in the production of a good on willingness to pay for the vast majority of subjects, with a mean willingness to pay of over £10 on a £30 good for a reduction in CEO-to-median worker pay ratio from 750 : 1 to 5 : 1.

Many of the cited studies also examine variability in ideal and estimated inequality as well as willingness to pay across demographics (e.g. age, political beliefs). They find that the variation is generally not more than by a factor of 2, and that the previous points describe fairly accurately a significant proportion of the population.

The cited studies focus uniquely on ‘disinterested’ opinions and preferences concerning income inequality, and hence on what has been called peoples’ ‘normative’ evaluation of inequality—whether such or such inequality is ‘right’ or justified, independently of how it affects them (Clark and d’Ambrosio, 2015). Another strand of the literature focusses on ‘comparative’ or positional evaluation, where individuals’ preferences are sensitive to how they compare to others (and hence where they lie on the income distribution). A large literature on such interdependent preference establishes their existence and significance in some situations (e.g. Fehr and Schmidt, 2003; Clark and d’Ambrosio, 2015). Whilst the former literature is the more relevant for the proposal in this paper—which is to inform consumers of the inequality involved in the production of a good, not of how they situate with respect to it—the latter literature of course provides some empirical corroboration of the relevance of inequality considerations for preferences.

Note finally that, beyond the realm of purchasing choices, there is evidence of potential impacts of information on attitudes to public policies or public issues: for instance, Kuziemko et al. (2015) provide survey data suggesting a sensitivity of concern about inequality, as well as redistributive preferences, to information.

3 The impact of inequality information: a model

The evidence just reviewed suggests that a significant proportion of the population may be willing to pay moderate amounts for a reduction in extreme levels of income inequality among those involved in the creation of the products they purchase. We shall say that such consumers exhibit *inequality aversion*. Note that individuals that are inequality averse in this sense *fully accept* inequalities that can be justified by considerations such as merit, desert or fairness; they are *only* averse to those inequality levels that cannot, in their opinion, be justified on such grounds. For a population of consumers, some of which are inequality averse, what impact would providing information on the income inequality across all those involved in the creation of a good or service, for every product on the market, have on income inequality? In this section, we show, on a simple model, that such an information intervention will generally reduce income inequality, will never cause it to rise, will be stronger the more inequality averse the population, and will be socially optimal. We first set out the model and establish the main theoretical results, before turning to what they say about the information intervention in Section 3.5.

3.1 Economic Model

We limit attention to a single good, which can be produced with varying amounts of inequality, and consider two markets: the market for the good of interest (potentially with inequality information) and the market for the production factors

(e.g. labour, capital) contributing to the creation of the good. Firms ‘recruit’ production factors on the latter market to produce the good, which they then sell to consumers on the former market.

Consumers The goods on the good market are identical except for their inequality levels, which take values in a positive real interval I . Consumers are price-takers, where the price of the good may only depend on the inequality in production, of which they are fully and correctly informed. Each consumer can purchase zero or one unit of the good, where the choice is whether to purchase, and at which inequality level. A consumer j ’s preferences are represented by a utility function $U_j : (I \cup \{\}) \times \mathbb{R} \rightarrow \mathbb{R}$ of the following quasi-linear form:

$$U_j(i, n) = \begin{cases} n + (v_j - \psi_j(i)) & i \in I \\ n & i = \{\} \end{cases} \quad (1)$$

where n is the quantity of the numéraire (e.g. money), v_j —a positive real number—is the ‘intrinsic’ value of (one unit of) the good for the consumer, independent of inequality considerations, and ψ_j —an increasing real function of I —is the disutility of obtaining the good with inequality score i . $\{\}$ represents the outcome of not obtaining the good, which is normalised to have utility 0. So, for instance, a consumer with endowment \hat{n} of the numéraire who pays $p(i)$ to purchase the good produced at inequality level i obtains utility

$$\hat{n} - p(i) + v_j - \psi_j(i) \quad (2)$$

The elements v_j and ψ_j are characteristics of consumer j . In particular, the latter reflects her willingness to pay to avoid (a given degree of) inequality. Given the choice between a unit of the good offered at inequality i' and a unit offered at higher inequality i'' , such a consumer would pay a premium of up to $\psi_j(i'') - \psi_j(i')$ for the low-inequality good.

In the light of the evidence set out in Section 2, we assume that ψ_j has the following form:

$$\psi_j(i) = \begin{cases} 0 & i \leq \theta_j \\ \eta_j(i - \theta_j) & i > \theta_j \end{cases} \quad (3)$$

Consumers are insensitive to inequality scores below a threshold θ_j : such scores are considered possibly justified by considerations such as fairness or merit. This *justifiable-inequality threshold* can capture the fact, found in the empirical data (Section 2), that most people accept inequalities below a certain ‘ideal’ level as potentially justified. They become sensitive above that point: this is reflected in the second clause of (3), where higher inequality leads to higher disutility. Higher η_j reflects more inequality aversion: individuals are willing to pay more to avoid inequality. Consumers with $\eta_j = 0$ are insensitive or neutral towards inequality: they are not willing to pay anything to reduce even extreme inequalities.

Although we focus on the simple preference form (1) in the bulk of the paper, our results hold under more general additively-separable preferences (Appendix A). In particular, they do not require the disutility of inequality to be linear in the inequality level above the threshold (but can incorporate a wide range of

shapes), nor that it be additive with respect to the utility for the numéraire (incorporating, for instance, inequality disutility that is relative to the price of the good). Interested readers are referred to Appendix A.1 for further information.

To focus on the effect of inequality aversion, we assume that all consumers have the same v (intrinsic value of the good) and θ (justifiable-inequality threshold), but may differ in inequality aversion η . More specifically, we assume that there are v, θ and $K > 1$ levels of η , $\eta_1 > \dots > \eta_K = 0$ and say that a consumer is of type j if her utility is as in (1), with ψ_j as in (3) with (v, θ, η_j) . We consider a continuum $P = [0, N] \subseteq \mathbb{R}_{\geq 0}$ of consumers, with measure N .⁵ A sequence $\mu = (\mu_1, \dots, \mu_n)$ of positive real numbers for which $\sum_{k=1}^K \mu_k = N$ is called the *inequality aversion distribution* of the population. Each such μ represents the distribution of inequality attitude in the population: μ_j consumers have inequality aversion η_j , under μ . For a given inequality aversion distribution μ , P is partitioned into intervals $[0, \mu_1], \dots, (\sum_{l=1}^{j-1} \mu_l, \sum_{l=1}^j \mu_l), \dots, (\sum_{l=1}^{K-1} \mu_l, N]$, with j th interval containing all and only consumers of type j (and accordingly, this interval is of measure μ_j). This model can thus account for the varying inequality attitudes across the population noted in Section 2; indeed, we shall consider how equilibrium income inequality is impacted by the inequality aversion distribution μ .

Finally, to focus on the effect of inequality aversion, we assume that all consumers have the same endowment of \hat{n} units of the numéraire. So each consumer of type j chooses (i, n) to maximise (1) under the budget constraint $\hat{n} \geq p(i) + n$; for large enough v_i , she basically maximises (2).

Production factors To produce the good, each firm must use two *types* of production factors, which we call L and H . Each production factor may admit a range of *levels*. Income inequality in the production process is driven by the different rates at which the types and levels involved are remunerated. This setup can be interpreted in terms of production requiring two sorts of workers—for instance, high vs. low skill or more vs. less well educated (e.g. managers vs. factory workers)—with each sort of worker admitting different talent or skill levels (e.g. differences in the talent levels of factory workers or artisans, or of managers). In this case, the relevant difference is that between their salaries. Alternatively, L could represent labour, and H capital, with the levels in the latter case representing differences in the attractiveness of financing conditions, supply of financing and so on. Here, the difference would thus be the rates of return on labour and capital. Since firms will not make profits in equilibrium (because we consider perfect competition), H could even be taken to represent the ‘investment’ of shareholders or owners, with levels reflecting their input and its advantages; in this case the difference will be that between salaries and dividend returns to shareholders or owners. For the purposes of the exposition, we adopt the first interpretation, and speak of labour, wage differences, and so on.

For simplicity, we assume that there is only one level for the low type L , but a non-degenerate range of high-type H skill levels, taking values in the real interval $[s_H, \bar{s}_H]$.⁶ Each production factor type and level will be remunerated

⁵Throughout, we adopt the Lebesgue measure on P .

⁶So the model developed here is naturally read as focussing on managerial pay in relation to a fixed (e.g. median) salary, or capital returns in relation to a fixed salary income. A

at a unit wage which depends only on the type and level. We normalise wages (and prices) and set the wage for L -type labour to 1; $w(s)$ denotes the (unit) wage for a H -type worker of skill or talent level s .

To remain relatively non-committal on the supply side of the production factor (labour) market, we only assume, for each $s \in [\underline{s}_H, \overline{s}_H]$, a labour supply function X_s , where $X_s(x)$ is the supply of H -type s -level labour when the wage offered for that level is x . We assume that X_s is continuous in x and s , and strictly increasing in x throughout the range where it takes non-zero values, for each s . To translate the fact that H -type workers are interpreted as high earners relative to L -type workers, we assume that $X_s(1) = 0$ for all s : no H -type worker would work for the L -type wage.⁷

Firms Firms recruit labour (production factors) on the labour market, and use them to produce goods, which are sold on the good market. We do not focus on firm market power in either market, and suppose that they are ‘price-takers’ in both markets, which operate under perfect competition.

To produce the good, each firm recruits one unit of L -type labour and one unit of H -type labour at a single skill level: its only choice is thus the level of the unit of H labour recruited. Firms are fully and correctly informed of each worker’s type and level.

In equilibrium, there will be a wage assigned to each skill level: by the law of one price it will be unique and continuous in skills. We model these assignments by a wage schedule—a continuous function $w : [\underline{s}_H, \overline{s}_H] \rightarrow \mathbb{R}_{\geq 0}$.

The inequality involved in the production of the good by a firm choosing H -type level s is simply defined as the ratio between the wage paid for the H -type labour with skill level s — $w(s)$ under the wage schedule w —and the wage paid to the L -type worker—namely 1. So inequality in this model is fully characterised by $w(s)$. The set of inequality levels is thus $I = \mathbb{R}_{\geq 1}$.⁸ Consumers in the good market are fully and correctly informed of the production inequality for the good offered by each firm, namely of $w(s)$ for the s chosen by the firm.

Production is modelled by a continuous, twice-differentiable production function $F : [\underline{s}_H, \overline{s}_H] \rightarrow \mathbb{R}_{> 0}$: $F(s)$ is the quantity of the good produced with one unit of L -type labour and one unit of H -type labour of skill level s . We assume that skill is favourable to production: $F' > 0$. Firms choose the recruited skill level to maximise profits, so solve:

$$\max_{s \in [\underline{s}_H, \overline{s}_H]} p(w(s)) \cdot F(s) - (w(s) + 1) \quad (4)$$

Firms compete on both markets under perfect competition, with free entry. So equilibrium for an inequality aversion distribution μ is defined as a set of prices $p^* : I \rightarrow \mathbb{R}_{> 0}$, specifying the price for each inequality score in the good market, a wage schedule $w^* : [\underline{s}_H, \overline{s}_H] \rightarrow \mathbb{R}_{\geq 0}$ specifying the wage for each skill

similar analysis holds for the opposite case of a single H level and a range of L levels—which is most naturally read as a focus on low wages or labour exploitation. The general points continue to hold if several levels are incorporated for both types.

