

Show me the money – the case for income transparency

Peter Dietsch*

Hoover fellow Avril – Juin 2004

Chaire Hoover, Université catholique de Louvain

peter.dietsch@umontreal.ca

“...he set about testing the ring to see if it really had this power, and always with the same result: according as he turned the bezel inside or out he vanished and reappeared. After this discovery he contrived to be one of the messengers sent to the court. There he seduced the Queen, and with her help murdered the King and seized the throne.” (Plato, *The Republic*, book II, 360 a,b)

Can we produce a justification of just behaviour that is not merely instrumental? Why should leading a just life as well as the actions it calls for be preferred over the perhaps more enticing alternatives? Thus, in general terms, runs the challenge set up by Glaucon in the early stages of Plato's *Republic*, a challenge that Socrates subsequently attempts to meet throughout this dialogue, which is often depicted as the cradle of political and moral philosophy. In particular, Glaucon advances three claims. First, norms of justice in a society come about not because its members value justice as such, but because most of them lack the power to secure unjust advantages for themselves and hence accept a system of rules of justice as a compromise. Second, illustrated by the story – quoted above – of the shepherd Gyges, who finds a magic ring that renders him invisible, no man will have the motivational strength to refrain from gaining unjust advantages if he can be sure to get away with it. And finally, again working on the assumption that the injustice is not detected, is it not obvious that the benefits of gaining unjust advantages by far outweigh the costs? Is it not obvious, in other words, that unjust behaviour pays?

All of these questions still preoccupy philosophy today, and the available answers are just as controversial as 2'500 years ago. To give a selective list of examples, Glaucon's first claim about the genesis of norms of justice can be viewed as a precursor to social contract theory. Concerning the second set of issues, utilitarians have posited a principle of *sympathy*, and John Rawls a *sense of justice*, as alternative motivational anchors of justice within social agents that will help them overcome the temptations of unjust action. As to the question whether injustice pays, a suitably construed Prisoners' Dilemma is frequently used to sanction an unjust path of action as individually rational, if socially suboptimal.

Any attempt to retackle all these issues at once would have to count as ambitious, to say the least. My objectives in this short paper are more limited. Taking the desirability of a just society for granted, and hence abstracting from one fundamental question of justification, I propose to learn a lesson from Glaucon's story about the ring of Gyges. Whereas Glaucon emphasises the fact that the vacuum of inaccountability that the ring produces for Gyges drives the latter towards unjust actions, the flipside of the story tells us that visibility increases accountability. No rule of justice will be effective if its infractions can be kept secret. Yet, the more those actions answerable to an existing system of justice are visible or public, the higher the pressure for

members of society to conform to that system. I will use this insight to argue for a complete transparency of wage incomes.¹ Though similar arguments for transparency can no doubt be made in other contexts, too, this paper exclusively focuses on income distribution.²

If transparency were only beneficial when it fosters compliance with an *existing* set of norms of justice, the case for it would be relatively weak. Therefore, I will also aim to show, under a second scenario, that transparency is likely to have an indirect, positive effect on the justice of income distribution even if no prior consensus exists about what constitutes a just such distribution.

In section 2, we will distinguish two different ways to justify income transparency, and conclude that the degree of transparency recommended varies depending on which justification we endorse. Finally, in section 3, we will consider, and find wanting, a series of objections to the project of rendering income information transparent.

1. The redistributive effects of publicity

There are few secrets people guard more anxiously than their pay details. Even among work colleagues, open discussion about the size of their paycheques is the exception. This strange taboo surrounding earnings is both puzzling and arguably presents a serious obstacle to distributive justice. If one's income is a fair reflection of one's contribution to the productive process, why be protective about this information?

The central hypothesis of the paper is straightforward: People will be more hesitant to walk away with an unfair share if their income is public information. Let us assume, realistically, that such a publicity requirement will not have any direct redistributive effects, i.e. it will not induce those on unfairly high incomes to hand a percentage of their revenue to those who are not paid their due under the *status quo*. However, income transparency can prove indirectly instrumental to the cause of distributive justice by enlarging the domain of politically feasible social policies. In other words, if the disproportionality between one's income and that of others is for all to see, this is unlikely to turn the person in question into an altruist who gives large sums to others or to charity, but it will increase her acceptance of redistributive policies.

Regarding this effect of publicity on the public acceptance of policies in the name of economic justice, I see two idealised scenarios. In one case, those privileged by the current distribution of earnings have been aware of the injustice all along, but have failed to act on this awareness. In the other case, they in fact believe the *status quo* to be just, which means there exists a deeper, moral, disagreement about what constitutes a just distribution of income. We will now treat these two scenarios in turn and see that a publicity requirement for incomes is beneficial from the perspective of justice in both cases.

* I would like to thank the members of the Chaire Hoover for the warm welcome and stimulating academic environment they provided during my three months in Louvain-la-Neuve. Thank you to Axel Gosseries and Christian Nadeau for their comments on this paper.

¹ Unless otherwise stated, 'income' refers to wage income throughout the paper.

