Basic Income and the Labour Market: Labour Supply, Precarious Work, and Technological Change

David A. Green *

October 19, 2023

Abstract

In this paper, I examine the relationship between a Basic Income (BI) as a policy tool and the functioning of the labour market. I focus on three key areas where a BI has been hypothesized to relate to labour markets: through altering work decisions; as a response to predicted changes in work arising from technological change; and as backstop that would allow workers to demand better working conditions and higher wages. I provide answers on the role or impact of a BI in each area in the context of the current Canadian labour market. But a key focus in the paper is on the ways we could alter our labour market models to provide a better basis for debating the impacts of policies like a BI in the context of a goal of moving toward a more just society.

*Vancouver School of Economics, UBC, and Institute for Fiscal Studies, david.green@ubc.ca. I am grateful to Steve Lehrer for a thorough set of comments on a first draft of this paper and to my students, past and present, whose work, insights, and friendship have shaped the way I look at economics.

1 Introduction

The evaluation of a Basic Income (BI) as a policy tool is intricately connected with its relationship to the functioning of the labour market. In this paper, I will examine some of the key connections between a BI and the labour market that have been raised as arguments either for or against a BI. My goal, though, is not to provide a definitive pronouncement on whether a BI is a good policy tool - though I will make some clear statements on some elements of its relationship to the labour market. Instead, I will use an investigation of claims about a BI as a sort of crow-bar to pry open our models of different aspects of labour markets. I will argue that our models and the empirical estimates that relate to them need to be enhanced to allow a more definitive evaluation of BI (and other polices that affect labour markets). This is particularly true when we view policies through what I view as the most appropriate lens - an evaluative lens centred on a goal of increasing the justice of our society. Of course, different people have different definitions of justice, and one of my goals is to set out ideas about how to decide on policies when people have different notions of justice.

I will focus my attention on three specific claims about the way a BI would interact with labour markets. One of the most visible connections (and perhaps a key reason that BI has risen to prominence in public discourse in recent years) is the claim, sometimes seen as emanating from Silicon Valley, that AI and robotics are going to either eliminate paid work as we know it or reduce all work to low wage, service jobs. Machines, it is argued, will take over most higher level functioning. They will write code on their own, directing other machines to manufacture, provide services, co-ordinate retail sales, etc.. In such a world, the owners of the machines will reap huge rewards and everyone else will be left with no work or bad jobs. Our redistribution system, predicated as much of it is on promoting or subsidizing work, will fail and we will necessarily have to shift to something like lump sum, unconditional (on work) transfers if we are to undo the vast inequalities that an unfettered market would generate.

Even if we don't believe this scenario in its worst form, providing an unconditional cash transfer for workers and non-workers has, BI proponents argue, the potential to raise wages and improve working conditions. The key idea is that people will have a secure backstop that will increase their bargaining power by allowing them to refuse jobs with bad working conditions. Firms will then need to respond with better working conditions and higher wages. Put in more specific terms, a BI would lead to better jobs.

The last connection between a BI and the labour market is, perhaps, the most obvious

and the one that is most often used by opponents of a Basic Income. It is the idea that more generous non-work tied benefits will lead to substantial withdrawals from paid work to the detriment of the functioning of the economy. In this outlook, though, all non-work is gathered together under the somewhat derogatory term of 'leisure', with its implication that not using our hours at paid work means not making a contribution. Instead, BI proponents would argue that caregiving (as an example) is also work - just work that often isn't rewarded with pay - and concerns about negative impacts on the economy reflect a mismeasurement of what we are truly producing and the way we truly want 'work' integrated with society. Moreover, they don't concede the point that a basic income would lead to a decrease in paid work because they see it as reducing the 'welfare wall' that consists of high effective tax rates on income when moving off income assistance benefits and with it, increasing incentives to work. In other words, a BI would cure the key work-related ills of the current welfare system.

Much of the basis of the discussion here will come from work done for the British Columbia Expert Panel on a Basic Income. I served on that panel along with the other panel members, J. Rhys Kesselman and Lindsay Tedds as well as Daniel Perrin (who was not an official panel member but contributed substantially to its conclusions in his role as an official consultant). The panel was created by the BC government in 2018 to investigate the question of whether BC should reform its transfer and support systems to centre them on a BI. The ultimate mandate of the panel, though, was wider than just examining a BI. We were also charged with examining the current transfer and support system in order to understand its shortcomings and benefits. In the end, we provided a report that recommended against centring BC's support system on a BI and, instead, provided a comprehensive set of reform ideas for the existing system. Of key interest for this paper, our work was based on over 40 research reports written for the panel by leading Canadian and international experts.[1] I will directly refer to papers by Chris and Craig Riddell, Kevin Milligan, Robin Boadway and Katherine Cuff but all of the research reports provide useful background for evaluating both BI and other policy tools regardless of a person's ultimate stance on the usefulness of those policies. We also benefited greatly from talking with vulnerable individuals who utilize BC's support systems and organizations that represent them.[2] I will use the outcomes of those discussions to guide the investigations in this paper.

In the remainder of the paper, I will first define a BI then move through a discussion of each of the labour market related issues I mentioned earlier. I will start with the most direct issue - paid employment labour supply - in section 3. I will then provide a somewhat

[1] All of the research papers along with the full report can be found at `https://bcbasicincomepanel.ca`.

[2] The work of the BC Poverty Reduction Committee preceded and overlapped with our panel and we benefited from access to their accounts of their extensive discussions with vulnerable people and groups.

brief discussion of current trends in the labour market and their implications for concerns about the impacts of new technologies. Both of those discussions point to the importance of evaluating the third labour related BI claim - that it can lead to an improvement in wages and the quality of jobs - and I focus on examining that claim in the fifth section. The sixth section contains conclusions.