⁷The focus on the consumer side rather than the labour side sets our model apart from several existing studies of income inequality which concentrate on features of workers or suppliers in the labour market; see Section 5 for further discussion of the related literature.

⁸Recall that by the assumption on X_s , the H -type workers are always paid more than the L -type ones, whenever they are employed.

level in the labour market, and $J^* : [\underline{s}_H, \overline{s}_H] \rightarrow \mathbb{R}_{\geq 0}$ specifying the number (mass) of active firms recruiting at skill level s such that:

1. All firms maximise profits (4);
2. All consumers maximise utility (1) under the budget constraint $\hat{n} \geq p(i) + n$;
3. Closure of the labour market (labour demand = supply): $J^*(s) = X_s(w^*(s))$ for every $s \in [\underline{s}_H, \overline{s}_H]$;
4. Closure of the good market (good demand = supply): for every $1 \leq j \leq K$:

$$\mu_j = \int_{S_j} F(s) X_s(w^*(s)) ds \quad (5)$$

where S_j is the set containing every $s \in [\underline{s}_H, \overline{s}_H]$ such that consumers of type j buy goods of inequality $w^*(s)$ in equilibrium;

5. Free entry condition: for every $s \in [\underline{s}_H, \overline{s}_H]$, $p^*(w^*(s)) \cdot F(s) - (w^*(s) + 1) = 0$.

To simplify the analysis, we assume throughout that v (the ‘intrinsic’ value of the good) is large enough that all consumers purchase the good in equilibrium.

3.2 Solving the model: wage schedules in equilibrium

We show in Appendix A.2 that any equilibrium wage schedule is characterised by a sequence of positive real numbers C_k, \dots, C_K , for $0 \leq k \leq K$, and a sequence $\underline{s}_H \leq \underline{s}_k \leq \overline{s}_k \leq \theta \leq \underline{s}_{k+1} \leq \overline{s}_{k+1} \leq \dots \leq \underline{s}_K \leq \overline{s}_K \leq \overline{s}_H$ with:

$$w^*(s) = \begin{cases} \frac{C_K F(s) - 1}{\eta_K F(s) + 1} & s \in [\underline{s}_K, \overline{s}_K] \\ \frac{C_{K-1} F(s) - 1}{\eta_{K-1} F(s) + 1} & s \in [\underline{s}_{K-1}, \overline{s}_{K-1}] \\ \dots & \dots \\ \frac{C_{k+1} F(s) - 1}{\eta_{k+1} F(s) + 1} & s \in [\underline{s}_{k+1}, \overline{s}_{k+1}] \\ C_k F(s) - 1 & s \in [\underline{s}_k, \overline{s}_k] \end{cases} \quad (6)$$

where these sequences satisfy:

$$\mu_j = \int_{\underline{s}_j}^{\overline{s}_j} F(s) X_s \left(\frac{C_j F(s) - 1}{\eta_j F(s) + 1} \right) ds \quad (7)$$

for every $k < j \leq K$ and

$$\sum_{j=1}^k \mu_j = \int_{\underline{s}_k}^{\overline{s}_k} F(s) X_s (C_k F(s) - 1) ds \quad (8)$$

The interested reader is referred to Appendix A.2 for a detailed discussion of this wage schedule; here, we simply report some properties that will be relevant below. First, the wage w^* can be shown to be strictly increasing in the skill

level s . Second, the equilibrium wage schedule varies according to the inequality distribution in the population, μ . Finally, there is ‘sorting’ or ‘self-selection’ in the following sense: more inequality averse consumers will always buy from firms employing lower skilled workers. (Roughly, consumer of type j will buy from firms employing workers of skill level between \underline{s}_j and \bar{s}_j in equilibrium, and this yields the wage schedule of form (6).) However, the skill levels from which a consumer of given inequality aversion will purchase depends on the whole inequality aversion distribution μ , and not just on her individual inequality aversion.

Although the focus here is not on the existence or uniqueness of equilibria, note that a straightforward extension of Proposition 5 in Appendix A.3 implies that any equilibrium is unique.

3.3 Inequality aversion and inequality

Based on the previous characterisation, we can look at how inequality in the labour market varies with the inequality aversion of consumers. We use a simple proxy of inequality in the labour market: the ratio of the maximum to minimum wage among all workers employed, which we call the *max-min wage ratio*. In this simple model, where the minimal wage is fixed at 1, and where, as noted above, the equilibrium wage schedule is strictly increasing in s , this ratio is equal to the wage of the highest skilled employed H -type worker.

We are interested in the effect of (increased) inequality aversion in the population. To this end, let us say that an inequality aversion distribution μ *Inequality Aversion Dominates* another distribution μ' if, for every $1 \leq j \leq K$, $\sum_{i \leq j} \mu_i \geq \sum_{i \leq j} \mu'_i$. Recalling that lower j correspond to higher inequality aversion, this means that the proportion of the population having inequality aversion higher than a certain level is higher under μ than μ' . This is the standard notion of First Order Stochastic Dominance, applied to inequality aversion. We have the following result (see Appendix A.3 for the proof).

Theorem 1. *If μ Inequality Aversion Dominates μ' , then the max-min wage ratio in equilibrium is lower under μ . Moreover, it is strictly lower whenever $\mu_j \neq \mu'_j$ for some type j which, in equilibrium under μ' , buys the good at an inequality level greater than θ .*

This theorem gives a glimpse of the potential power of informing consumers about the inequality involved in the creation of goods. Any inequality-aversion increasing shift in the population—as long as it involves some consumers which are sensitive to inequality, in the sense that they purchase goods with inequality above the justifiable-inequality threshold—will reduce the income inequality in the labour (production factor) market. So, for instance, even if the proportion of inequality neutral consumers in the population remains the same, but inequality averse consumers become more inequality averse, this will drive down inequality across the board—and in particular the gap between the highest and lowest incomes.

Theorem 1 is based on a simple insight. Consumers with little or no inequality aversion tend to prefer cheaper, higher inequality goods, so they support a demand for highly productive workers, which *ceteris paribus* can produce goods at lower unit cost. When there are many such consumers in the population, this translates to a significant demand for higher skilled or more talented workers,

and drives up their wages. By contrast, inequality averse agents are willing to pay for lower inequality products, and so buy from firms employing lower skilled H -type workers. When a population has more of such consumers, this shifts the labour demand towards low skilled H -type workers and away from higher skilled ones. This deflated demand leads to a drop in the highest wages (and an increase in mid-range wages), and hence less income inequality.

3.4 Social Efficiency

One might worry that introducing inequality information may lead to social inefficiencies.⁹ Considering efficiency under the perfect competition, perfect inequality-information model set out above will allow evaluation of the extent to which inequality information can result in social inefficiencies. Indeed, since there is perfect competition, one might expect the First Welfare Theorem to hold, so that the equilibrium is Pareto optimal. Whilst it does not follow from the standard version of this theorem,¹⁰ it will nevertheless turn out that there is social efficiency: the equilibrium is Pareto optimal.

To consider the issue of social efficiency, we first define allocations in our markets. An allocation (in the goods market) is a pair consisting of

- a (measurable) function $c : P \rightarrow (I \cup \{\}) \times \mathbb{R}_{\geq 0}$ specifying, for each consumer $x \in P$, the inequality level of the good received (or $\{\}$ if no good is received), $c_1(x)$, and the quantity of the numéraire obtained, $c_2(x)$.
- a (measurable) function $q : I \rightarrow \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0}$ specifying, for each inequality level $i \in I$, the total quantity of the good produced with inequality i , $q_1(i)$, and the total cost of that production in numéraire terms, $q_2(i)$.

Moreover, the production allocation q must be generated from an assignment, to each skill level in the labour market, of the amount of firms hiring at that skill level and the wage offered, i.e.:

- there exists a (measurable) function $r : [s_H, \bar{s}_H] \rightarrow \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0}$, specifying for each skill level $s \in [s_H, \bar{s}_H]$, the amount of firms hiring workers of skill level s , $r_1(s)$, and the wage offered to workers of skill level s , $r_2(s)$, such that:

$$\begin{aligned} - q_1(i) &= \int_{\{s:r_2(s)=i\}} F(s)X_s(i)ds \text{ and} \\ - q_2(i) &= \int_{\{s:r_2(s)=i\}} (i+1)X_s(i)ds \text{ for all } i \in I. \end{aligned}$$

For an allocation to be *feasible*, it must satisfy the market clearing conditions, namely: for each $i \in I$

$$\int_{c^{-1}(i \times \mathbb{R}_{\geq 0})} dx = q_1(i) \tag{9}$$

(i.e. the total amount of good consumed at inequality level i is equal to the total amount produced) and

$$\int_C c_2(x)dx = N\hat{n} - \int_I q_2(i)di \tag{10}$$

⁹See Section 5 for a brief discussion of some related economic literature.

¹⁰For instance, standard versions usually involve the possibility of consuming more than one unit of the good or labour.

(i.e. the total amount of numéraire consumed is equal to the total endowment minus the production costs) and, for all $s \in [s_H, \bar{s}_H]$

$$X_s(r_2(s)) = r_1(s)$$

(ie. the total labour supply at every skill level equals the total labour demand). Any competitive equilibrium generates a feasible allocation in this sense.

We adopt the standard notion of Pareto optimality for a continuum of consumers (e.g. Hammond, 1979). A feasible allocation is *Pareto optimal* if there is no alternative feasible allocation under which no consumer has strictly lower utility and a set of consumers of strictly positive measure have strictly higher utility.

Theorem 2. *Any allocation generated by a competitive equilibrium is Pareto optimal.*

So informing about inequality in a perfectly competitive market leads to a socially efficient outcome in terms of the consumers' preferences, and in particular their preferences concerning the inequalities involved in the production of the goods they consume.

3.5 The impact of inequality information on inequality

The previous results suggest that, for a population with some inequality aversion, informing consumers about the inequality involved in the creation of the products purchased can both have a moderating impact on inequality, and a positive effect on social efficiency. A market where there is no inequality information can be modelled as one where consumers' purchasing choices are made entirely on the basis of prices; in other words, consumers act as if they were inequality neutral ($\eta = 0$). If, as the literature surveyed in Section 2 suggests, there is a proportion of the population that are not inequality neutral—i.e. for which $\eta > 0$ —then such consumers would in general be willing to purchase at a higher price to reduce inequality. Theorem 1 tells us that providing inequality information in such cases will have a moderating effect on inequality: on making the information universally available, income inequality in the labour market will fall.

Moreover, given the inequality averse consumers' preferences, the introduction of firms producing with lower inequality and supplying these consumers would constitute a Pareto improvement with respect to the no information $\eta = 0$ case. So a market lacking inequality information provision is not Pareto optimal. By contrast, a competitive market with full inequality information provision is (by Theorem 2). This means that, although there may be productivity losses on information provision—because not all production is assigned to the highest skilled workers—it nevertheless is a Pareto improvement because the impacted consumers are willing to pay the increased costs brought about by lower productivity in order to obtain lower inequality.

These results thus show that information provision does have an impact on inequality. However, they give little information on the size of this impact, an issue to which we now briefly turn.