² Cf. for instance Geoffrey Brennan and Philip Pettit's argument for "Unveiling the Vote" (1990). Abolishing the secrecy of the ballot, so their claim goes, will increase people's disposition to vote according to what they believe to be in the public interest.

Scenario 1

Under the first scenario, two questions arise. First, it seems legitimate to wonder why, if the privileged have been aware of the injustice of current distributive arrangements, they still failed to change them. Second, we will want to know how to evaluate this omission from a moral point of view.

On the latter question, moral theories have traditionally offered two kinds of reply. (cf. Nagel, 1991: 49) Consequentialist theories, which admit only *agent-neutral* reasons to morality, will classify the omission as morally wrong. They would have advocated a realignment of incomes all along. Individual rights theories, on the other hand, which evaluate conduct on the basis of its conformity to certain *agent-relative* reasons, will in the extreme not object to the omission by the privileged at all. They might for instance say that, though the distributive system as such is to be considered unjust, this does not entail an individual responsibility to change it. On this premise, the compliance of the privileged with the rules of an unjust system may be classified as morally unobjectionable.

Neither of these, admittedly caricatured, alternatives is satisfactory. Whereas the agent-neutral approach imposes excessive moral demands on the privileged in our example, the agent-relative approach lets them off the hook too lightly.³ One proposal to reconcile agent-neutral and agent-relative considerations, and thereby overcome their respective weaknesses, has been put forward by Thomas Nagel in his book *Equality and Partiality*. Nagel maintains that both an impersonal and a personal viewpoint are constitutive of our moral make-up. This conceptualisation allows a more sophisticated evaluation of the omission by the privileged. Asserting that, from their personal point of view, there exists no obligation for the privileged to go beyond the call of institutionalised duty, is compatible with acknowledging that the institutional structure as such is unjust and that, from the impersonal point of view, the privileged are under *some* obligation to change it. There exists, in other words, a certain negative responsibility on the part of all members of the society in question to ensure that the social arrangements they live under are just. For completeness' sake, I should mention that an alternative way for the privileged to act on the impersonal point of view consists in donating the difference between their actual income and the sum they would earn under a just social arrangement to those hard done by under the *status quo*. As indicated before, we will bracket this variant here.

Within the framework of Nagel's moral theory, the past failure of the privileged to act on their awareness of the injustice can be explained as follows. Despite acknowledging both the personal and the impersonal point of view, the former clearly dominates motivationally. One might say that the prevailing circumstances do not yield a sufficiently strong incentive for the privileged to act on the more demanding impersonal perspective.

This is where our proposed publicity requirement comes in. It will tilt, so the claim goes, the motivational balance in favour of the impersonal point of view, and thereby provide an incentive to act upon it that was absent before. To be sure, publicity does not lead to a change in moral attitude. From the perspective of the individual moral agent, for instance of a member of the group unjustly privileged by the present income distribution, the personal viewpoint of merely conforming to the system and the impersonal viewpoint of changing it stay exactly the same. Publicity merely changes the extent to which these two components of one's moral perspective are

³ For a detailed discussion of the problem of excessive demands in non-ideal theory, cf. Liam Murphy (2000).

acted upon. This shift is triggered by the fundamental human desire to see one's actions appraised by other members of the community.⁴

In other words, the reverse of the Gyges-effect kicks in. The more my actions are observable by other members of society, the more I will tend to conform to the norms of justice of the community. At least, the conceptual framework set out over the last few paragraphs makes a reaction of this sort look very plausible. However, given the argument I want to make in this paper, the claim just put forward ultimately has to count as an empirical one. Whether or not making incomes transparent is a good idea depends on whether it will *actually* bring people's behaviour more in line with the demands of justice. The answer to this question, rather unsatisfactorily, once again is "It depends."

One argumentative strategy at this point would be to try and find examples of situations where income figures have indeed been made public, and look at the distributive effect of these measures. Yet, precisely because the distributive effects of transparency depend on situational factors, simply citing examples that confirm my hypothesis would be inconclusive if not outright misleading. Instead, what would be required is an analysis of the situational factors that determine what kind of distributive effects transparency is likely to have. One reasonable, albeit very general, proxy to capture the relevant situational factors, I suggest, is the level of social cohesion that a society enjoys.⁵ If the ties of civil society are relatively strong, then transparency will increase their motivational effectiveness. If they are weak, then transparency is likely to trigger tension rather than solidarity. I believe that the social environment in most Western European countries, for instance, is such that transparency would have a positive effect on distributive justice. On the other hand, increasing transparency in post-Soviet Russia, where in some instances the mere fact of differential income based on effort led to resentment, would probably have been a bad idea.

These remarks are admittedly speculative, and could be seen as nothing more than hand-waiving. To some extent, we will address this criticism in more detail in section 3, where we will look at some specific objections to the proposal of income transparency. For now, and subject to the qualification just introduced, we can conclude that we have good reason to expect a publicity requirement to have a positive effect on income distribution when a background consensus exists about what constitutes a just such distribution.