2 Defining Basic Income

To begin, we need to establish a definition of basic income (BI). At first glance, this seems simple, but many different policy proposals have been attached to the 'basic income' label. Still, most of them share three main features: it is a cash transfer that is unconditional and universally available. The fact that it is a cash (as opposed to an in-kind) transfer is of central importance since BI proposals embody what have been described as either left or right-Libertarian sentiments. Giving people money, it is argued, allows them to make their own choices without intrusive over-sight. The unconditional feature tends to specifically refer to payments not being conditional on any work-related requirements (e.g., requirements to search for work, get training, etc.). The universal feature means that anyone can access the benefit when needed, i.e., not only if they fall into a specific group such as people living with a disability. It is worth noting that every BI proposal put forward for Canada violates either unconditionality or universality to some degree. They all have some proposed conditions for accessing the benefits (such as being a permanent resident) and/or do not apply universally. Many proposals specifically exclude people living with disabilities, for example, since their needs are more complex than could be met with a cash benefit of an amount that could be at a useful level for others.

BI proposals come in two main forms. The first is sometimes called a Universal Basic Income (UBI) and would involve sending a cheque of the same amount to every adult resident of Canada (and, possibly, a pro-rated cheque for each child). These schemes are, on the face of it, prohibitively expensive and virtually none of the BI proposals made in the Canadian context take this form. The other form is as an income-contingent cash transfer. In this form, everyone has access to a basic guaranteed income if they have no other income source and then that transfer is reduced according to a tax-back rate for each dollar of other income the person receives. The combination of the amount of the basic guarantee and the tax-back rate will imply a break-even income above which individuals will not receive any transfer.

A very wide set of claims have been made for the beneficial effects of a BI. Proponents argue that having a guaranteed backstop would reduce stress (with associated health bene-

fits), allow people to switch careers, allow people to invest in human capital, support people in starting a new business, create the conditions for better health and education outcomes for children, and many more. In our report for the BC government (**?**) and in an expanded version applied to conditions in all of Canada in our recent book (**?**), the BC Basic Income Panel members presented intensive investigations of these claims. Here, I will focus on just one set of claims - those related to labour market outcomes, diving deeper into the impact of a BI than was possible in the broader discussions in the report and book.

3 Basic Income and Labour Supply

To establish the competing arguments about how a BI would affect paid employment labour supply, it is helpful to set out a simple static labour supply model.[3] In such a model, individuals choose whether and how much to work based on a comparison of their preferences with their choice set of feasible combinations of hours of work and consumption of a single good. Tax and transfer systems affect the shape of the budget constraint and, through that, the incentives that individuals face in their work decisions. In figure 3, I set out the budget constraints under two systems: an income assistance (IA) system like those in place in most provinces; and a BI. In the figure, hours of 'leisure' (i.e., hours used in anything other than paid work) are read from left to right along the horizontal axis. The individual has a maximum of A hours to divide between work or leisure, so that hours of work are read from A (where all hours are used in leisure and so there are zero hours of work) to the left. The vertical axis shows income (or consumption given that the consumption good is the numeraire).

If there are no taxes and transfers, the person represented by the figure has no income when they don't work. For each hour that they work, they earn their exogenously set wage, with the result that their budget constraint is given by the solid line AF. Under the IA system, the person gets a transfer benefit given by AC. If they work and generate earnings, they are allowed to keep their earnings without affecting their IA benefit until their earnings hit the 'earnings disregard' ($500 per month in BC). This happens at point B, with the vertical difference between B and C being the value of the disregard. After that, the transfer is reduced by a dollar for each dollar of earnings until all of transfer is taxed away, which happens at the 'break even point', D. In comparison, under a BI system, the transfer benefit the person gets when not working is AC', which is larger than the benefit under the IA system to reflect the fact that this is the case with most Canadian BI proposals. Under the

[3]The exposition here is a shortened version of the discussion in Chapter 10 of **?**.

BI system, the transfer is reduced as the person earns money in the labour market but the effective tax rate is less than 100%. The combination of the BI transfer and the tax back rate result in a break even point given by D'. As drawn here, the BI system is more generous, providing a higher income when the person is not working and also while working, up to the D' break even point.