3.6 Impact estimation via Gabaix-Landier calibration

The impact of information provision will depend on several factors, such as the distribution of inequality aversion in the population, the distributions of productivity and the supply of various skill levels in the labour market. Drawing on the functional forms (and calibrations) from Gabaix and Landier (2008), we now perform some rough back-of-the-envelope calculations to get an idea of an approximate expected size of the impact.

We consider only economies where all consumers have the same inequality aversion η , and compare how wage inequality changes as the inequality aversion increases from the $\eta = 0$ case modelling no inequality information. In order to propose functional forms for productivity (F) and labour supply (X_s), we draw on Gabaix and Landier’s 2008 paper on CEO salaries, which is relevant since the focus of this application is on high-paid workers. Their model is a superstar assignment model in which managers, of differing talent levels, are assigned to firms, of different sizes. (Consumers are absent from their model.) The skill levels $[s_H, \bar{s}_H]$ in our model correspond neatly to the set of managers (ranked by talent) in theirs; we normalise to one unit of skill levels, $[s_H, \bar{s}_H] = [0, 1]$. In their model, firms are of different sizes, with the largest being matched (in equilibrium) with the most talented managers. Since one proxy for firm size—which is fairly reasonable (Gabaix and Landier, 2008, Sect III.A), and the only one they consider that can be directly mapped into our model—is sales, our model can be mapped into their terms by considering firms hiring more skilled workers (and hence producing and selling more) as larger. So their specification of firm earnings in equilibrium as a function of the manager employed provides a reasonable proxy for the production function F (relating skill level to production) in our model. Plugging in their equations and calibrations, we obtain $F(s) \simeq A \frac{\bar{s}_H - (\bar{s}_H - s)^{2/3}}{\bar{s}_H - s}$ for some constant A .¹¹ Finally, we need a form for the labour supply function. Gabaix and Landier (2008, Sect. IV.A) estimate the density of the skill distribution at the upper end of the spectrum as proportional to $(T(\bar{s}_H) - T(s))^{\frac{1}{2}}$, where T is their talent function. Normalising, this gives an indication of the size of the workforce with a given talent level. To get the labour supply function, we need a representation of the propensity of a manager with a given skill level to work for a specific wage. Taking account of the assumption (Section 3.1) that each manager will not accept a salary below 1 (the salary of low paid workers), we assume that the propensity for a manager (of skill s) to accept a wage w is proportion to the ratio between how high it is above 1 and the manager’s productivity (in a matched firm), i.e. it is proportional to $\frac{\max\{w(s)-1,0\}}{F(s)}$. Moreover, since CEO salaries are roughly three orders of magnitude lower than firm market capitalisation (a proxy for firm size; Gabaix and Landier, 2008, Section IV.B), we take the constant of proportionality to be 1000. This gives $X_s(w(s)) = B \frac{1000 \max\{w(s)-1,0\}}{F(s)} \frac{(T(\bar{s}_H)-T(s))^{\frac{1}{2}}}{\int_{s_H}^{\bar{s}_H} (T(\bar{s}_H)-T(s))^{\frac{1}{2}} ds}$. We

¹¹Firm earnings in their model are given by $S + S^\gamma.C.T$, where $S(n)$ is the size of firm n , C is a constant, $T(m)$ is the talent of manager m and $\gamma \simeq 1$. We assume that C is large enough that they can be approximated by $S.T$. Plugging in the calibrated approximate forms for S and T ($S(n) \propto n^{-\alpha}$, with $\alpha \simeq 1$ and $T(m) \propto m^{2/3}$ at the tails of the talent distribution), renormalised to our skill space $[s_H, \bar{s}_H]$ (where, unlike theirs, higher skilled managers correspond to larger s) and set such they are always positive, and using the fact that in equilibrium, firm n is matched to manager n , yields the functional form in the text.

plug these functionals into the equations from Section 3.2, setting the remaining proportionality coefficients to 1. Finally, taking $N = 100$ —the population of consumers to be 100 times larger than that of managers—we can calculate estimates of the max-min wage ratio for various values of η ; they are plotted in Figure 2. Note that, by the choice of the value of N (and the coefficients), the max-min ratio at $\eta = 0$ is around 1100, which is the order of current highest CEO-to-median wage ratio among S&P500 firms (AFL-CIO, 2019).¹²

Figure 2 suggests that even moderate amounts of inequality aversion can have significant effects on the degree of salary inequality. For instance, if everyone had an η of 0.005—which corresponds to them being willing to pay €5 to fully eliminate an inequality of factor 1000 in income—then under full information, the inequality in equilibrium is thirty times lower. Note that this is on the low end of inequality aversion levels found in experiments (Mohan et al., 2015; Hill and Lloyd, 2020). For comparison, a CEO-to-median wage ratio of 50 is below the 10th percentile among S&P500 firms, according to the data in AFL-CIO (2019).

4 Implementation

4.1 When informing consumers about inequality could work

The previous analysis suggests that blanket inequality reporting in the (goods) market could both reduce inequality and improve social efficiency. It is worth stressing some central conditions for the success of such an intervention. A first, which has already been discussed and is supported by existing studies (Section 2), is a degree of inequality aversion among at least a section of the population. A second, more basic, condition is that all consumers understand the inequality report provided: if the report ‘doesn’t speak to’ a consumer, she will not be able to incorporate it into her purchasing decisions. Thirdly, of course, consumers need to trust the source of inequality information: with no possibility of ‘on the spot’ verification, they need to adopt the reports as their beliefs about the goods on offer. That said, much of the information required to calculate the inequality involved in the creation of a product is available and in principle certifiable, either publicly (e.g. in statements of public companies) or to tax authorities (e.g. in tax returns).

These conditions may have consequences for the options for implementing the current proposal.

4.2 Two phases of implementation

Information works—if and when it does—due to market mechanisms, so no information about individual preferences or assumptions about social ones are required. All that is needed is information about current inequalities—much

¹²The form displayed in this graph—and in particular the significant drop in inequality as η increases from 0—continues to hold under different assumptions about the proportionality coefficients, and parameters such as N , as well as C in the previous footnote, although the max-min wage ratio values, and in particular the ‘match’ to S&P500 data, may depend on parameter choice. In that sense, the qualitative conclusion of this exercise—that information provision promises to have a significant effect on income inequality—is independent of the parameter assumptions made here.




Figure 2: Max-min wage inequality vs. η , for markets where all consumers have inequality aversion η .

Calculated using (6) and (7), with $[s_H, \bar{s}_H] = [0, 1]$, $F(s) = \frac{\bar{s}_H - (\bar{s}_H - s)^{2/3}}{\bar{s}_H - s}$, $X_s(w(s)) = \frac{1000 \max\{w(s) - 1, 0\}}{F(s)} \frac{(T(\bar{s}_H) - T(s))^{\frac{1}{2}}}{\int_{\bar{s}_H}^{\bar{s}_H} (T(\bar{s}_H) - T(s))^{\frac{1}{2}} ds}$ and $N = 100$.

of which is contained in tax returns, company statements and so on. Given this, we separate the challenge of implementation into two parts: information collation and information provision.

Whilst a large part of the information required to calculate the inequality involved in the creation of a good or service is available to governments, at least in many developed countries, there is no centralised place where the inequality associated to every good or service on the market is collated and made publicly available in an objective, transparent way. Effectiveness of the intervention—and in particular, the previously noted need for trust in the information provided—depends on oversight of information collation and verification being assigned to a publicly trusted body. This could be a government body, but need not be.¹³ It could, for instance, be a specifically created organisation, drawing on the competences of academic, governmental, non-governmental and business actors. The IPCC (Intergovernmental Panel on Climate Change) is a cross-national source of objective, trusted information about climate change, and similar non-governmental structures have emerged to document the nutri-

¹³Indeed, whilst government collaboration would be useful, insofar as they often can verify relevant earnings information, a single government will often not have all of the required information, especially for international companies and goods with cross-border production lines.

tional properties of foodstuffs.¹⁴ A similar body could be created for inequality.

The second phase of implementation involves the provision or ‘administering’ of inequality information. For the sake of effectiveness—and in particular given the importance of consumers being able to understand and use the information—the information needs to be delivered in an appropriate, easy-to-use, easy-to-understand, accessible form at the point of purchase. Moreover, as emphasised above, it should be provided for all products, which means that voluntary labelling (like CSR or eco-labels) cannot properly implement the proposal. Again, one possibility would involve government involvement: regulation could force all firms to inform customers of the inequality involved in a good or service, much as nutritional information for foodstuffs and origin labels for a range of products are mandatory in many markets. That said, current technology affords the possibility of a regulatorily lighter option. One could develop a mobile phone application which presents, on scanning the barcode of a product, its inequality report. Such an app could be run by any organisation; what counts is that the inequality figures reported are drawn from the database established and certified as described above.¹⁵ Note that similar point-of-purchase applications giving nutritional information about foodstuffs exist, and have been claimed to be effective.¹⁶

4.3 Inequality measures

These brief remarks still leave several questions open: most pressingly, perhaps, which measure of inequality should be reported? There is a rich literature proposing and studying inequality measures (e.g. Lambert, 2001; Chakravarty, 2009; Cowell, 2011). However, many of these measures have not been developed with a mind to capturing the inequality involved in the creation of a product (as opposed, for instance, to the inequality in revenue across or between countries or professions) nor with an aim to public information (by contrast, for example, with informing public policy or furthering economic knowledge). It is thus worth setting out the principles that should guide the development of an appropriate measure of inequality for presentation to consumers at the point of purchase. Three seem to follow from the previous success conditions (Section 4.1):

Objectivity Given the importance of public trust in the inequality reports, the measures used should be based as far as possible on verifiable, objective information. An interesting lesson is provided by CSR measures. These have frequently been based on a range of factors, weighed in perhaps non-codified ways, often involving subjective judgements or non-verified self-reports; their credibility has sometimes been undermined as a result. This consideration pleads for inequality measures which rely as little as possible on difficult-to-access information, unverifiable company self-reports or subjective judgements.

Exhaustivity To impact income inequality ‘across the board’, the inequality reported to consumers should ideally encompass everyone involved in

¹⁴An example is Open Food Facts, <https://world.openfoodfacts.org/>.

¹⁵Indeed, one could even imagine several apps presenting the same information, just as there are several weather apps drawing their information from the same source.

¹⁶A notable example is <https://yuka.io/en/>.

guaranteeing the existence of the good or service on the market. This includes notably the stages of financing, conception, management, production, transport, marketing and sale. To achieve the widest reach possible—as well as to limit the possibility of manipulation—it should include sub-contractors and suppliers as well as financiers and support staff, and take account of all remuneration.¹⁷

Conceptual simplicity Presumably, inequality aversion is related to people’s basic intuitions and opinions on social justice. People will only incorporate inequality information into their decision making if they can connect it to such intuitions—if they can understand it. Inequality thus needs to be presented in a way that can be grasped easily and quickly, without specialist knowledge or theoretical baggage. It must also be clear and unambiguous in meaning, to avoid undermining trust. These considerations plead in favour of conceptual simplicity in the inequality measure—perhaps at the price of precision, analytical power or classical properties emphasised in the existing literature (e.g. satisfaction of various transfer principles, or decomposability; Chakravarty, 2009; Cowell, 2011).¹⁸

Despite the strength of these constraints, there are inequality measures that satisfy them fairly well. One simple example is the *max-min ratio*: the ratio between the highest and lowest hourly revenue among all those involved in the financing, conception, management, production, transport, marketing and sales of the good. It is objectively calculable on the basis of tax returns, salary slips and the like. It is not information-demanding, only requiring the maximum and minimum salary for every firm involved in production. It is, by definition, exhaustive. Finally, it is conceptually easy to understand and grasp: we all know what it means for top management or financiers to earn 1000 times more than factory workers. That said, it is not perfect: for instance, it only looks at the extremes of the distribution, ignoring what happens in between.¹⁹ As intimated above, a trade-off will ultimately need to be made, and approximations such as these may be the price to pay for enough conceptual simplicity to allow consumers to effectively connect with the inequality information. Further experimental research may be able to determine to what extent this is indeed the case.