Scenario 2

Unfortunately, the assumptions made under our first scenario suffer from a fatal flaw. Granting the Nagelian picture that posits both a personal and an impersonal point of view as constitutive elements of our moral make-up, who said that the impersonal points of view will converge across persons? In fact, people differ not only in the interests that inform their personal point of view, but they also disagree about what would constitute a just distribution of income from an impartial perspective. Among those privileged by what we stipulated to be an

⁴ In their recent book *The Economy of Esteem* (2004), Geoffrey Brennan and Philip Pettit claim that this desire to be viewed in a positive light by others represents a major force of human motivation, and one that has been neglected by contemporary social theorists. As an illustration, and one that speaks to our context of transparency, they cite the "results of a study of New York public washrooms which revealed that whereas only about 35 per cent of women washed their hands after using the toilet when there was no one else present in the washroom, nearly 80 per cent did so when there was someone else there." (2004: 32)

⁵ Social cohesion needs to be distinguished from, and does not entail, moral consensus. The relevance of this distinction will become clear in the discussion of Scenario 2.

unjust distribution of income, we will find quite a few who maintain that the *status quo* is in fact just. They may, for instance, defend a desert-based theory of distributive justice, whose desert basis is the share each member of the community contributes to the production of wealth; on the basis of this theory, they believe the income they receive to be justified. Other theories of justice combined with a similar configuration of beliefs could be employed to the same effect.

Under these conditions, the proposed publicity requirement for income information *prima facie* looks ineffective. Even if it increases the relative motivating force of the impersonal point of view, the latter will not induce any change, since, as far as the privileged are concerned, they endorse the present distribution of income as just. Does this imply that the beneficial effects of income transparency are limited to the unrealistically idealised scenario 1?

I will attempt to convince you that this conclusion would be premature. Albeit in a more roundabout way, making income information public can still enlarge the domain of redistributive measures that are politically feasible. Imagine a population with two equally strong subgroups. For simplicity, the members of the privileged group P earn €5'000 a month, whereas the members of the low income group L earn €2'000 a month. Under the *status quo*, no one has an exact idea of what either their peers or members of the other group are earning, though a glance at the standard of living of members of the P-group sometimes makes representatives of the L-group wonder whether they are being short-changed. Now a publicity requirement for incomes is introduced and the cards are being put on the table. It soon emerges that the Ps think that the income structure is justified as it stands, whereas the Ls believe that the income gap of €3'000 exaggerates the differential contribution to production of the two groups.⁶ The Ls appreciate that the Ps' contribution is relatively more important, but would deem an income difference of €2'000 appropriate, calling for an adjustment of the income structure to €4'500 and €2'500 respectively.⁷

If nothing else, the publicity requirement will encourage a public debate – or intensify an existing one – between the two groups about what constitutes a just distribution of income and why. If it turns out that the better arguments in this debate are on the side of the Ls, the Ps will have a harder time justifying the magnitude of their income premium. The door for a compromise will be pushed open.

But hang on, since when do we determine wages by consensus? In practice, the relative wages of the two groups are likely to be settled by a bargaining process, and if the Ps hold a sufficiently strong bargaining position to obtain their €5'000 a month, the question is settled. It seems as if my consensus model has little practical relevance. Even in those countries where wage negotiations have traditionally been more consensus-oriented, like the Scandinavian countries or Germany, a more confrontational and pure form of bargaining is taking over.

In order to avoid being drawn into an ideological sling-match between free-marketeters and trade unionists, let us look at these alternative models from a relatively safe theoretical distance. Jon Elster tells us that “[w]hen a group of equal individuals are to make a decision on a matter that concerns them all and the initial distribution of opinion falls short of consensus, they can go about it in three different ways: arguing, bargaining, and voting.” (Elster, 1998: 5) Two additions to this classification are worth mentioning, both of which are also highlighted by

⁶ Assume that everyone in this particular society subscribed to a theory of justice like the one mentioned above: Income should be distributed according to desert, where the desert basis is given by the individual shares contributed to the production of wealth.

⁷ For simplicity, we assume that income distribution is a zero-sum game. The discussion of incentive effects further down relaxes this assumption somewhat.

Elster. First, in practice, these decision procedures are rarely instantiated in their pure form. Second, if we decide in favour of a certain configuration of decision procedures in a particular case, we will need to justify our choice. Whether we maintain that wages should be set by bargaining, arguing, or a particular mixture of the two – voting, after all, is an unlikely candidate in our context – we will need to buttress our case with arguments. I interpret my proposal of a publicity requirement for income information not as a call to move from bargaining to arguing, but as one to move from ‘more bargaining and less arguing’ to ‘less bargaining and more arguing.’ In support of this gradual shift towards a more consensus-oriented model of wage determination, I cite my conjecture that it will lead to a more just distribution of income.

To be sure, a publicity requirement does not preclude that wages are primarily determined by bargaining. Yet, it forces all those participating in the bargaining process to justify their position in public. This has two important consequences. First, the publicity requirement acts as a filter for the arguments that are admitted to the bargaining process in the first place. For example, it would not be acceptable for the Ps in our example to justify the current income structure by appealing to the fact that it has been in place for more than three decades. More controversially, neither could they point to a neighbouring country N, where the income differential between the Ps and the Ls is even more pronounced. This is not to deny that, under conditions of mobility, the situation in country N will in effect substantially strengthen their bargaining position. Yet, this fact will be redundant in a context of justification. We will briefly come back to this aspect under the heading of incentives later in the text.