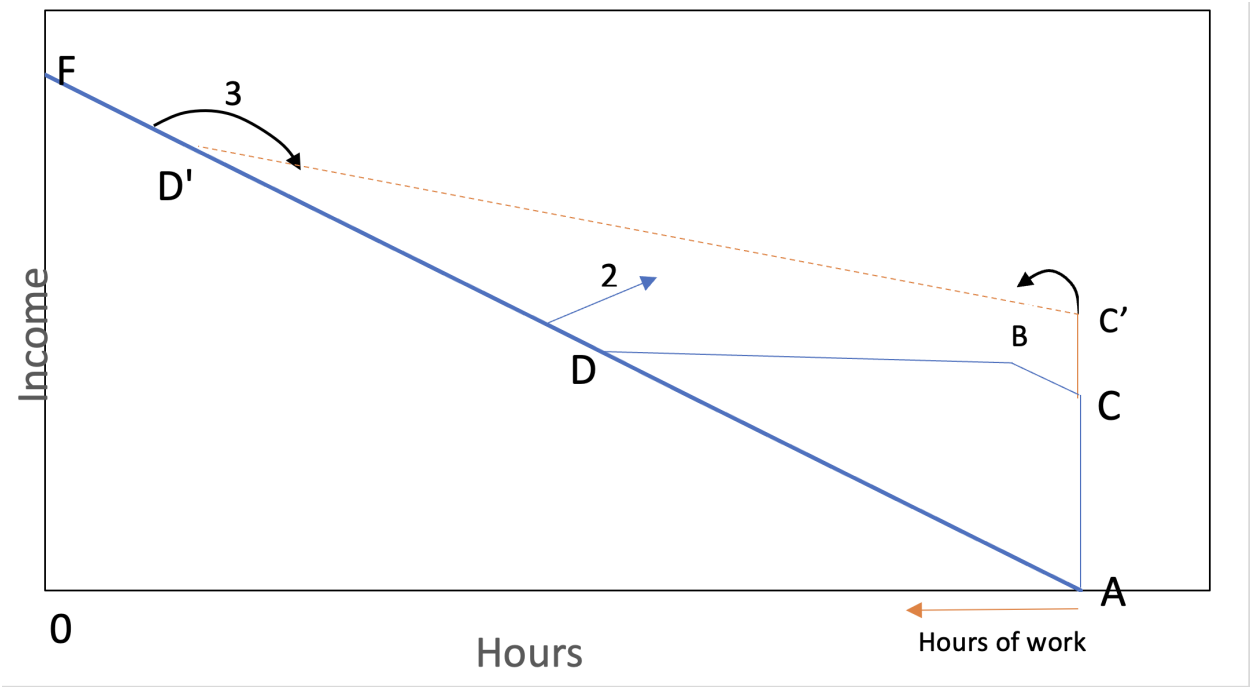

Figure 1: Work and Income with IA and BI

The work incentives generated by the two systems are represented by the arrows in the figure. Relative to the situation without any transfers, both systems generate an incentive to withdraw from work altogether (i.e., for the person to choose the point C or C', depending on the system they face). In addition, for a person working on the budget constraint to the left of D under the IA system, the more generous BI system incentivizes them to work less (arrow 2 in the figure). For such a person, both the income effect (arising because the BI system is more generous) and the substitution effect (arising because they face a tax back rate under the BI system but not under IA since they are past the break even point for IA) push them to work less. Finally, as arrow 3 indicates, some people who were formerly working beyond the BI break-even point may cut their hours of work in order to access some of the benefits. It is the sum of these forces that BI critics point to as the basis of their concern that a BI could lead to substantial declines in hours of work. On their side, BI proponents argue that the 100% tax back rate facing people looking to move into work under the IA system (often

called the 'welfare wall') is reduced to the lower tax back rate under the BI system and that will induce some people to move into paid work (represented by the arrow at C'). The figure indicates that both of these contrasting incentives exist. Which would dominate (and hence whether moving from IA to a BI would increase or decrease hours of paid work) will depend on how much the actual budget constraints under the two systems differ and on preferences for work versus consumption in the population. The latter are captured in labour supply elasticities, and I turn to discussing what we know about them in Canada, next.

3.1 Labour Supply Elasticity Estimates

Paid work responses to transfer system incentives can be captured in four elasticities: ϵ_{hw} (the elasticity of hours of work, h, conditional on working with respect to the wage, w, - the intensive margin elasticity); ϵ_{hy} (the elasticity of hours of work conditional on working with respect to non-labour income, y); ϵ_{Pw} (the elasticity of participation with respect to the wage - the extensive margin elasticity); and ϵ_{Py} (the elasticity of participation with respect to non-labour income). An assessment of the potential impact of shifting to a BI on hours of work requires estimates of all four of these elasticities.

The most comprehensive, recent elasticity estimates for Canada are found in **?**, using Survey of Labour and Income Dynamics data. They find quite small hours and participation elasticities for married women, fitting with more recent estimates from the US. Importantly, though, they provide estimates for different quantiles of the hours distribution showing that ϵ_{hw} is near zero for women with hours of work that place them at the top of the hours distribution and quite high (on the order of 0.2) for women near the bottom of the distribution.

Estimates related to a BI or transfer systems that have similarities to a BI are of particular interest for us. In a set of papers, Chris and Craig Riddell have re-examined the evidence from the MINCOME BI experiment run in Manitoba in the 1970s and key BI experiments run in Denver and Seattle in the same decade.(**?**,**?**) A key point in their analysis is that it is important to consider the BI in the context of the system it replaces. In particular, they argue that a substantial proportion (possibly two-thirds or higher) of single mothers in the MINCOME experiment were previously in receipt of income assistance. This is important because the increase in the income guarantee associated with the shift to the BI was not necessarily large (indeed, in the Denver experiment they examined, the guarantee actually declined relative to welfare benefits). That means we are considering a situation like figure 3 but with C' close to C, where what the shift to a BI really implies is a decline in the tax-back rate of benefits. This fits with claims by BI proponents that a BI could increase

labour supply at the extensive margin by reducing the welfare wall.(**?**) Indeed, in both the MINCOME and Denver experiments, they find that the BI treatment led to an increase in participation of single mothers who were either previously in receipt of welfare or were likely to have been in receipt. In the MINCOME case, they found a 14% increase in the probability of working and a 30% increase in hours worked. On the other hand, in the Denver experiment, the BI treatment led to a decrease in hours worked for single mothers who were previously working (fitting with arrow 3 in figure 1).(Riddell and Riddell(forthcoming)) Similarly, married women (who were much less likely to be in receipt of welfare benefits in the MINCOME experiment) decreased their hours of work in the presence of the BI. [4] This provides a more nuanced take on the overall conclusion from **?** that MINCOME had small effects on hours worked.

For more recent results, we can turn to papers examining the introduction of various child-linked benefits. The Universal Child Care Benefit (UCCB), introduced in 2006 then made substantially more generous in 2015, and the Canada Child Benefit (CCB), introduced in 2016, are two key programmes which are, in essence, basic incomes but restricted to families with children. **?** examines the impact of the introduction of the UCCB in 2006 on the labour supply of married women with children, finding sizeable negative effects. In contrast, **?** find no effect of the introduction of either the UCCB or the CCB on the labour supply of either married or single mothers at either the intensive or extensive margin, and Koebel and Schirle(2016) find positive effects for divorced women but no effects for never-married lone mothers.