¹⁷This criterion ties into a current debate about the extent to which rise in income inequality is due to an increase in inequality within firms or between firms (e.g. Mueller et al., 2017; Song et al., 2019). Exhaustivity recommends looking across all firms involved in creating a good (a dimension ignored in the simple model above): so the inequality reported will be that across all such firms, combining the inequality within constituent firms and that between them.

¹⁸For instance, for all their analytical power and theoretical defence, measures such as the Gini index, or attempts to put a monetary value on inequality (e.g. quantifying the ‘objective cost’ of inequality), may not be properly understood or trusted by the typical consumer.

¹⁹Note however that the min-max ratio implies upper and lower bounds for other, more standard inequality measures (e.g. Gini, Theil), so a decrease the min-max ratio implies that the worst case inequality under any other measure has also decreased.

5 Discussion

5.1 Inequality information and some related economic literature

There is a significant literature in economics on the scale and sources of inequality, and the importance and role of information. We now make some brief remarks on connections with, and differences from, parts of this literature.

For consumers sensitive to the inequality involved in the creation of a good, the level of inequality is a ‘credence quality’ in the sense of Darby and Karni (1973)—it is expensive to judge, even after purchase. Such qualities naturally give rise to information asymmetries (firms typically know more about the inequality involved in production than consumers). Accordingly, much focus has been on the market effects of such asymmetries, and signalling techniques that firms could use to differentiate themselves; for instance, voluntary labelling policies (e.g. eco-labels, CSR) have been analysed in such a perspective (e.g. Baksi and Bose, 2007; Crifo and Forget, 2015). By contrast, universal inequality reporting implies complete removal of the information asymmetry and the associated market effects.

Under universal information, the inequality involved in production is a factor of differentiation in the goods market (e.g. Tirole, 1988). As noted in Section 3.2, consumers self-select according to the inequality levels of the goods, and hence the skill levels of the workers employed to produce them. Note also that in equilibrium goods are sold at cost and consumers with different degrees of inequality aversion prefer different inequality levels, so this is not a case of vertical differentiation in the sense of Shaked and Sutton (1983). Under free entry, as we have shown, markets remain competitive.

Policy proposals concerning income inequality—including evaluation of the need for policy—are sometimes suggested by mechanisms that are purportedly responsible for it. The literature on potential mechanisms driving the rise in income inequality is too large to survey here. For illustration, one part focuses on the rise of CEO salaries in the past decades, with suggestions that it could be due to incentivisation considerations in the face of moral hazard, managerial entrenchment, or the structure of the firm-CEO matching at the upper tail of the talent distribution (e.g. Jensen and Murphy, 1990; Hall and Murphy, 2003; Rosen, 1981; Gabaix and Landier, 2008; Tervio, 2008). Rather than tapping into a theory about why there is upward pressure on income inequality, the current proposal focuses on a potential reason why the counterweight downward pressure is so weak: namely that inequality is not incorporated into the market and the consumption decisions of those who care about it. An analogy with pollution may be enlightening here. To the question ‘Why has air pollution increased so much over the last two centuries?’ one can cite upward pressures, such as technological change or population growth, as well as the lack of potential downward pressures, such as the fact that pollution is an externality in many markets. Whilst much of the aforementioned literature on inequality examines the (analogue of the) former sorts of reasons, the current proposal is inspired by the reasons of the latter sort.

Although perhaps neglected recently, the conception of inequality as an externality has a long history in economics (Thurow, 1971). However, to the best of our knowledge, the specific information intervention proposed here, which

follows easily from such a perspective, has not been extensively studied to date. Indeed, the current proposal—to ‘internalise’ the inequality externality by universal information provision—differs from more classic interventions targetting externalities, such as Pigouvian taxes. It is most closely related with property-rights or tradeable-permit approaches. Indeed, one could reframe the proposal in terms of a particular allocation of special ‘inequality-in-production’ property rights. To produce a good with a given inequality level, a firm must acquire a permit to employ that inequality level in the production of that good. Since they specify the inequality level allowed, lets call these *specified* permits. Such permits are *non-amalgamable*: two ‘medium-inequality’ permits for a good only allow the firm to produce two units at that inequality level; they do not permit it to produce one unit of the good at a higher inequality level. Whilst only specified permits can be traded, each consumer is allocated, for each unit of good purchased, a *specifiable* inequality-in-production permit for that good: a ‘blank deed’ that she must ‘fill in’ with the inequality level to which it gives rights, before selling it on the market. So the nature of the permit—the inequality level to which it gives a right—is determined by the consumer prior to sale. It is clear that this market for goods and inequality-in-production permits is basically equivalent to the market set out and studied above—the inequality level at which a consumer purchases the good maps into the inequality level she puts on the permit she sells; the price at which the consumer purchases the good at a certain inequality level is the result of paying the market price for the good and receiving the proceeds of the sale of her specified inequality permit. So the equilibrium is the same, and the results carry over.

This reframing brings out several points on which the proposal differs from typical property-rights or tradeable-permit approaches to externalities. First, there is a simple allocation mechanism: according to good purchase. In particular, unlike standard marketable-permit (or ‘cap and trade’) approaches in, say, environmental policy, there is no need for a social planner (or regulatory authority) to decide on the optimal aggregate amount of inequality. Second, the permits here are non-amalgamable, in contrast to carbon markets for instance, where a firm can buy lots of permits from different actors to pollute more in the production of the same quantity of the good. Third, the allocation of specifiable permits that must be specified before sale has not, to our knowledge, been proposed or explored previously. These differences all contribute to clarifying each consumer’s *responsibility* for inequality: she alone specifies the inequality level on the permit she sells, and is ensured that it will result in the production of at most one unit of the good at that level. As such, they tie into two previously mentioned contrasts with typical approaches, that are worth recalling.

One is the reliance on consumer preferences *over the inequalities in the production of the goods they purchase* (or, under the property-rights reframing, the inequality on their specified permit), rather than their preferences concerning the *overall level of inequality in society*. The level of inequality (or equality) in society could be considered as a public good over which consumers have preferences, and hence as a ‘nondepletable externality’, and standard analysis of property-rights or tradeable-permit allocation could thus be applied, with all the familiar related issues (e.g. Baumol and Oates, 1988).²⁰ Consumer preferences

²⁰In particular, a pure property-rights approach (with no aggregate inequality quota) will fail to be optimal (e.g. Mas-Colell et al., 1995); an approach involving tradeable permits and

over the overall inequality in society play no role in our analysis, and inequalities associated to purchased goods are closer to private goods, a fact which is central to the social efficiency of the information provision intervention. Of course, as emphasised, the intervention only has an impact if consumers' preferences are sensitive to these latter inequalities: existing empirical evidence, documented in Section 2, suggests that a significant proportion of them are.

The other is the reliance on consumer inequality preferences, *and nothing else*: in particular, there is no role for the social planner (beyond ensuring the proper functioning of the market). By contrast, a tradeable-permits approach to inequality would require the social planner to appropriately determine the optimal aggregate amount of inequality (in the model above, the aggregate quantity of goods produced at each inequality level). Such a quota reflects the sorts of inequality versus productivity tradeoffs mentioned in the Introduction, and naturally poses the question of how the social planner is equipped, or sufficiently well-informed, to correctly set these values. As noted previously, the current proposal avoids such issues.

5.2 Information beyond income inequality

The contrast with existing approaches to externalities poses the question of whether the approach set out here could be applied to externalities other than income inequality. Certainly much of the theoretical analysis (Section 3) seems extendable to other externalities, and would seem to hold if income inequality were replaced by capital-to-labour share of proceeds (hence connecting into the wealth vs. income inequality debate), the lowest wages paid by the firm (or some other indicator of the degree of offshoring, dumping or unfair wages to low-paid workers), or income inequalities across gender or race in the firm, to name but a few examples. Similar theoretical points—impact if there is consumer sensitivity to these issues, and social efficiency gains—would thus hold for information provision on these issues. That said, for any issue, the potential for the proposal to be effectively applied to a given externality will ultimately depend on the extent to which the conditions noted in Section 4.1 hold. Recall that one of these conditions was the sensitivity in consumer preferences to the levels of the externality associated with goods in the market. Another is sufficient understanding of the information to be able to incorporate it into purchasing decisions. Whilst there may be reason to suspect that these conditions hold for the previous examples, other cases are less straightforward.

One interesting possibility where they could matter would involve universal information provision on the global-warming-related impacts stemming from the production of a good. Such information possibilities already exist, for instance in carbon footprint reporting,²¹ though they are voluntary in many sectors and regions. Information interventions in this domain will face a significant challenge concerning consumer understanding. Whilst in the case of income inequality, the intervention can tap into existing intuitions and opinions about social justice, it is less clear whether people have sufficiently developed views about atmospheric processes to incorporate, for instance, CO₂ emissions data

an aggregate quota may fail to be optimal in the presence of uncertainty (e.g. Weitzman, 1974).

²¹A carbon footprint is usually defined as the total emissions caused by an individual, event, organization, or product, expressed as carbon dioxide equivalent.

into purchasing choices. Of course, recognising a challenge does not mean considering it insurmountable. It may well be possible to develop ‘global warming impact measures’ that are easily understandable, whilst also satisfying the other conditions set out above, such as objectivity. Moreover, understandability is a relative concept, depending on common knowledge in the community; one might thus expect that improvements in climate awareness and education may enhance the effectiveness of previously incomprehensible information.

6 Conclusion

This paper proposes universal information provision as a means of moderating income inequality, and improving social efficiency concerning it. Current empirical evidence suggests a significant mismatch between current levels of income inequality and peoples’ perceived and ideal levels; moreover, it suggests that many would be willing to pay more for a reduction in the income inequality involved in the creation of the goods they purchase. Tapping into these facts, we show on a simple model that informing all consumers about the inequality involved for each good on the market will lead to a drop in income inequality, even for only moderate levels of aversion to inequality among a limited section of the population. The more the population are willing to pay to avoid large inequality levels, the less the income inequality effectively achieved in equilibrium. Moreover, information provision re-establishes social efficiency, incorporating in particular the inequality dimensions of consumer preferences.

We also argue that implementation of the proposal is far from daunting, setting out some principles for the choice of inequality measure to report, and suggestions for information collation and provision.

This paper only focuses on the economic dimension of the proposed intervention—and in particular its impact on inequality levels in equilibrium—but this is not the only one. In particular, as with any information intervention, there is a potential political dimension. Inequality information can correct misperceptions, which, as noted above (Section 2), are widespread. It can improve awareness of the issue. Moreover, to the extent that it relates inequality levels to consumer choice, it involves an empowerment of citizens on this issue.