Before, we need to spell out the second, and crucial, consequence of the publicity requirement. The gradual shift from bargaining to arguing about the income structure significantly changes the mode of interaction between the respective impersonal points of view, or theories of justice, held by the participating parties. Whereas it is an essential feature of a bargaining process to conceal and frequently even misrepresent one’s true preferences or moral attitudes, a process of arguing pulls these pieces of information out in the open. If, as we have assumed, the debate starts out from a position short of consensus, some of these attitudes are bound to change in the process. “The transformation of preferences through rational deliberation is the ostensible goal of arguing.” (Elster, 1998: 6)⁸ Assuming, in admittedly idealistic fashion, that the less reasonable perspectives on justice are more likely to be weeded out in this process, the publicity requirement indirectly leads to a relative convergence of opinion endorsing a just distribution of income. Similar to scenario 1, this is an empirical claim, and the same qualifications apply.

We will come back to the question in what sense exactly the result of this deliberation process can be classified as just. For now, we can conclude that the publicity requirement, introduced under conditions of moral disagreement about distributive issues, will trigger a process of deliberation, which, by transforming the least justified perspectives on justice, will pave the way for a more just distribution of income. In our example introduced above, if the arguments of the Ls indeed prove to be the more robust, the outcome of this process would see the Ps earning a sum between €4’500 and €5’000, whereas the Ls get between €2’000 and €2’500. The stronger the case of the Ls that the *status quo* is unjust, the closer their income will rise towards €2’500; the

⁸ For a view that is sceptical of these alleged benefits of publicity, cf. Daniel Naurin (2003). He suggests that publicity will indeed favour arguing over bargaining, but that the kind of arguing that results is a purely formal one. Arguing of this kind remains superficial. Yet, “...why would we want publicity if it produces hypocrisy rather than deliberation and sincere public-regardingness?” (2003: 26) This objection poses a problem not only for my proposal, but for any account that relies on public deliberation. We bracket this problem here.

stronger the bargaining position of the Ps, taking into account the arguments presented in the debate, the closer their income will stay to €5'000.

Let me add two sets of comments to this second, more realistic, scenario about the consequences of introducing a publicity requirement for income information. The first comes back to the controversial question of whether the story I have told is “incentive-compatible.” I suggested, recall, that in the process of deliberation about the just income structure, the Ps could not appeal to the fact that in a neighbouring country, the income differential between two equivalently positioned groups was even larger. Two things need to be said to clarify this statement. First, as emphasised before, this is not to deny that in effect, the situation described would indeed boost the bargaining power of a member of group P who is prepared to dislocate for higher pay. Yet, second, it is equally clear that pay across the border cannot count as an argument in trying to answer the normative question about adequate pay at home. If the process of deliberation triggered by the publicity requirement works as sketched, and if the arguments of the Ls indeed prove to be the more robust, then the moral attitudes of the Ps will be transformed so that they themselves would regard an appeal to the income structure across the border as off limits. As a concrete illustration that this scenario is not utopian, consider for instance the relative restraint exercised by Scandinavian managers in their wage demands. Whether the adjustment of the income structure following the publicising of incomes is incentive-compatible depends on whether the participating parties are prepared to put in the same amount of work after the changes have taken place. If the deliberation process transforms the perspective on justice of the Ps, their incentives will be preserved despite the adjustment of the wage structure. Having said that, this question of the incentive-compatibility, especially in the context of adjustments of the income structure, deserves more attention than we can devote to it in this paper.

The second comment picks up on the conceptual patronage behind the process of deliberation that kicks in under our second scenario of making incomes transparent. The process is modelled on the concept of deliberative democracy.⁹ The fundamental idea is to embed the voting process constitutive of democracy in a public discussion about the social choices to be made. Using Elster’s classification introduced above, the decision procedure hence represents a mixture of arguing and voting. Similarly, in our context of income distribution, the decision procedure of bargaining is complemented by a process of public deliberation.¹⁰

With such a process of public deliberation in mind, the hopes we place in transparency are fundamentally different from those we have under our scenario 1 above. Back then, we assumed that a consensus about what constitutes a just distribution of income lingered under the surface of egotistic actions already. The task of transparency consisted in exposing this consensus and closing the gap between people’s moral convictions and their actions. Such an optimistic picture of a consensus about moral views is typical of, if not exclusive to, theories which take the principles of justice to be objective in some sense. Transparency, under such circumstances, can be viewed as a catalyst for promoting these principles.

⁹ For one of deliberative democracy’s main contemporary defenders, cf. for instance Jürgen Habermas and his *Between facts and norms: contributions to a discourse theory of law and democracy* (1996).