These estimates, of course, focus on women with children. But a key point of concern in the current policy environment is single adults without children, who are the group with the highest poverty rates. Evidence that is potentially relevant for evaluating the impact of a BI for this group can be found in **?**, who examine the elimination of a previous age-related benefit restriction in the social assistance system in Quebec in 2008. In the prior SA system, young people under age 30 were given a reduced benefit, with the benefit jumping by 175% to the full benefit on their 30th birthday. Lemieux and Milligan employ an RD design to examine this effect, finding sizeable declines in employment at age 30. Their estimates imply an elasticity of participation with respect to income of between -.026 and -.13, a large range that runs from what are typically considered small at the bottom end (effects that don't imply a substantial reduction in total work hours) to large at the top end.

[4]The amount by which they did this was only about twice the very small hours effect for men but because married women worked few hours on average at the time, this amounted to a 40% decline in average hours worked).(**?**).

3.2 Incentive Effects from Introducing a Basic Income

One of the conclusions from the empirical literature on the impact of introducing a BI is that estimated effects will naturally depend on how much of a change the specific BI variant being considered represents relative to existing systems. To make this point in a context that is relevant for the Canadian case, I consider the effects of introducing a BI with a $20,000 guaranteed income combined with a 50% tax back rate. This is a scheme that has been suggested by several proponents since the income guarantee is approximately the poverty line for a single adult and, so, this BI could constitute a substantial move toward eliminating poverty. I compare this to the set of provincial and federal transfer systems (including tax credits) currently in place in BC. In making the comparison, I assume that only IA is eliminated in the move to a BI - all other transfers and taxes remain in place. This is based on the exercise in **?** in which they examined each tax and transfer system in place in BC in 2019, arguing that virtually none of them could be effectively replaced with a BI apart from IA. That is, there are no substantial tax or programme savings available from introducing a BI.

To summarize the incentive effects of shifting from IA to BI, I use the participation tax rate (PTR) - the average tax a person would pay if they moved from not working to working at a job that paid earnings of y. Formally, the PTR is defined as:

$$PTR(y) = \frac{f(y) - f(0)}{f(y)} \tag{1}$$

where, $f(y)$ is disposable income and reflects all the taxes paid and transfers received for a person with earnings of y (and can be larger or smaller than y). I plot the PTR under the BC IA system and the alternative BI system for a single adult without children in figure 2. The solid line corresponds to the current system and is taken from **?**. It is zero for the first $6,000, corresponding to the annualized value of the monthly earnings disregard under the BC IA system. It then moves rapidly to nearly 70% before declining gradually. Note that the marginal tax rate goes to 100% after the earnings disregard, as in figure 1, but the fact that the PTR is an average up to y and there are zeros from the earnings disregard being averaged in means that the PTR stays below 100%. The rapid rise in the PTR to very high tax levels in figure 2 is the embodiment of the welfare wall. The PTR for the BI on its own is given by the lighter line that is horizontal at 50% up to $40,000 (the break even income under this scheme) and then declines gradually - again, reflecting the fact that the PTR is an average. The dashed line corresponds to removing IA from the current system and replacing it with the BI scheme.

A key point from comparing the BI complete system line to the line representing the PTR for the current system is that one is not clearly better than the other. The BI system has a much higher PTR over the first $6,000 (50% instead of 0%) then is lower up to almost $20,000 before being substantially higher beyond that point. One might argue that a BI could be created that also had an earnings disregard, negating the advantage of the current system at low earnings levels, but this would introduce a work conditioning element that goes against the very nature of a BI. Regardless of this point, it is difficult to make a case that a shift to a BI would represent a substantial reduction in the welfare wall and every reason to believe it would generate incentives to reduce labour supply over an extended range of earnings. And this is even without bringing in the higher taxes that would have to be assessed to pay for the more generous BI scheme.