A Extensions and proofs

A.1 General inequality-attitude preferences

In the interests of generality, we prove the results for a model that is more general than that presented in the bulk of the paper. The firm and worker structures are the same; the only difference is the use of more general forms of utility function for consumers. More specifically, each consumer has a utility function of the following additively separable form:

$$U_j(i, n) = \begin{cases} v_j(n) + (v_j - u_j(\mathcal{I}(i))) & i \in I \\ v_j(n) & i = \{\} \end{cases} \quad (11)$$

where v_j , as in the model in Section 3.1, is the ‘intrinsic’ value of the good, v_j —the utility function over the numéraire—is strictly increasing and twice differentiable, u_j —the disutility of inequality—is a (weakly) increasing, twice differentiable function, and $\mathcal{I}(i)$ is the ‘justifiable-threshold-corrected’ inequality, given by:

$$\mathcal{I}(i) = \begin{cases} 0 & i \leq \theta \\ i - \theta & i > \theta \end{cases} \quad (12)$$

We assume without loss of generality that $u_j(0) = 0$ (the disutility of no inequality is zero). Note that the same θ is involved for all consumers, though consumers with higher justifiable-inequality thresholds can be modelled by u_j functions which take the value zero up to a certain (higher) level (this is accommodated since these utility functions are not assumed to be strictly increasing).

Clearly, the utility function presented in Section 3.1 (equation (1)) is the special case where v_j is the identity and $u_j(x) = \eta_j x$. However, the functional form (11) is considerably more flexible, accommodating a range of ‘shapes’ of the disutility of inequality, including an higher sensitivity to inequality increases at higher inequality levels, as well as the (opposite) more acute sensitivity to changes at low inequality levels. Note also that using log utilities, this form encompasses inequality disutility that is relative to wealth or the price of the good (i.e. multiplicative), rather than absolute (additive).²²

Following the case treated in the paper, we assume that all consumers have the same v (intrinsic value of the good); moreover, we assume that they share the same (not necessarily linear) utility for the numéraire, $v = v_j$. We say that consumer j is *more inequality averse* than k if, for every pair of inequality levels $i_1 > i_2$, the difference in j ’s disutility between them is larger than for k : i.e. $u_j(i_1) - u_j(i_2) \geq u_k(i_1) - u_k(i_2)$. When this holds with strict inequality for some i_1, i_2 we write $u_j >_{I.A.} u_k$.²³ More inequality averse consumers obtain a sharper

²²For instance, taking v_j and u_j to be the appropriate multiple of logarithms, it is clear that the utility function

$$U_j(i, n) = \begin{cases} An^\alpha \mathcal{I}(i)^\beta & i \in I \\ n^\alpha & i = \{\} \end{cases}$$

characterises preferences belonging to the family represented by (11) and to which the results proved below apply.

²³Note that, were u_j and u_k strictly increasing, this condition would be equivalent to u_j being obtainable from u_k by a (strictly) convex transformation.

jump in disutility from any increase in inequality. Clearly, for the special case considered in the paper, higher η_j implies more inequality aversion in this sense.

To formulate the generalised version of Theorem 1, we assume that all consumers are ordered according to inequality aversion: i.e. there exists $u_1 >_{I.A.} \dots >_{I.A.} u_K$ where u_K is the constant function taking the value zero, such that each consumer's utility over inequality is given by one of these functions.²⁴ A consumer of type j has (dis)utility for inequality u_j . Inequality aversion distributions and other related notions are defined as in Sections 3.1 and 3.3 with this notion of consumer type. Given that, the formulations of Theorems 1 and 2 in the context of this general model is the identical to those in Sections 3.3 and 3.4.

We first present an analysis of the wage schedule and then prove Theorems 1 and 2 under this general model. The versions for special case presented in the paper follow immediately.

A.2 Derivation of equilibrium wage schedule

General case

We first derive the equilibrium wage schedule under the general utility form (11). The average cost of production of a good of inequality $w(s)$ is $\frac{w(s)+1}{F(s)}$. By standard reasoning, in equilibrium if there is any demand for goods at inequality level $w^*(s)$, then $p^*(w^*(s)) = \frac{w^*(s)+1}{F(s)}$.²⁵ So, in equilibrium, wages and prices are connected via the production function.

For a consumer of type j faced with prices p and wages w , the FOC conditions for an interior solution above the threshold θ are given by:

$$-\bar{v}'(p(w(s))) \frac{dp}{d(w(s))} = -\bar{v}'(p(w(s))) \left(\frac{1}{F(s)} - p(w(s)) \cdot \frac{1}{F(s)} \cdot \frac{F'(s)}{w'(s)} \right) = -u'_j(w(s) - \theta) \quad (13)$$

where $\bar{v}(x) = v(\hat{n} - x)$. Plugging in the form of equilibrium p^* , this can be rewritten as:

$$(w^*)'(s) = \frac{F'(s)(w^*(s) + 1)}{F(s) \left(1 - \frac{u'_j(w^*(s) - \theta)}{\bar{v}'(\frac{w^*(s)+1}{F(s)})} F(s) \right)} \quad (14)$$

Because F, v, u_j are twice differentiable, $\bar{v}'(x) < 0$, $u'_j(x) \geq 0$, and F is bounded away from zero, the functional on the right hand side (considered as a functional of s and $w^*(s)$) is uniformly Lipschitz continuous in s and $w^*(s)$. Hence, by the Picard-Lindelöf Theorem, for any initial value for w^* , there exists a unique, continuously differentiable solution $w^*(s)$ for the initial value problem given by (14) and the initial value. We write such solutions as functions $G(C_j, u_j)$ where

²⁴In the special case considered in the paper, this assumption is automatically satisfied, since the η_j are ordered.

²⁵Suppose that $p^*(w^*(s)) > \frac{w^*(s)+1}{F(s)}$ in equilibrium: then a firm entering the market and recruiting at skill level s would make a strictly positive profit, violating the free entry condition. On the other hand, if $p^*(w^*(s)) < \frac{w^*(s)+1}{F(s)}$, firms recruiting at skill level s would make strictly negative profits, and hence drop out of the market (i.e. this would be a violation of the free entry condition).

C_j is a constant (real number) encoding the initial value. So, in equilibrium, the wage schedule has the form

$$w^*(s) = G(C_j, u_j)(s) \quad (15)$$

for all s servicing consumers of type j . Note that, since $F' > 0$ and v and u_j are increasing (strictly in the former case), it follows from (14) that $(w^*)'(s) > 0$ for all s : $G(C_j, u_j)(s)$ is strictly increasing in s .

Whenever a consumer with disutility for inequality u_j purchases the good with inequality below the threshold θ , she is minimising price under the condition that the inequality is below θ , so across s servicing such customers, we have $v \left(\hat{n} - \frac{w^*(s)+1}{F(s)} \right) + v$ constant. Therefore, in equilibrium:

$$w^*(s) = CF(s) - 1 \quad (16)$$

for some constant C .

Furthermore, higher skilled workers service less inequality averse consumers in equilibrium. To see that, consider $j < k$ with consumers with inequality utility functions $u_j >_{I.A.} u_k$ purchasing goods at inequality above the threshold θ and suppose for reductio that s services j but not k and t services k but not j , with $s > t$. Since j prefers the good produced by firms employing s to that produced by firms employing skill level t , we have:

$$\bar{v}(p^*(w^*(s)) - \bar{v}(p^*(w^*(t))) > u_j(\mathcal{I}(w^*(s)) - u_j(\mathcal{I}(w^*(t))))$$

whereas since k prefers the good produced by firms employing t to that produced by firms employing skill level s :

$$\bar{v}(p^*(w^*(s)) - \bar{v}(p^*(w^*(t))) < u_k(\mathcal{I}(w^*(s)) - u_k(\mathcal{I}(w^*(t))))$$

If $w^*(s) > w^*(t)$, then it follows from the two inequalities that $u_k(\mathcal{I}(w^*(s)) - u_k(\mathcal{I}(w^*(t))) > u_j(\mathcal{I}(w^*(s)) - u_j(\mathcal{I}(w^*(t))))$, contradicting the fact that $u_j >_{I.A.} u_k$. If $w^*(s) \leq w^*(t)$, then $p^*(w^*(s)) = \frac{w^*(s)+1}{F(s)} < \frac{w^*(t)+1}{F(t)} = p^*(w^*(t))$, since F is strictly increasing. But then firms employing s produce goods which are cheaper and have less inequality than those employing skill level t , and hence are preferred by all consumers; this contradicts consumer k 's preferences. So, for all consumers of types $j < k$ purchasing goods at inequality above the threshold θ , if they are serviced by s and t respectively in equilibrium, then $s \leq t$: the higher skilled workers service the less inequality averse consumers.

Given this, the equilibrium wage schedule is characterised by a sequence of positive real numbers C_k, \dots, C_K , for $0 \leq k \leq K$, and a sequence $\underline{s}_H \leq \underline{s}_k \leq \bar{s}_k \leq \theta \leq \underline{s}_{k+1} \leq \bar{s}_{k+1} \leq \dots \leq \underline{s}_K \leq \bar{s}_K \leq \bar{s}_H$ with:

$$w^*(s) = \begin{cases} G(C_K, u_K)(s) & s \in [\underline{s}_K, \bar{s}_K] \\ G(C_{K-1}, u_{K-1})(s) & s \in [\underline{s}_{K-1}, \bar{s}_{K-1}] \\ \dots & \dots \\ G(C_{k+1}, u_{k+1})(s) & s \in [\underline{s}_{k+1}, \bar{s}_{k+1}] \\ C_k F(s) - 1 & s \in [\underline{s}_k, \bar{s}_k] \end{cases} \quad (17)$$

Moreover, by the closure of the good market, these sequences satisfy:

$$\mu_j = \int_{\underline{s}_j}^{\overline{s}_j} F(s) X_s (G(C_j, u_j)(s)) ds \quad (18)$$

for every $k < j \leq K$, and

$$\sum_{j=1}^k \mu_j = \int_{\underline{s}_k}^{\overline{s}_k} F(s) X_s (C_k F(s) - 1) ds \quad (19)$$

Note that, for any j such that $\mu_j = 0$, $\underline{s}_j = \overline{s}_j$.