¹⁰ It is worth making explicit that there are two different kinds of publicity or transparency at stake now. On the one hand, there is the accessibility of income figures, which lies at the heart of this paper. On the other hand, there is the *public* character of the process of deliberation that is constitutive of the kind of democracy and collective decision-making championed by Habermas, Elster, and others. Transparency of the former kind appears to be a precondition of a meaningful deliberative democracy.

The second scenario, on the other hand, entails a *procedural* view about the epistemological status of principles of justice. Transparency not only encourages just action, but it plays a crucial role in the determination of what is to count as a just action in the first place. In the present context, what constitutes a just distribution of income is not determined *ex ante*, but arrived at through a process of public reason. At the very least, if one makes the realistic assumption that full consensus will always prove elusive, the range of eligible principles will be narrowed down.

As an illustration of how transparency fuels this process of public reason, consider the role publicity plays in the theory of justice of John Rawls. He points out that the public recognition of a conception of justice is crucial in making it stable, by bringing about the corresponding sense of justice (cf. Rawls, 1999: 154). The ‘bringing about of the corresponding sense of justice’ mirrors the transformation of people’s perspective on justice, or their impersonal point of view.¹¹

To conclude the discussion of our two alternative scenarios, reality presumably lies somewhere in between. Transparency of incomes, in other words, can be instrumental both in encouraging compliance with existing principles of justice and in defining new ones.

2. Transparency of what, and why?

So far, a cynical remark that this paper does not practice what it preaches would not be entirely misplaced. I have not said anything about the precise meaning of the phrase “transparency of incomes” – a phrase, after all, that can be interpreted in a number of ways. The most important distinction to be drawn here is that between transparency of the structure of wages – or pay scale – on the one hand, which we will call *anonymous* transparency, and *individual* transparency on the other, which means that I have access to the figures on my neighbour’s paycheque. Individual transparency, it almost goes without saying, is the more demanding of the two.¹²

The position we take on these issues fundamentally depends on the kind of justification we give for transparency of incomes in the first place. In this section, we will distinguish two such justifications, and establish how far-reaching a notion of transparency they can legitimise. Subsequently, in the final section of the paper, the distinction made here will prove useful in evaluating a number of objections that can be made against the project of rendering income information public.

First, transparency of incomes can be justified by appeal to individual rights. On this account, we start from the premise that the productive process in our society is a co-operative venture for mutual advantage.¹³ Members of society specialise in different tasks, and are remunerated according to how much their particular activity is

¹¹ On a more speculative interpretation of Rawls, one could even find a reflection of the two scenarios distinguished in this paper in his work. *A Theory of Justice* adopts the perspective of scenario 1. Assume a well-ordered society that has agreed on a particular set of principles of justice. The public recognition of these principles serves to ensure that people’s sense of justice – in Nagelian terms, their impersonal point of view – actually matches those principles. The initial moral disagreement of scenario 2, on the other hand, takes centre stage in Rawls’ *Political Liberalism*. In this context, the primary function of publicity can be interpreted as that of maximising the common ground between conflicting moral positions.

¹² Our example under scenario 2 in the last section dodged this issue by assuming that there are only two different wages: €5’000 for the Ps, €2’000 for the Ls. This simplistic assumption will now be relaxed.

¹³ Cf. for instance Rawls (1999).

valued either by others through the market mechanism, or by a third-party judgement, as for instance in the case of public service jobs. Once we accept that the strong interdependence at the heart of our economic system turns the productive process into a co-operative venture, which is hard to deny, it seems reasonable to ascribe a right to each participating individual to know how their contribution to this venture is valued relative to that of others.

Yet, it is unlikely that this right can be made strong enough to support a form of individual transparency of incomes. The ‘right to know’ about my status within the productive process can be satisfied by pointing to the remuneration of *kinds* of contributions. As long as I know what the typical doctor, baker, or lawyer earn in relation to me as philosopher, my ‘right to know’ is not infringed. In contrast, a right to have a look at my neighbour’s paycheque or other individualised income information can hardly be justified in this manner. In sum, the individual rights justification of income transparency emerges as endorsing the *status quo*. After all, information about the typical incomes of various professions is readily available in our society. On an individual rights account, then, it will be hard to justify any level of income transparency that goes beyond the *status quo* in our societies today.

We now turn to the second possible justification of income transparency, which is consequentialist in nature. You will note that advocating income transparency on the basis of its beneficial effects better captures the motivation implicit in this paper so far. In section 1, we argued that transparency will rub off positively on distributive justice. We now ask what kind of transparency is required to achieve this aim.

As a preliminary remark in this context, it is worth pointing out that the promotion of distributive justice is not the sole positive consequence of income transparency. Economic theory highlights the importance of the availability of price information for the smooth functioning of our economy. Wage incomes represent one, central, category of such prices. Even though the assumption of perfect information made in neoclassical economics represents an idealisation, it is nonetheless true that relative prices form the prime source of information for the decision making of economic agents. Whether I choose to train as a lawyer, a doctor, or a teacher at least in part depends on the financial prospects these career options offer.