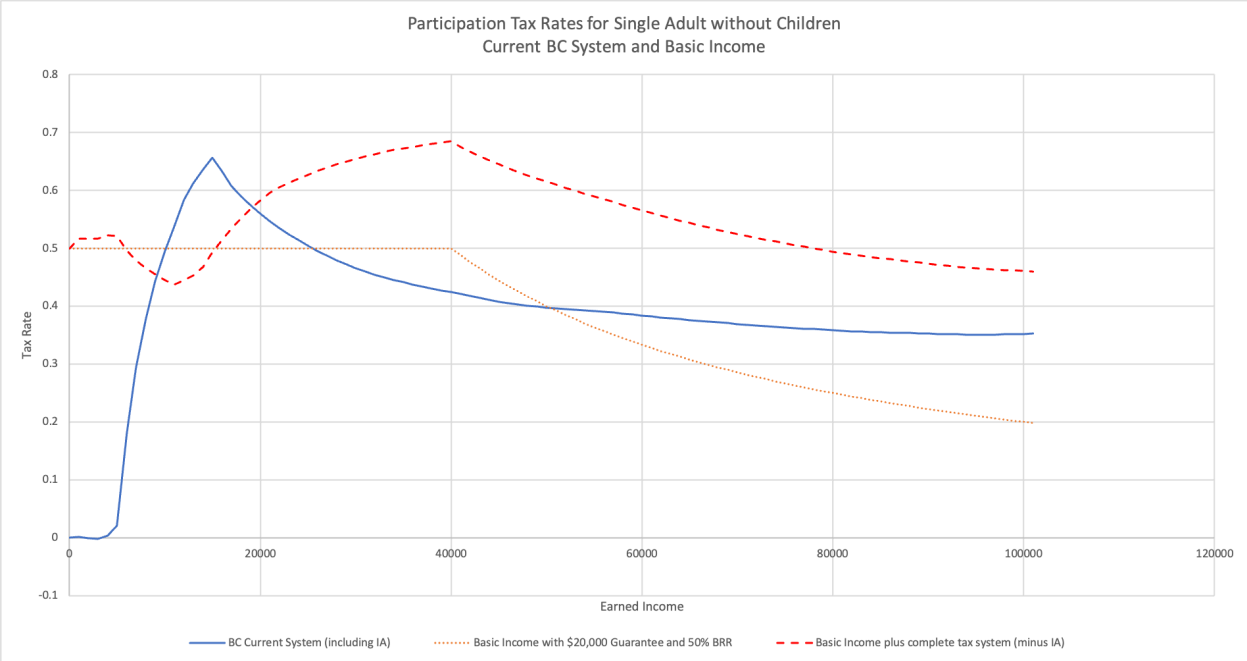

Figure 2: Participation Tax Rates for Single Adult without Children

3.3 A Prediction of Hours Effects of Shifting to a BI

In **?**, we use the range of elasticity estimates for Canada combined with effects of shifting to a BI on marginal and participation tax rates to provide a range of estimates of net effects for paid hours of work. As in figure 2, we use a BI system with a $20,000 income guarantee and 50% tax back rate, removing only the IA system when we introduce a BI. We generate the different implied marginal and participation tax rate effects for two main types of people.

The first consists of people in receipt of IA benefits in the current system (in order to assess claims about welfare wall reduction effects for this group). The second is all other people, breaking them down into single adults without children, lone parents, couples with children, and couples without children (to account for differences in how existing systems such as the CCB affect these groups). We use both BC IA caseload data and the 2016 Census to get the number of people in each group and their initial hours of work distributions. For each person in these data, we calculate counterfactual hours of work based on the combination of our computed tax rate changes and the elasticity estimates that are relevant for the specific group. We focus on 19 to 64 year olds under the argument that BI type programmes already exist for children and seniors. We assume, based on the arguments in **?** that hours effects are zero for individuals earning over $40,000 a year.

The first key conclusion from this exercise is that a shift to this BI would have minimal effects on the hours of work or participation decisions of people on IA. In part, this follows from figure 2, which shows that the effective tax rates facing this group do not change dramatically. Overall, we predict that introducing a BI would imply a decline in total annual hours worked in the province of between 0.6% and 1.7% (depending on the specific elasticity estimates used). This is non-trivial but also doesn't represent a hugely destructive effect.

It is important to highlight why our predicted hours effects are so small. BI proponents sometimes argue that labour supply effects will be small because people simply want to work.(**?**) In our terms, this is an argument that labour supply elasticities are all small. To the extent this is true, we could expand the generosity of a BI scheme substantially without fear of substantial labour supply effects. That, though, is not the source of the small effects here. Some estimated elasticities are, indeed, small, but recall that **?** find large elasticities at the bottom of the hours distribution and we use their estimates here. Our small estimates stem from two sources. The first is that there are offsetting responses for different groups - increased hours for some currently on IA and decreased hours for working mothers. The second is that the incentives from moving to a BI are not dramatically different from those in the current system, with better paid work incentives in some ranges and worse in others. This echoes the argument in **?** and **?** that a BI integrated (as it must be in the real world) into the existing system of other taxes and benefits is not necessarily a simpler approach with lower tax walls. It is a complex system in its own right that is not obviously better or worse than existing systems in terms of labour supply incentives.

Taken together, these factors imply that there could be versions of a BI that could have large, negative effects on paid labour supply. The main versions being proposed in the

Canadian context, though, do not seem to suffer from that problem. Finally, it is worth repeating the obvious point that even for those who appear to react to a BI by cutting their work, there is reason to believe this may not be a bad outcome, since the main group for whom this is true is mothers with young children.(**?**)

3.4 Optimal Tax Systems and Labour Supply

The fact that a BI is more accurately viewed as a variant on transfer systems that we already know than something radically different can be seen by thinking of the problem of systems design through the lens of the optimal tax literature. In that literature, the goal is to design tax and transfer systems that maximize a social welfare function or pursue some goal like minimizing the poverty rate while taking account of the combination of incentivized labour supply responses and incomplete information on the part of the government. The classic reference, of course, is **?**, who examines this question in a static context in which the government wants to redistribute toward the lowest ability workers but cannot directly observe ability. The solution is couched as a balance between the equity goal of helping the least able and the efficiency goal of restricting losses of production that would arise if other workers were to reduce their work hours in order to get the benefits. The specific form of the solution is one in which a benefit is paid to those who don't work combined with a tax schedule with very high tax rates at lower levels of income, lower rates in the middle of the income distribution and (reflecting redistributive goals), high tax rates at the top end of the distribution. The high tax rates at the bottom are required to focus benefits on the least able. With them, the break even point of the benefit system is lower so that fewer workers face direct incentives to work less. Importantly, there is no concern with the least able not working. Not having them work has very little efficiency cost and delivering benefits to those who do not work is a way of targeting support to them.