Moreover, for any $k < j < j' < K$, if $\overline{s}_j = \underline{s}_{j'}$, then by continuity of w^* (the law of one price), C_j and $C_{j'}$ are related by

$$w^*(\overline{s}_j) = G(C_j, u_j)(\overline{s}_j) = G(C_{j'}, u_{j'})(\overline{s}_j) \quad (20)$$

(i.e. $G(C_j, u_j)$ and $G(C_{j'}, u_{j'})$ solve their respective initial value problems with the same initial value). By contrast, for $k < j < K$ with $\mu_j > 0$, if $\overline{s}_j < \underline{s}_{j+1}$, then $\overline{s}_j = \sup \{s : X_s (G(C_j, u_j)(s)) > 0\}$. This is because, for any s such that $X_s (G(C_j, u_j)(s)) > 0$, if $\overline{s}_j < s$ a firm would be able to enter the market, hire workers with skill s at wage $G(C_j, u_j)(s)$, and sell to consumers with inequality utility function u_j . So, in equilibrium, \overline{s}_j must be greater than or equal to the supremum such s ; but since above the supremum there is no labour supply, $\overline{s}_j = \sup \{s : X_s (G(C_j, u_j)(s)) > 0\}$. In other words, \overline{s}_j is determined by C_j , u_j and the functional form G as the maximal skill level for which there is positive labour supply under this wage pattern. A similar argument establishes that $\underline{s}_{j+1} = \inf \{s : X_s (G(C_j, u_j)(s)) > 0\}$ whenever $\mu_{j+1} > 0$. Similar arguments establish that $\overline{s}_K = \sup \{s : X_s (G(C_K, u_K)(s)) > 0\}$ and $\underline{s}_k = \inf \{s : X_s (C_k F(s) - 1) > 0\}$ whenever there is positive demand for the good at the corresponding inequality levels. So, in equilibrium, the wage schedule is entirely characterised, modulo the functional form G , by the inequality aversion distribution μ and the sequence C_k, \dots, C_K .

Note finally that by (20) and the discussion below it, as well as the fact that each $G(C_j, u_j)(s)$ is strictly increasing in s , w^* is strictly increasing in the skill level s .

Special case: utility of form (1)

In the special case presented in the bulk of the paper, the FOC (13) simplifies to

$$\frac{dp}{d(w(s))} = \frac{1}{F(s)} - p(w(s)) \cdot \frac{1}{F(s)} \cdot \frac{F'(s)}{w'(s)} = -\eta_j \quad (21)$$

For equilibrium p^* and w^* , this can be solved analytically as:

$$w^*(s) = \frac{C_j F(s) - 1}{\eta_j F(s) + 1} \quad (22)$$

for all s servicing such consumers. The utility obtained by the consumer is $\hat{n} + v + \theta \eta_j - C_j$.²⁶ As noted (and as can be verified directly from $(w^*)'(s) = \frac{(C_j + \eta_j) F'(s)}{(\eta_j F(s) + 1)^2}$), the wage is strictly increasing in s .

²⁶Plugging in the form of p^* , (21) implies that $\frac{(w^*)'(s)}{w^*(s)+1} = \frac{F'(s)}{F(s)(1+\eta_j F(s))}$; solving this differential equation yields (22). Explicitly, it implies that:

Plugging this into the general solution form derived above, we obtain that the equilibrium wage schedule is characterised by a sequence of positive real numbers C_k, \dots, C_K , for $0 \leq k \leq K$, and a sequence $\underline{s}_H \leq \underline{s}_k \leq \overline{s}_k \leq \theta \leq \underline{s}_{k+1} \leq \overline{s}_{k+1} \leq \dots \leq \underline{s}_K \leq \overline{s}_K \leq \overline{s}_H$ with:

$$w^*(s) = \begin{cases} \frac{C_K F(s) - 1}{\eta_K F(s) + 1} & s \in [\underline{s}_K, \overline{s}_K] \\ \frac{C_{K-1} F(s) - 1}{\eta_{K-1} F(s) + 1} & s \in [\underline{s}_{K-1}, \overline{s}_{K-1}] \\ \dots & \dots \\ \frac{C_{k+1} F(s) - 1}{\eta_{k+1} F(s) + 1} & s \in [\underline{s}_{k+1}, \overline{s}_{k+1}] \\ C_k F(s) - 1 & s \in [\underline{s}_k, \overline{s}_k] \end{cases} \quad (23)$$

Moreover, by the closure of the good market, these sequences satisfy:

$$\mu_j = \int_{\underline{s}_j}^{\overline{s}_j} F(s) X_s \left(\frac{C_j F(s) - 1}{\eta_j F(s) + 1} \right) ds \quad (24)$$

for every $k < j \leq K$ and

$$\sum_{j=1}^k \mu_j = \int_{\underline{s}_k}^{\overline{s}_k} F(s) X_s (C_k F(s) - 1) ds \quad (25)$$

Note that, for any j such that $\mu_j = 0$, $s_j = \overline{s}_j$. Of course, the other properties of the general solution (e.g. characterisation by μ and the sequence C_k, \dots, C_K) are inherited in this special case.

$$\begin{aligned} \int \frac{(w^*)'(s)}{w^*(s) + 1} ds - \int \frac{F'(s)}{F(s)(1 + \eta_j F(s))} ds &= C \\ \int \frac{dw^*}{w^* + 1} - \int \frac{dF}{F(1 + \eta_j F)} &= C \\ \int \frac{dw^*}{w^* + 1} + \int \frac{d(\eta_j + \frac{1}{F})}{(\eta_j + \frac{1}{F})} &= C \\ \ln(w^*(s) + 1) - \ln(w^*(\underline{s}) + 1) + \ln(\eta_j + \frac{1}{F(s)}) - \ln(\eta_j + \frac{1}{F(\underline{s})}) &= C \\ w^*(s) + 1 &= \frac{D}{(\eta_j + \frac{1}{F(s)})} \\ w^*(s) &= \frac{(D - \eta_j) F(s) - 1}{\eta_j F(s) + 1} \end{aligned}$$

where $D = e^C (w^*(\underline{s}) + 1) \cdot (\eta_j + \frac{1}{F(\underline{s})})$ and \underline{s} is the lowest skill level servicing consumers of inequality aversion η_j . Plugging this $w^*(s)$ into the consumer's utility function yields:

$$\begin{aligned} u_j(i) &= \hat{n} + v - p(i) - \psi_j(i) \\ &= \hat{n} + v - \frac{w^*(s) + 1}{F(s)} - \eta_j (w^*(s) - \theta) \\ &= \hat{n} + v + \eta_j \theta - \frac{w^*(s)(1 + \eta_j F(s)) + 1}{F(s)} \\ &= \hat{n} + v + \eta_j \theta - (D - \eta_j) \end{aligned}$$

A.3 Proof of Theorem 1

We now prove Theorem 1 under the general utility form for consumers, (11). The statement is precisely as in Section 3.3.

Let the equilibrium wages under μ' be as in (17), satisfying (18) and (19), for the sequences C_k, \dots, C_K and $\underline{s}_H \leq \underline{s}_k \leq \bar{s}_k \leq \theta \leq \underline{s}_{k+1} \leq \bar{s}_{k+1} \leq \dots \leq \underline{s}_K \leq \bar{s}_K \leq \bar{s}_H$. If $k = K$ and all wages are below θ , then this same wage schedule satisfies the equilibrium conditions under μ , and has the same inequality. We henceforth suppose that not all wages are below θ , so $k < K$. We consider the case where $\underline{s}_j = \bar{s}_{j-1}$ for all $k < j \leq K$, where the $\bar{s}_j, \underline{s}_j$ are as in (6): an argument similar to that below holds for μ_j for the highest j for which $\underline{s}_j \neq \bar{s}_{j-1}$, hence establishing the other case. Moreover, we assume that $\mu'_K > 0$: again, if this is not the case, the same argument can be run starting from the highest j such that $\mu'_j > 0$. Now consider the following construction.

Definition 3. For an inequality aversion distribution μ and a $\bar{s} \in [\underline{s}_H, \bar{s}_H]$, define the sequences D_1, \dots, D_K and $\underline{s}_H \leq \underline{t}_1 \leq \bar{t}_1 = \underline{t}_{1+1} \leq \bar{t}_{1+1} = \dots \leq \underline{t}_K \leq \bar{t}_K \leq \bar{s}_H$ generated by \bar{s} with respect to μ inductively by:

- $\bar{t}_K = \bar{s}$, D_K is the unique C satisfying²⁷

$$\sup \{t : X_t(G(C, u_K)(t)) > 0\} = \bar{s} \quad (26)$$

and \underline{t}_K is the unique t satisfying

$$\mu_K = \int_{\underline{t}_K}^{\bar{t}_K} F(s) X_s(G(D_K, u_K)(s)) ds \quad (27)$$

- given D_{j+1} and \underline{t}_{j+1} with $G(D_{j+1}, u_{j+1})(\underline{t}_{j+1}) > \theta$, let $\bar{t}_j = \underline{t}_{j+1}$, define D_j as the constant in the solution of (14) for u_j with initial value:

$$G(D_j, u_j)(\bar{t}_j) = G(D_{j+1}, u_{j+1})(\bar{t}_j)$$

and define $\underline{t}_j = \max \{t_{j1}, t_{j2}, t_{j3}\}$ where:

- t_{j1} is the maximal t satisfying

$$\mu_j = \int_{\underline{t}_j}^{\bar{t}_j} F(s) X_s(G(D_j, u_j)(s)) ds \quad (28)$$

if such a t exists, and \underline{s}_H if not;

- t_{j2} is the unique t satisfying $G(D_j, u_j)(t) = \theta$ if such a t exists,²⁸ and \underline{s}_H if not;
- $t_{j3} = \inf \{t : X_t(G(D_j, u_j)(t)) > 0\}$.

- if D_{j+1} and \underline{t}_{j+1} are such that $G(D_{j+1}, u_{j+1})(\underline{t}_{j+1}) = \theta$, then $\bar{t}_j = \underline{t}_{j+1}$, D_j is defined by

$$D_j F(\bar{t}_j) - 1 = G(D_{j+1}, u_{j+1})(\bar{t}_j) = \theta$$

and $\underline{t}_j = \inf \{t : X_t(D_j F(t) - 1) > 0\}$. In this case, set $l = j$ and the induction (construction of the sequences) is complete.

²⁷Such a C is unique because of the uniqueness of the solutions of initial value problems given by (14) and the fact that X_s is continuous and strictly increasing for each s .

²⁸Such a t is unique because $G(D_j, u_j)(s)$ is continuous and strictly increasing in s .

Consider the sequences generated by $\overline{s_K}$ with respect to μ , and call them D_l, \dots, D_K and $\underline{s_H} \leq \underline{t_l} \leq \overline{t_l} = \underline{t_{l+1}} \leq \overline{t_{l+1}} = \dots \leq \underline{t_K} \leq \overline{t_K} \leq \overline{s_H}$. By the definition, they generate a wage schedule w' according to (17). Moreover, by construction, this wage schedule is such that $G(D_K, u_K)(\overline{t_K}) = G(C_K, u_K)(\overline{s_K})$ — the maximum wage is the same as under the equilibrium for μ' —and it satisfies the market closure condition in the goods market (under μ) for all inequality aversion levels greater than k .

Proposition 4. $\sum_{j=1}^K \mu_j \leq \sum_{j=l+1}^K \int_{\underline{t_j}}^{\overline{t_j}} F(s) X_s(G(D_j, u_j)(s)) ds + \int_{\underline{t_l}}^{\overline{t_l}} F(s) X_s(D_l F(s) - 1) ds$, with strict inequality whenever $\mu_j \neq \mu'_j$ for some $l < j \leq K$.