However, it is plain to see that this argument for income transparency from economic efficiency once again only extends to the anonymous version. In fact, resting one’s evaluation of various career options on pieces of *individual* wage information may be rather misleading, particularly under the typical conditions of a heterogeneous labour market. Modelling my expectations of what doctors earn on the revenues of a very successful heart surgeon, for instance, would have to count as optimistic.

So much for why economic efficiency requires transparency. Yet what about the argument from distributive justice? Can we show that for the positive impact on distributive justice to occur, a more demanding form of transparency is necessary than what we already have under the *status quo*? Can we make a case for individual rather than anonymous transparency?

I believe the answer is yes. In making this case, I will borrow from a recent paper by Jeremy Waldron, in which he defends “The Primacy of Justice” (2003). Waldron invites us to consider what our evaluation of the distributive justice or injustice of a social arrangement is based on. The answer, he convincingly argues, is that we look for distributive information across *individuals*. If someone gives us the total utility of a population, or provides the latest GDP (gross domestic product) figures for a country, this information is irrelevant from the perspective of distributive justice. Importantly, the same holds true for average utility or GDP per head. These pieces of information do not tell us anything about how the individuals fare within the population or country in question. Yet it is such information about individuals that distributive justice is essentially concerned with.

For Waldron, transparency of individual outcomes is not merely conducive to distributive justice, but it is a constitutive element of making evaluative statements about distribution. Transferring this lesson to the context of our paper, for the members of society to be able to make distributive judgements at all, information about incomes necessarily has to go down to the individual level. Anonymous transparency, similar to average utility or GDP per head, would be insufficient for this purpose. In other words, we need to go beyond occupational income information that tells us the *typical* wage of, for instance, a doctor. To make distributive judgements, we have to “ensure that the system is made responsive to microdifferences in individual circumstances,” (Waldron, 2003: 277) and we therefore need access to individual income information.

This argument gains even more momentum in societies like ours where occupational income information is in fact a very poor indicator of individual incomes. Not only do incomes within professions vary considerably – think of the successful heart surgeon as opposed to the general practitioner whose clientele mainly relies on the public health care system; but it is also very common that incomes differ where one would assume them to be the same – think of the two employees of the same firm who do exactly the same job for the same nominal wage, but one of whom also gets to live in a company flat and to drive a company car. Such fringe benefits often make the wage discrepancies look less considerable than they in fact are. There may be a justification for some of these differences, and in cases where there is, introducing transparency may not change much. Yet, the mere fact that occupational income information does not reveal fringe benefits represents an additional reason to call for individual income information that includes them.

In sum, for transparency to have a positive effect on distributive justice, it necessarily needs to be individual rather than anonymous in kind. This conclusion holds independently of whether there exists a prior moral consensus about what constitutes a just distribution at the time when transparency is introduced.¹⁴

A boundary problem?

Before turning to a series of objections that can be mounted against the case for income transparency as such, we need to briefly discuss a different kind of worry, namely one that concerns the scope of transparency. In order to make distributive judgements, do we not need to know a lot more than mere income figures? In particular, do we not require more information on both the content of people’s jobs and on how well they are performing in them?

Arguably, these questions take up Robert Nozick’s classic argument that the assessment of distributions from the perspective of justice crucially depends on how the distribution in question came about. Nozick rightly insists that we cannot assess a distribution without knowing something about its history. This opens up the following danger for my argument: Can we coherently call for a transparency in the realm of incomes without also demanding information about job descriptions and performance? If not, it would seem that advocating transparency of incomes gets us onto a slippery slope towards an unpleasantly intrusive system of work surveillance.

I believe this threat can be defused. Let us look at job descriptions and performance in turn. There exists a fundamental difference between job descriptions and income information in that the former are not plagued by the same kind of opacity that motivates this paper in the case of the latter. In other words, of course we need information about the content of people’s jobs to make distributive evaluations. Yet, such information is

¹⁴ Cf. the two scenarios distinguished in section 1.

available already. One may have to look or ask for it, but it is certainly not surrounded by the same aura of secrecy as income information. If I want to find out what the daily schedule of the heart surgeon, teacher, or airline employee actually looks like, I will not run into any significant barriers. Hence, on this front, my account does not advocate anything more or less intrusive than the *status quo*.

Prima facie, the case of information on job performance looks trickier. Once again, such information is indeed indispensable to make distributive judgements. Yet, especially when working in a team, shirking is fairly easy to cover up. Our evaluation about whether a certain income is appropriate to any given job will always have to rely on proxies to assess how well the job is actually done. Fortunately, however, my case for income transparency is not the only approach to face this problem. The immense literature on the principal-agent-relationship in economic theory is testimony to its pervasiveness. So what does this imply for my case in favour of income transparency? Well, once we acknowledge that performance evaluation happens on the basis of proxies, we can once again make a case that such information is readily available. One can find out, for instance, what kinds of hours the surgeon or the teacher put in, or what their reputation is among their clients and colleagues. In contrast to the case of job content, the degree of transparency with respect to performance that prevails under the *status quo* may not seem satisfactory. Yet, this is due to the nature of the principal-agent relationship, a relationship that by definition enjoys limited transparency.