? present a thorough and insightful examination of a BI through the lens of the optimal income tax literature in a contribution to the BC BI Panel research series. In that paper and **?**, the authors point out that the focus on optimizing social welfare functions (SWFs) in the optimal tax literature is potentially restrictive. I want to consider, instead, the broader goal of moving Canada toward being a more just society. In the context of that question, an approach based on SWFs is restrictive in the sense that issues related to ideals of justice can only enter to the extent that some set of individuals see them as affecting their utility. The philosopher John Rawls criticizes these Welfarist approaches on the basis that there is nothing in them preventing the individual utilities that are their fundamental building blocks from being abhorent and, as a result, generating abhorent implications.**?**

The simplest example is a situation with a standard Utilitarian SWF and a majority of people in the society favouring enslaving the minority. But this is just one example of issues with Welfarist approaches. Many philosophers have weighed in on the problems with these approaches.[5]

Of course, if we want to focus on a broad goal of moving toward a more just society, we need to specify what notion of justice we are considering since there are many and varied options.[6] I would argue that virtually all theories of justice place considerable emphasis on providing everyone with the bases of self and social respect. In John Rawls' famous theory of justice, for example, the individuals behind the veil of ignorance who are debating over the principles that should guide the key institutions of society focus on impacts on what Rawls calls the 'social primary goods' - the things that people are believed to value regardless of whatever else they value. Key among these are the bases of self respect.(**?**) Martha Nussbaum and Amartya Sen similarly put substantial weight on the bases of self respect in their Capabilities theory of justice.(**?**, **?**) And the list goes on. Rawls defines self respect as a 'confident sense of (our) own worth,' arguing that such a sense will allow people to pursue their goals and stand in a position of equality with others. Importantly, our self-respect necessarily has a social element to it since we look to others to understand even how to value ourselves. Thus, I will talk about the combination of self and social respect.

In **?**, we develop the practical implications of this idea, arguing for a specific set of principles to use in considering policy options. Key among these are adequacy of income levels, balancing individual autonomy with supporting communities (since we find both our self and social respect in communities); and focussing on the process of implementation of policies as well as their end goals (since the respect that people experience when accessing supports are a key feature of how policies affect their lives.) Considering those principles leads in a different direction from optimal tax theory. As part of the BC Basic Income Panel, we had the opportunity to talk with people who were in need of government support. We also benefited greatly from the work of the BC Poverty Reduction Committee that preceded us and who created a record of their considerable consultations. What emerged from those discussions was a picture of a Kafkaesque system. The system is constructed along the lines of optimal tax theory with high tax rates at low income. But in the optimal theory, those high tax rates aren't really relevant to the benefit recipients because they aren't expected

[5]Apart from John Rawls, others include Ronald Dworkin (**?**), Martha Nussbaum (**?**), and Amartya Sen (**?**).

[6]I have argued elsewhere that the apparent promise of the Welfare Theorems that we can separate efficiency from justice and make policy prescriptions without reference to any theory of justice is false. To the extent that is the case, deciding on a notion of justice to use in our deliberations is unavoidable even if we wanted to focus our attention on issues that are seemingly more about efficiency.**?**

to work. In the actual implementation, these people are pressured to search for work but face very high tax rates if they find it. Moreover, the jobs they are likely to find tend to be short and temporary but the system - with its emphasis on work - is constructed to make it difficult to get back on benefits. That creates a perverse incentive not to take a job even while being shown clearly that respect is to be found moving off of benefits and into work. This is particularly difficult since the benefit recipients and others in need of support who were part of these conversations often, themselves, expressed a desire for work, likely as an element of their own search for self and social respect.

I see all of this as a problem with our approach toward thinking of work in economic models. In our main models, people value leisure and would always prefer a life of living on benefits without work (if their earnings options place them near the benefit level of income). Supporting such a life for the least productive people in society is considered a policy success in a system in which walls are built to prevent others from joining them. My sense from the discussions with people who might actually need the benefits is that a more accurate model would be one in which people have a basic desire for working (for reasons of self respect, feelings of self-efficacy, and social connection) but would also, in any moment, prefer more leisure. This fits with Akerloff and Kranton's ideas about identities and their relationship to working.(**?**) At the same time, many of the people who are in need of support face multiple barriers to work (health issues, poor work records, insufficient housing, etc.) that mean they face low paying, short duration jobs. That is, we shouldn't think in terms of models in which they can be moved into permanent working states but ones in which they will repeatedly find themselves in need of support mixed with jobs they take in their search for self-respect and connection. The obvious trade-off is that we would want to create a system that provides support in this sporadic work pattern without creating incentives to take up sporadic work patterns for people who might not otherwise face them. I don't know of any papers that take this perspective when thinking about designing transfer policies. The closest is **?** who develop an optimal tax model in which people differ in their abilities in the market and non-market sector. They argue for an optimal system that includes wage subsidies at low wages but on a limited basis. But part year work patterns in their model are a choice rather than a feature of the labour market options facing individuals who need support, and there is no element related to a desire to work.

Returning to the main theme of this paper, a perspective that places more emphasis on supporting self and social respect and their relationship to work has implications for how we think about a BI. On one hand, a BI does well under this perspective. At the very least, it implies providing benefits without the judgement associated with work requirements. In that sense, it is closer to the recommendations of optimal tax theory. On the other hand,

it doesn't pay enough direct attention to the relationship between respect and work. People could (and might) use a BI as monetary support for their desire to find work and get training. But this is left entirely up to them. Surely, a system that provides direct supports would be more effective. A specific example is found in the Community Employment Innovation Project (CEIP), a federal government experiment that involved enlisting a set of six Cape Breton communities to design community projects. IA and EI recipients were then given the opportunity to shift off their regular benefits and, instead, take up jobs with a regular wage carrying out the community determined projects. The idea of the programme was to both develop community capacity and networks and to provide useful work experience (and networking) for those in receipt of benefits. This is an approach that works directly on the goals of community building and provides people on benefits with the opportunity to find a sense of self-efficacy and contribution while receiving support.[7] As a general statement, the conclusion that a BI could provide a backstop for various kinds of changes but that there are more direct ways to generate those changes was something that we, as a panel, found in multiple policy areas.