Proof. First of all, since $\overline{t_K} = \overline{s_K}$, it follows from the equilibrium solution under μ' (Section A.2) and Definition 3 that $D_K = C_K$. Moreover, $\mu'_K \geq \mu_K$ implies that $\underline{s_K} \leq \underline{t_K}$, where the latter inequality is strict whenever the former is. Since $\underline{s_K} \geq \inf \{s : X_s(G(C_K, u_K)(s)) > 0\}$, it follows that the same holds for $\underline{t_K}$. We first show that $G(D_{K-1}, u_{K-1})(s) > G(C_K, u_K)(s)$ for all $\underline{s_H} \leq s < \underline{t_K}$. By definition, $G(D_{K-1}, u_{K-1})(\underline{t_K}) = G(C_K, u_K)(\underline{t_K})$. Moreover, since $u_{K-1} >_{I.A.} u_K$, it clearly follows that $u'_{K-1}(x) \geq u'_K(x)$, for all x . Hence, for every s such that $G(D_{K-1}, u_{K-1})(s) = G(C_K, u_K)(s)$, since these functions solve (14) with u_{K-1} and u_K respectively and this initial value, it follows from the differential equation and the aforementioned ordering of u'_{K-1} and u_K that $G(D_{K-1}, u_{K-1})'(s) < G(C_K, u_K)'(s)$. So $G(D_{K-1}, u_{K-1})'(\underline{t_K}) < G(C_K, u_K)'(\underline{t_K})$. It follows by a standard argument that there exists no $\underline{s_H} \leq s < \underline{t_K}$ with $G(D_{K-1}, u_{K-1})(s) = G(C_K, u_K)(s)$,²⁹ so $G(D_{K-1}, u_{K-1})(s) > G(C_K, u_K)(s)$ for all such s , as required. Since X_s is strictly increasing, for each s , it follows that, if $\underline{s_K} < \underline{t_K}$, then

$$\int_{\underline{s_K}}^{\underline{t_K}} F(s) X_s(G(D_{K-1}, u_{K-1})(s)) ds > \int_{\underline{s_K}}^{\underline{t_K}} F(s) X_s(G(C_K, u_K)(s)) ds$$

Now note that by the uniqueness of the solutions defining $G(D_{K-1}, u_{K-1})$, if $\underline{s_K} = \underline{t_K}$, then $G(D_{K-1}, u_{K-1}) = G(C_{K-1}, u_{K-1})$. If $\underline{s_K} < \underline{t_K}$, then by the previous observation, $G(D_{K-1}, u_{K-1})(\underline{s_K}) > G(C_K, u_K)(\underline{s_K}) = G(C_{K-1}, u_{K-1})(\underline{s_K})$. However, if there exists $s < \underline{s_K}$ with $G(D_{K-1}, u_{K-1})(s) = G(C_{K-1}, u_{K-1})(s)$, then since these functions solve the same differential equation with the same initial value (at s), by the uniqueness of the solution they must be identical, contradicting the strict inequality at $\underline{s_K}$. Hence there is no such s , and $G(D_{K-1}, u_{K-1})(s) > G(C_{K-1}, u_{K-1})(s)$ for all $s \leq \underline{s_K}$. Hence, for every $s' < \underline{s_K}$

$$\int_{s'}^{\underline{s_K}} F(s) X_s(G(D_{K-1}, u_{K-1})(s)) ds \geq \int_{s'}^{\underline{s_K}} F(s) X_s(G(C_{K-1}, u_{K-1})(s)) ds$$

²⁹For reductio, suppose there exists $\underline{s_H} \leq s < \underline{t_K}$ with $G(D_{K-1}, u_{K-1})(s) = G(C_K, u_K)(s)$, and let s be the largest such one. By the previous fact, $G(D_{K-1}, u_{K-1})'(s) < G(C_K, u_K)'(s)$. By the intermediate value theorem, there exists $\sup \{t : G(D_{K-1}, u_{K-1})(t) < G(C_K, u_K)(t), G(D_{K-1}, u_{K-1})'(t) < G(C_K, u_K)'(t)\} < s' < \inf \{t : G(D_{K-1}, u_{K-1})(t) > G(C_K, u_K)(t), G(D_{K-1}, u_{K-1})'(t) < G(C_K, u_K)'(t)\}$ with $G(D_{K-1}, u_{K-1})(s') = G(C_K, u_K)(s')$. However, by construction $G(D_{K-1}, u_{K-1})'(s') \geq G(C_K, u_K)'(s')$, contradicting the fact that $G(D_{K-1}, u_{K-1})'(s) < G(C_K, u_K)'(s)$ for all s with $G(D_{K-1}, u_{K-1})(s) = G(C_K, u_K)(s)$.

with strict inequality whenever $s_K < t_K$. Since $\mu_{K-1} + \mu_K \leq \mu'_{K-1} + \mu'_K$ it follows that $\underline{t_{K-1}} \geq \underline{s_{K-1}}$, with strict inequality whenever either $\mu_K < \mu'_K$ or $\mu_{K-1} + \mu_K < \mu'_{K-1} + \mu'_K$.

The same argument implies, for every $k < j < K$, that if $G(D_{j+1}, u_{j+1})(\underline{t_{j+1}}) \geq G(C_{j+1}, u_{j+1})(\underline{t_{j+1}})$ and $\underline{t_{j+1}} \geq \underline{s_{j+1}}$ with strict inequality (in both inequalities) whenever there exists $\hat{j} \geq j+1$ with $\sum_{j'=\hat{j}}^K \mu_{j'} < \sum_{j'=\hat{j}}^K \mu'_{j'}$, then $G(D_j, u_j)(\underline{t_j}) \geq G(C_j, u_j)(\underline{t_j})$ and $\underline{t_j} \geq \underline{s_j}$ with strict inequality whenever there exists $\hat{j} \geq j$ with $\sum_{j'=\hat{j}}^K \mu_{j'} < \sum_{j'=\hat{j}}^K \mu'_{j'}$. Hence, by induction, $G(D_{k+1}, u_{k+1})(\underline{t_{k+1}}) \geq G(C_{k+1}, u_{k+1})(\underline{t_{k+1}})$ and $\underline{t_{k+1}} \geq \underline{s_{k+1}}$, with strict inequality whenever there exists $j > k$ with $\mu_j \neq \mu'_j$.

If $\mu_j = \mu'_j$ for all $j > k$, then $\sum_{j'=1}^k \mu_{j'} = \sum_{j'=1}^k \mu'_{j'}$ and the sequences D_{k+1}, \dots, D_K and C_{k+1}, \dots, C_K are identical, as are $\underline{t_{k+1}} \leq \overline{t_{k+1}} = \dots \leq \underline{t_K} \leq \overline{t_K}$ and $\underline{s_{k+1}} \leq \overline{s_{k+1}} = \dots \leq \underline{s_K} \leq \overline{s_K}$. So $G(D_{k+1}, u_{k+1})(\underline{t_{k+1}}) = \theta$, and hence $D_k F(\underline{t_{k+1}}) - 1 = \theta$, so $D_k = C_k$, $\underline{t_k} = \underline{s_k}$ and $\sum_{j'=1}^k \mu_{j'}$ satisfy (8). So the inequality in the Proposition holds with equality.

If there exists $j > k$ with $\mu_j \neq \mu'_j$, then $G(D_{k+1}, u_{k+1})(\underline{t_{k+1}}) > G(C_{k+1}, u_{k+1})(\underline{t_{k+1}}) = \theta$ and $\underline{t_{k+1}} > \underline{s_{k+1}}$. By definition, for all $\underline{D_j}$ with $k \geq j > l$, and $G(D_j, u_j)(s) > \theta \geq C_k F(s) - 1$ for all $\underline{t_j} < s \leq \overline{t_j}$, so $\sum_{j=l+1}^k \mu_j = \sum_{j=l+1}^k \int_{\underline{t_j}}^{\overline{t_j}} F(s) X_s (G(D_j, u_j)(s)) ds > \int_{\underline{t_{l+1}}}^{\overline{t_k}} F(s) X_s (C_k F(s) - 1) ds$. Moreover, since $F' > 0$ and $C_k F(\overline{t_k}) - 1 = \theta$, $C_k F(\overline{t_l}) - 1 < \theta = D_l F(\overline{t_l}) - 1$, so $D_l > C_k$, and hence $D_l F(s) - 1 > C_k F(s) - 1$ for all $s \leq \overline{t_l}$. But then the t' with $\sum_{j=1}^l \mu_j = \int_{t'}^{\overline{t_l}} F(s) X_s (D_l F(s) - 1) ds$ must be such that $t' > \underline{s_k}$. However since, for every $t \leq \overline{t_l}$, if $X_t (C_k F(t) - 1) > 0$, then $X_t (D_l F(t) - 1) > 0$, $\underline{t_l} = \inf \{t : X_t (D_l F(t) - 1) > 0\} \leq \inf \{t : X_t (C_k F(t) - 1) > 0\} = \underline{s_k}$. So $\sum_{j=1}^l \mu_j < \int_{\underline{t_l}}^{\overline{t_l}} F(s) X_s (D_l F(s) - 1) ds$. This, and the definition of the sequences (notably (27) and (28)) establishes the strict \square

Proposition 5. Let $\pi(\overline{s}) = \sum_{j=l+1}^K \int_{\underline{t_j}}^{\overline{t_j}} F(s) X_s (G(D_j, u_j)(s)) ds + \int_{\underline{t_l}}^{\overline{t_l}} F(s) X_s (D_j F(s) - 1) ds$ where D_l, \dots, D_K , $\underline{s_H} \leq \underline{t_k} \leq \overline{t_k} = \underline{t_{l+1}} \leq \overline{t_{l+1}} = \dots \leq \underline{t_K} \leq \overline{t_K} \leq \overline{s_H}$ are the sequences generated by \overline{s} with respect to μ (Definition 3). Then π is strictly increasing in \overline{s} (i.e. $\pi(\overline{s}) > \pi(\hat{s})$ whenever $\overline{s} > \hat{s}$).

Proof. Consider $\overline{s} < \hat{s}$, and let D_k, \dots, D_K , $\underline{s_H} \leq \underline{t_k} \leq \overline{t_k} = \underline{t_{l+1}} \leq \overline{t_{l+1}} = \dots \leq \underline{t_K} \leq \overline{t_K} \leq \overline{s_H}$, $\hat{D}_l, \dots, \hat{D}_K$, $\underline{s_H} \leq \hat{\underline{t}}_l \leq \hat{\overline{t}}_l = \hat{\underline{t}}_{l+1} \leq \hat{\overline{t}}_{l+1} = \dots \leq \hat{\underline{t}}_K \leq \hat{\overline{t}}_K \leq \overline{s_H}$ be the sequences generated by \overline{s} and \hat{s} respectively with respect to μ . By the argument in the proof of Proposition 4, $l \geq k$, and for each $j > k$, $G(\hat{D}_j, u_j)(\hat{\underline{t}}_j) > G(D_j, u_j)(\underline{t}_j)$ and $\hat{\underline{t}}_j > \underline{t}_j$. By the argument at the end of the Proposition, it thus follows that $\sum_{j=l+1}^k \int_{\underline{t_j}}^{\overline{t_j}} F(s) X_s (G(\hat{D}_j, u_j)(s)) ds + \int_{\hat{\underline{t}}_l}^{\hat{\overline{t}}_l} F(s) X_s (\hat{D}_l F(s) - 1) ds > \int_{\underline{t_k}}^{\overline{t_k}} F(s) X_s (D_k F(s) - 1) ds$. Since, by construction of the sequences, $\int_{\underline{t_j}}^{\overline{t_j}} F(s) X_s (\hat{D}_j F(s) - 1) ds = \int_{\underline{t_j}}^{\overline{t_j}} F(s) X_s (D_j F(s) - 1) ds$ for all $j > k$, the result follows. \square

Since F , X_s and $G(\bullet, \bullet)$ are continuous in s , the sequences constructed in Definition 3 are continuous (pointwise) in \bar{s} , and π in Proposition 5 is continuous in \bar{s} . It follows from Propositions 4 and 5 that any equilibrium under μ is such that the highest hired skill level $\bar{t}_K < \bar{s}_K$. Moreover, by Proposition 5, any such equilibrium is unique; and by the previously noted continuity of π , there exists such an equilibrium. It follows from the form of the solution (Section A.2) that the highest wage is higher under μ' than under μ . This establishes Theorem 1.