In sum, we may draw the following conclusions. First, there exists a difference between income figures on the one hand and job descriptions and performance on the other; whereas the former are surrounded by secrecy, there are no significant barriers to the flow of information on the latter two. Second, this makes it both coherent and possible to limit the scope of our call for more transparency to incomes. Transparency does not necessarily slide into intrusiveness.

3. Contemplating concealment

I see three principal objections to my proposal. Respectively, they argue that we should not push for income transparency because it interferes with people's rights, because it has negative consequences, or because transparency already exists and has little of the desired effect. The purpose of this section is to demonstrate that none of these objections stand up to scrutiny. In doing so, we will rely on the two possible justifications of transparency – based on a 'right to know' and a consequentialist argument respectively – distinguished in the previous section.

To see where the first objection is coming from, imagine a proposal to make income information public were tabled in the parliament of any European country today. It takes little political savvy to predict the outcry such a step would provoke from civil rights groups. Information about incomes, one can reasonably anticipate them to maintain, is private and should hence be protected. In particular, privacy can be invoked at two different levels. On the one hand, put yourself in the shoes of the recipient of any given salary; she will claim that this insight into her overall disposable income is nobody else's business. On the other hand, think about the decision makers in companies, who will argue that the income structure they decide on for their employees is a matter of private management discretion. In the name of both of these perspectives, civil rights groups might appeal to a liberal ideal to justify a concealment of wage information. Borrowing a phrase from Thomas Nagel, they will submit that "no more should be subjected to the demands of public response than is necessary for the requirements of collective life." (Nagel, 1998: 14) Should we accept this line of reasoning?

I think the answer is a clear no. Think back to the rights-based justification of transparency set out in the previous section. Contributing to the productive process of society is a paradigm example of a *public* activity, where “public” in this context refers to the interdependence of the parties involved in the productive process.¹⁵ Doing a particular job amounts to contributing one tessera to a much bigger mosaic. However, as we saw in the previous section, the public character of work is unlikely to be able to justify more than *anonymous* transparency. Could it be that the case for concealment is more successful in defeating *individual* transparency? I believe not. Though it may not be possible to justify individual transparency in a rights-based framework, once we accept the public character of work, a right to conceal personal income information would look equally odd. With respect to individual transparency, a rights-based approach simply emerges as little informative.

We must not let ourselves be swayed by the concealment of income information to date, which is a rather contingent kind of convention. As a parallel case and precedent to my call for income transparency, consider John Stuart Mill’s case in favour of public voting.¹⁶ Emphasising the fact that other members of society take a justified interest in this important public duty to be performed diligently, he urged to abolish the box of confidentiality around the ballot.¹⁷ Whereas in the case of voting, a convincing reply to Mill can be based on likely negative consequences from a public ballot, like the exercise of pressure on individual voters, the discussion of our next objection will suggest that parallel drawbacks are absent in the case of income transparency.

The second objection to income transparency more directly tackles the consequentialist argument in its favour that lies at the heart of this paper. Might a publicity requirement for income information have negative consequences that outweigh the described benefits from the perspective of justice? Several candidates spring to mind. First, it could be said that revealing all income information will prove demeaning for those on the lowest incomes. But would this really be any more demeaning than maintaining a hypocritical silence about the magnitude of income differentials? I doubt it. Besides, if our predictions about the re-ad-“justment” encouraged by transparency are correct, the process of public deliberation would not only reduce the gap between high and low incomes, but also give the remaining differences a legitimacy that would reduce the danger of a demeaning effect. Second, will a society whose members wear their remuneration on their sleeves not be much more prone to the sentiment of greed? This, in the end, is once again an empirical question, which we can only give a speculative answer to. Given that the symptoms of hugely different incomes like big houses, fancy cars, or even

¹⁵ Recall that the scope of my argument is limited to labour income. Civil rights groups may have a stronger stance when it comes to revealing information about wealth rather than income, but I do not have to discuss this question here.

¹⁶ As mentioned earlier, Mill’s case is taken up again, albeit in slightly modified form – in the contemporary literature by Brennan and Pettit (1990).

¹⁷ “In any political election, even by universal suffrage (and still more obviously in the case of a restricted suffrage), the voter is under an absolute moral obligation to consider the interest of the public, not his private advantage, and give his vote, to the best of his judgment, exactly as he would be bound to it if he were the sole voter, and the election depended on him alone. This being admitted, it is at least a *prima facie* consequence that the duty of voting, like any other public duty, should be performed under the eye and criticism of the public; every one of whom has not only an interest in its performance, but a good title to consider himself wronged if it is performed otherwise than honestly and carefully. Undoubtedly neither this nor any other maxim of political morality is absolutely inviolable; it may be overruled by still more cogent considerations. But its weight is such that the cases which admit of a departure from it must be of a strikingly exceptional character.” (Mill, 1865: 201-2)

yachts, are already on display, I would not expect openness about the underlying spending power to make a significant difference in this respect. Finally, there is the fear that unveiling the extent of wage discrepancies could lead to losses in economic productivity as those who feel treated unfairly reduce their work effort. As the psychological analysis of economic relationships tells us, shirking represents one possible response to the perception of being treated unfairly.¹⁸ To gauge the effect of income transparency on work effort, we have to distinguish two effects. In the short term, there is the question to what extent income transparency will change people's perception of the fairness of economic relationships. If transparency triggers an outrage at the injustice that had previously been kept under the cover, then a reduction in work effort is indeed likely. If, on the other hand, transparency merely confirms a pre-existing impression of injustice, then the short-term consequences of transparency may prove negligible. In the long term, however, transparency may in fact boost productivity rather than dampen it. Provided the argument made in this paper that transparency will lead to an ad-"justment" of the wage structure is correct, equity theory would predict those who benefit from the reduction in injustice to raise their work effort.