4 Labour Markets and Technological Change

One of the arguments that has placed a BI at the centre of policy debate is the claim that we are facing a future with either no work or with the majority of people being in precarious work because of emerging effects of AI and robotics. Indeed, it is common to hear the argument that work is already becoming increasingly precarious (where precarious work is defined as short term, contract work, part time work, or own-account self-employment). Based on that, proponents argue that a BI will be necessary since it will become increasingly difficult to distribute what society produces through either wages or policies related to them such as wage subsidies or employment insurance. For example, Philippe Van Parijs and Yves Vanderborght, who are strong proponents of a BI, state,'the expectation that meaningful work will be lacking easily leads to the conviction that the growing jobless population must be provided with some means of livelihood' (**?**).

In Chapter 15 of **?**, we evaluate the claim that the end of work - or, at least, the end of good jobs - is nigh. Anyone paying even cursory attention to labour force statistics will know that rumours of the death of work are greatly exaggerated. The current employment rate for people age 25 to 54 is higher than at any time over the last 30 years. In addition,

[7] **?** examine the CEIP, showing that it did not lead to an ongoing increase in employment after the end of the programme but it had persistent effects on social networks up to 2 years afterward.

claims that precarious work is taking over our labour market are not easily supported. The proportion of jobs that are the opposite of precarious - the proportion of jobs that are full time, permanent (as opposed to contract), paid (as opposed to self-employed) is above where it was at in 1989 for men and approaching the level for that year for women. This is not at all to say that precarity is not a problem. Based on the LFS data, nearly a third of workers are in precarious employment, with higher proportions reported by Iglika Ivanova and Kendra Strauss in their study of precarious work in BC based on a survey targeted at understanding such work(**?**). Moreover, the level of such work appears to have stayed at similar levels for decades, prompting Andrea Noack and Leah Vosko to say that the labour market is in a state of "persistent precarity"(**?**). Moreover, there are some troubling trends with respect to worker vulnerability and having a say about their work environment, including the decline in unionisation rates (especially in the private sector) and the apparent rise of 'fissured' work. Fissured work refers to situations in which the worker is not directly employed by the firm where they physically work. A standard example is janitors who work for a janitorial services firm but do their work at a law office. There is evidence that these work arrangements cause declines in wages (**?**) and it seems likely they reduce the application of labour regulations(**?**).

The description of our situation as persistent precarity instead of increasing precarity is important. If we face increasing precarity driven by technological forces beyond our control then we may reach a point where wage based distribution of the product of society is not useful, pointing to the need for a BI. In that case, the relevant policy question becomes, should we introduce a BI now in order to prepare for a poor job future or should we wait until that future arrives. With persistent precarity, in contrast, there is reason to look backward in order to ask when this level of precarity started (or whether it has always been with us). The trends in Canadian data point back to the deep recession of the early 1990s. It was at that point that measures of precarity rose and measures such as unionisation were in decline. Measures of precarious employment stayed at heightened levels from that point until 2019. In COVID, those measures fell, as workers in precarious work arrangements were less likely to be able to work from home. We might have expected a return to the pre-COVID levels of precarity with the full return to employment but, instead, the tight labour markets of the last few years have allowed measures of precarity to stay at their lower levels and, in fact, are at levels that are on par with the late 1980s. From a policy perspective, the questions that arise are then, how do we lower precarity more and how do we support the types of tight labour markets that appear to be to the advantage of workers looking for better jobs? Or, more directly for the purposes of this paper, is a BI an effective tool for either lowering precarity or tightening labour markets? In the next section, I will examine the related claims that a BI allows workers to walk away from bad jobs, forcing firms to offer

jobs that are better both in terms of wages and working conditions.

But could we still be headed for a future with bad jobs and low wages even if current trends aren't pointing in that direction? After all, what is being considered is an historic technological change that could have effects that are in the nature of suddenly moving over a threshold. To assess that possibility, current trends aren't necessarily relevant and, so, we need to turn to theory. Indeed, there has been considerable debate about what theory tells us about the long term impact of AI and robotics.(e.g., **?**, **?**). The arguments in that debate hinge on factors such as whether the production function is constant returns to scale, whether we should characterize technological change as having a total factor productivity form or a factor augmenting form; the relative elasticities of supply of labour and capital; and the role of technology in changing market concentration.

I find convincing the arguments in **?** that labour is ultimately the more fixed factor and, so, will not - as a whole - be the loser in technological development. Though, of course, technological change could make the distribution of wages and work even more unequal. But I want to focus attention on a different part of the literature that examines the impacts of technological change. That is a stream of papers that argues that the technological change that actually affects production (as opposed to remaining as potential ideas on the shelf) is endogenous. It responds, in particular, to incentives that arise out of relative factor supplies and other forces that affect factor costs. The specific mechanisms tend to take two forms: induced changes in invention (**?**) or induced adoption of previously invented technologies (**?**, **?**). In either case, increases in the proportion of the labour force with higher education can have no effect on educational wage differentials or even imply increases in them - the opposite of the prediction in a simpler neoclassical model. And the evidence across a range of economies fits better with these models than simple models in which technological change arises as an exogenous force.

This conclusion is important for our purposes because it shifts the question from whether a BI is a necessary response to an inevitable process to whether a BI is a helpful tool in directing technological change along better paths. **?** argue that factor augmenting technological change can take either labour-substituting forms (automation) or labour-complementing forms (as AI is doing in some realms) and that we should design policies that favour the latter forms. What role a BI would play depends on whether it would raise wages and affect the nature of work. If it does then it would increase the cost of labour and push toward the labour-substitution form of technological change. Thus, this line of reasoning, too, points toward the need for an evaluation of claims by proponents that a BI would allow workers to hold out for better wages and working conditions.