A.4 Proof of Theorem 2

Let (c^*, q^*) be the allocation arising from an equilibrium in this market, and let p^* be the equilibrium price vector (giving a price of 1 to the numéraire, and the equilibrium prices to the other commodities, i.e. the good at various inequality levels). Suppose, for reductio, that (c, q) Pareto dominates (c^*, q^*) —it yields higher utility for all consumers and strictly higher utility for a set of consumers of strictly positive measure. It follows that c and c^* differ on a set of strictly positive measure.

For each consumer x of type j , by the utility maximisation (point 2. in the definition of equilibrium), if $U_j(c(x)) > U_j(c^*(x))$, then the budget constraint is not respected: $p^*(c_1(x)) + c_2(x) > \hat{n}$. Moreover, since the utility function (1) is strictly increasing in n , and so is locally non-satiated, if $U_j(c(x)) \geq U_j(c^*(x))$, then $p^*(c_1(x)) + c_2(x) \geq \hat{n}$. So, under the allocation (c, q) :

$$\int_P p^*(c_1(x)) + c_2(x) dx > N\hat{n}$$

At prices p^* , any firm producing the good by hiring the H -type worker at skill level s for wage i makes profits $p^*(i) \cdot F(s) - (i + 1)$. Consider a firm producing the good at inequality level i under (c, q) , and let s be the skill level hired; so q is generated by r with $r_2(s) = i$. For any firm producing the good under (c^*, q^*) with inequality level i , it will hire with skill level $w^{*-1}(i)$ (since q^* is generated by r^* with $r_2^*(s) = w^*(s)$ for all s).³⁰ Since, under the equilibrium allocation, firms maximise profits (point 1. in the definition of equilibrium), it follows that $p^*(i) \cdot F(w^{*-1}(i)) - (i + 1) \geq p^*(i) \cdot F(s) - (i + 1)$ for every s such that $r_2(s) = i$. Since this holds for all firms and all inequality levels, we have that

$$\begin{aligned} \int_I (p^*(i) \cdot F(w^{*-1}(i)) - (i + 1)) \cdot X_{w^{*-1}(i)}(i) di &\geq \int_I \int_{\{s: r_2(s)=i\}} (p^*(i) \cdot F(s) - (i + 1)) \cdot X_s(i) ds di \\ &= \int_I (p^*(i) \cdot q_1(i) - q_2(i)) di \end{aligned}$$

where the equality results from the conditions relating q and r . Since the left hand side is zero (because $p^*(w^*(s)) = \frac{w^*(s)+1}{F(s)}$), the right hand side is non-positive. So, under (c, q)

$$\int_P c_2(x) dx + \int_P p^*(c_1(x)) dx > N\hat{n} \geq N\hat{n} + \int_I (p^*(i) \cdot q_1(i) - q_2(i)) di$$

whence

$$\int_P c_2(x) dx > N\hat{n} - \int_I q_2(i) di \quad (29)$$

³⁰Since w^* is strictly increasing, $w^{*-1}(i)$ is a well-defined, single value.

because, by the market clearing condition for each i , $\int_{c^{-1}(i \times \mathbb{R}_{\geq 0})} p^*(c_1(x)) dx = p^*(i) \cdot q_1(i)$ and hence $\int_P p^*(c_1(x)) dx = \int_I p^*(i) \cdot q_1(i) di$. However (29) violates the market clearing condition for the numéraire. So (c, q) is not a feasible allocation, establishing the result.

References

- Acemoglu, D. and Autor, D. (2011). Chapter 12 - Skills, Tasks and Technologies: Implications for Employment and Earnings. In Card, D. and Ashenfelter, O., editors, *Handbook of Labor Economics*, volume 4, pages 1043–1171. Elsevier.
- AFL-CIO (2019). Company Pay Ratios. <https://aflcio.org/paywatch/company-pay-ratios>.
- Almås, I., Cappelen, A. W., Sørensen, E. Ø., and Tungodden, B. (2010). Fairness and the Development of Inequality Acceptance. *Science*, 328(5982):1176–1178.
- Andorfer, V. A. and Liebe, U. (2012). Research on Fair Trade Consumption—A Review. *Journal of Business Ethics*, 106(4):415–435.
- Antunez, K. and Papuchon, A. (2019). Les français plus sensibles aux inégalités de revenus et plus attachés au maintien des prestations sociales. synthèse des résultats du baromètre d’opinion 2018 de la drees. Technical Report 35. <https://drees.solidarites-sante.gouv.fr/IMG/pdf/dd35.pdf>.
- Baksi, S. and Bose, P. (2007). Credence goods, efficient labelling policies, and regulatory enforcement. *Environmental and Resource Economics*, 37(2):411–430.
- Baumeister, R. F., Bratslavsky, E., Finkenauer, C., and Vohs, K. D. (2001). Bad is stronger than good. *Review of general psychology*, 5(4):323–370.
- Baumol, W. J. and Oates, W. E. (1988). *The Theory of Environmental Policy*. Cambridge university press.
- Chakravarty, S. R. (2009). Inequality, polarization and poverty. *Advances in distributional analysis*. New York.
- Clark, A. E. and d’Ambrosio, C. (2015). Attitudes to income inequality: Experimental and survey evidence. In *Handbook of Income Distribution*, volume 2, pages 1147–1208. Elsevier.
- Cowell, F. (2011). *Measuring Inequality*. Oxford University Press.
- Crifo, P. and Forget, V. D. (2015). The economics of corporate social responsibility: A firm-level perspective survey. *Journal of Economic Surveys*, 29(1):112–130.
- Darby, M. R. and Karni, E. (1973). Free competition and the optimal amount of fraud. *The Journal of law and economics*, 16(1):67–88.
- DREES (2018). Le baromètre d’opinion de la Drees. <https://www.data.gouv.fr/fr/datasets/le-barometre-dopinion-de-la-drees/>.
- Fehr, E. and Schmidt, K. M. (2003). Theories of fairness and reciprocity-evidence and economic applications. In Dewatripont, M., Hansen, L. P., and Turnovsky, S. J., editors, *Advances in Economics and Econometrics*, volume 1 of *Econometric Society Monographs*, pages 208–257.

- Gabaix, X. and Landier, A. (2008). Why has CEO pay increased so much? *The Quarterly Journal of Economics*, 123(1):49–100.
- Hall, B. J. and Murphy, K. J. (2003). The trouble with stock options. *Journal of economic perspectives*, 17(3):49–70.
- Hammond, P. J. (1979). Straightforward Individual Incentive Compatibility in Large Economies. *The Review of Economic Studies*, 46(2):263–282.
- Hill, B. and Lloyd, T. (2020). Are People Willing to Pay for Reduced Inequality? HEC Paris Research Paper Series Forthcoming, HEC Paris. <https://papers.ssrn.com/abstract=3711785>.
- Jensen, M. C. and Murphy, K. J. (1990). Performance pay and top-management incentives. *Journal of political economy*, 98(2):225–264.
- Kahneman, D., Knetsch, J. L., and Thaler, R. H. (1991). Anomalies: The Endowment Effect, Loss Aversion, and Status Quo Bias. *The Journal of Economic Perspectives*, 5(1):193–206. <http://www.jstor.org/stable/1942711>.
- Kahneman, D. and Tversky, A., editors (2000). *Choices, Values, and Frames*. Cambridge University Press.
- Kiatpongsan, S. and Norton, M. I. (2014). How Much (More) Should CEOs Make? A Universal Desire for More Equal Pay. *Perspectives on Psychological Science*, 9(6):587–593.
- Kohut, A. (2013). Economies of emerging markets better rated during difficult times. *Pew Research Center*. <https://www.pewresearch.org/wp-content/uploads/sites/2/2013/05/Pew-Global-Attitudes-Economic-Report-FINAL-May-23-20131.pdf>.
- Kuziemko, I., Norton, M. I., Saez, E., and Stantcheva, S. (2015). How elastic are preferences for redistribution? Evidence from randomized survey experiments. *American Economic Review*, 105(4):1478–1508.
- Lacombe, C. (2019). Les patrons du CAC 40 gagnent en moyenne 277 fois le smic. *L’Obs*. <https://www.nouvelobs.com/economie/20191106.OBS20769/les-patrons-du-cac-40-gagnent-en-moyenne-277-fois-le-smic.html>.
- Lambert, P. (2001). *The Distribution and Redistribution of Income: Third Edition*. Manchester University Press.
- Mas-Colell, A., Whinston, M. D., and Green, J. R. (1995). *Microeconomic Theory*. Oxford University Press, USA.
- Mohan, B., Norton, M. I., and Deshpandé, R. (2015). Paying up for fair pay: Consumers prefer firms with lower CEO-to-worker pay ratios. *Harvard Business School Marketing Unit Working Paper*, (15-091).
- Mohan, B., Schlager, T., Deshpandé, R., and Norton, M. I. (2018). Consumers Avoid Buying From Firms With Higher CEO-to-Worker Pay Ratios. *Journal of Consumer Psychology*, 28(2):344–352.

- Mueller, H. M., Ouimet, P. P., and Simintzi, E. (2017). Wage inequality and firm growth. *American Economic Review*, 107(5):379–83.
- Osberg, L. and Smeeding, T. (2006). “Fair” inequality? Attitudes toward pay differentials: The United States in comparative perspective. *American sociological review*, 71(3):450–473.
- Piketty, T. (2013). *Le Capital Au XXIe Siècle*. Le Seuil.
- Rosen, S. (1981). The economics of superstars. *The American economic review*, 71(5):845–858.
- Rozin, P. and Royzman, E. B. (2001). Negativity bias, negativity dominance, and contagion. *Personality and social psychology review*, 5(4):296–320.
- Shaked, A. and Sutton, J. (1983). Natural Oligopolies. *Econometrica*, 51(5):1469–1483.
- Song, J., Price, D. J., Guvenen, F., Bloom, N., and Von Wachter, T. (2019). Firming up inequality. *The Quarterly journal of economics*, 134(1):1–50.
- Tavernier, J.-L. (2019). Tableaux de l’Economie Française. Technical report, INSEE. <https://www.insee.fr/fr/statistiques/fichier/3696937/TEF2019.pdf>.
- Tervio, M. (2008). The difference that CEOs make: An assignment model approach. *American Economic Review*, 98(3):642–68.
- Thurow, L. C. (1971). The income distribution as a pure public good. *The Quarterly Journal of Economics*, pages 327–336.
- Tirole, J. (1988). *The Theory of Industrial Organization*. MIT press.
- Weitzman, M. L. (1974). Prices vs. quantities. *The review of economic studies*, 41(4):477–491.