Summing up this second objection, then, the potential benefits of income transparency seem to outweigh any negative consequences. Similar to some of the arguments made in section 1 of the paper, this conclusion involves a fair amount of empirical claims. Ultimately, the proof of the pudding will be in the eating. The best we can do in advance is to shed some light on the factors that determine what kind of distributive consequences transparency is likely to have, and to investigate whether or not these factors obtain. There certainly remains much empirical work to be done here, yet I hope that this paper has contributed to clarifying some of the conceptual issues at hand.

Finally, let us look at the third objection to making income figures publicly available. Rather than taking issue with income transparency as such, this objection questions its effectiveness in promoting distributive justice. This claim can be broken down into the assertion that income transparency exists already, and the observation that it evidently lacks the beneficial impact on distributive justice I have tried to attribute to it. As far as the first part of this claim is concerned, it is true that *some* income transparency exists already. Executive salaries in companies listed on the stock exchange, for instance, need to be approved by their shareholders, and are therefore just as accessible to the public as the outcomes of unionised wage bargaining. Yet, if transparency comes in degrees, we are still moving through some rather opaque territory. In order to make my proposal somewhat more concrete, let me suggest that genuine transparency would be achieved when salaries of employees are displayed on a company's website. This would be a huge step forward compared to the incomplete information on the matter that one currently has to dig out from the back pages of annual reports.

As soon as one admits that today's transparency is a partial one at best, the observation that it is ineffective from the point of view of distributive justice looks misplaced. Though cases of executive pay packages being voted down are still very rare, they are on the increase.¹⁹ For that part of the income iceberg that is above the surface, these cases count as evidence for the effectiveness of transparency as promoting justice. The task now consists in uncovering the rest of the iceberg. A public debate on the distribution of income takes time to develop. We have entered this process already, but there still remains a long way to go.

¹⁸ For a characterisation of this argument as well as an overview of the literature on 'equity theory,' cf. Alan Garner (1986). The psychological analysis of economic relationships typically aims to provide explanations of economic behaviour where orthodox microeconomics fails.

¹⁹ Cf. for instance "Fat cats cornered" in *The Economist*, 20th May 2003.

Conclusion

If the idea of income transparency is as effective as it is simple, then it seems bewildering that we have not implemented it long ago. The answer, as so often, may lie in the politically explosive character of the proposal. Promoting it is likely to attract stiff resistance from all those who anticipate to lose out from the measure, or who fear to become targeted by resentment. In part, their resistance is informed by precisely those convictions about what constitutes distributive justice in our society that a process of public reason would envisage to transform. Overcoming this dilemma and pushing for a publicity requirement demands a healthy dose of political leadership. The arguments presented in this paper suggest that it is worth a try.

Bibliography

- Geoffrey Brennan and Philip Pettit, *The Economy of Esteem*, OUP, 2004.
- Geoffrey Brennan and Philip Pettit, “Unveiling the Vote”, *British Journal of Political Science* 20/3 (1990), 311-33.
- *The Economist*, “Fat cats cornered”, 20th May 2003.
- Jon Elster (ed.), *Deliberative Democracy*, Cambridge, CUP, 2003.
- C. Alan Garner, “Equity in Economics Relationships”, *Journal of Economic Behaviour and Organization* 7 (1986), 253-64.
- Jürgen Habermas, *Between facts and norms: contributions to a discourse theory of law and democracy*, translated by William Rehg, Cambridge, Polity Press, 1996.
- John Stuart Mill, John Stuart Mill (1865), *Considerations on Representative Government*, Longman, Green, Longman, Roberts & Green.
- Liam Murphy, *Moral demands in nonideal theory*, Oxford, OUP, 2000.
- Thomas Nagel, *Equality and Partiality*, Oxford, OUP, 1991.
- Thomas Nagel, “Exposure and Concealment”, *Philosophy and Public Affairs* 27/1 (1998), 3-30.
- Daniel Naurin, “Does Publicity Purify Politics?”, *Journal of Information Ethics* 12/1, (2003), 21-33.
- Robert Nozick, *Anarchy, State, and Utopia*, New York, Basic Books, 1974.
- Plato, *The Republic*, translated with introduction and notes by Francis MacDonald Cornford, 1975.
- John Rawls, *A Theory of Justice*, 2nd edition, Oxford, OUP, 1999.
- Jeremy Waldron, “The Primacy of Justice”, *Legal Theory* 9 (2003), 269-94.