5 Transforming the Labour Market

To this point, we have considered whether a BI would affect work decisions in the current labour market and how it would interact with technological changes that will affect the labour market. In both cases, questions arose that relate to the third claim about a BI and labour markets: that a BI would raise wages and alter the very nature of jobs. A key claim in favour of a BI is that it allows people to walk away from bad jobs and, because it does, it serves to discipline the market and reduce the extent of bad jobs. Evelyn Forget, a leading Canadian BI proponent, for example, argues that 'Some people, particularly those in low-paid and unpleasant jobs, may choose to work less if they have access to a basic income. If the work they abandon is necessary work that must be done, their refusal to work will encourage employers to improve the terms of employment.'(**?**, p. 90) Similarly, Annie Lowrey, author of *Give People Money* argues that a Universal Basic Income would give employees enough security to have bargaining power, arguing, 'Why take a crummy job for 7.25 an hour when you have a guaranteed 1,000 dollars a month to fall back on?'

It seems to me that it is natural to think of these claims in the framework of a search and bargaining model. In those models, workers have some difficulty finding job vacancies and firms posting vacancies have some difficulty finding matching workers. As a result, when they do meet, if the value of what they will produce together is greater than the sum of their next best options (their 'outside' options) then they form a match and bargain over a wage that serves to divide up the surplus of the match (its value over and above the sum of their outside options).

Shifting from IA to a BI in this context has several implications. First, most BI proposals involve an increase in the benefit received when not working, which I will denote as b. Second, IA systems in Canada tend to have very high tax back rates (100% after an earnings disregard that tends to be small). In contrast, BI systems have lower tax back rates (50% is a not-uncommon value). Though, as we have seen, the overall PTR structure in the two systems is not dramatically different. Third, IA systems have conditionality in accessing benefits that makes entry into the programme uncertain. In that sense, b should be multiplied by some probability less than 1 when thinking about its value as an outside option to a worker. In comparison, getting the BI is guaranteed. The question is, does shifting from an IA to a BI system lead to a disproportionate closure of 'bad' jobs - jobs with distasteful or otherwise problematic characteristics - and does it lead to increased wages? In this section, I will investigate these questions with standard models. Part of my goal is to use the interesting questions being raised by BI proponents as a crowbar to pry up parts of standard models in order to talk about where they could be improved in order to form a more complete

conclusion.

I will start with a standard Diamond-Mortensen-Pissarides (DMP) model, with some adjustments. In particular, I will assume that there are a fixed set of job 'types' (indexed by j) defined by two dimensions: productivity (ϵ_j) and the cost of making the job 'clean' (c_j). Thus, every job produces an amount of the output good (given by ϵ_j) when it is filled by one worker and that job can be done in a way that is 'dirty' (i.e., is in some way distasteful to the worker) or 'clean'. To start, I will assume there is a fixed number of firms that can open a vacancy of each type and define η_j as the proportion of firms of type j. Firms decide whether to open a vacancy and whether to operate as a clean or a dirty firm.

On the other side, workers are homogeneous in terms of their skill (so any worker will produce ϵ_j if they are assigned to a job of type j) and in terms of their distaste for dirty jobs. Later, I will discuss the effect of allowing workers to have heterogeneous tastes over job cleanliness, but standard discussions of BI effects imply that there is a commonly accepted notion that some jobs are distasteful. If that were not the case then we could not talk meaningfully about a BI improving the set of jobs in operation since there would be differences in opinion on whether any jobs that are eliminated are good or bad. So, our model needs to include some common assessment of job quality and it is simpler to assume that everyone shares the same assessment. In particular, workers have a distaste for dirty jobs that can be captured in an equivalent monetary value, τ.

As is standard in DMP models, workers and firms meet with probabilities less than 1 given by a matching technology that is expressed as a constant returns to scale function of the number of unemployed workers searching for jobs (U) and the number of empty vacancies (V). Workers can only meet vacancies when they are unemployed (a useful, extreme assumption in this case because we want to focus on the value of a BI in allowing them to prolong their job search while unemployed, turning down less desirable jobs while they do so) and workers and firms meet randomly. That is, workers can't find out which are the better firms and line up for jobs at them. I will return to this assumption later.

When a worker and a firm meet, they bargain a wage to divide up the surplus value of the match. Note that with homogeneous workers, the firm knows whether posting either a clean or a dirty job of its productivity type will provide a positive surplus (i.e., will be a job that both it and the worker will want to proceed with). It won't post vacancies that do not have a positive surplus and, so, any meeting of a worker with a vacancy will proceed to a match that will operate.

From the worker's side, the per-period value of a match with a firm of type j which is offering a clean job is $w_{jc} - T(w_{jc})$ - the wage that will be bargained when the job is clean

minus the tax liability at that wage. $T(w_{jc})$ can take negative or positive values and differs in form for IA versus BI systems. In an IA system, for wages below b (i.e., below the break-even point), $T(w_{jc}) = -b + w_{jc}$, i.e., the value to the worker when working remains at b regardless of the wage up to that point. If there is an additional utility benefit to not working (d) then no worker will want a job of this type. This implies a cut-off value $\epsilon^*(c)$ that is the productivity at which the bargained wage which would just equal b. For $\epsilon < \epsilon^*(c)$ firms won't open vacancies since they know workers won't accept the associated wages. For a BI system, $T(w_{jc}) = -b + \theta w_{jc}$, where, θ is the tax back rate on the guaranteed benefit amount. Here, too, there will be a break-even level of ϵ, which I will call ϵ^{**}, which corresponds to the wage $w^{**} = \frac{b}{\theta}$. In this case, though, firms with productivity below ϵ^{**} can still operate.

If, instead, the firm is offering a dirty job then the per-period value to a worker is $w_{jd} - \tau - T(w_{jd})$, where w_{jd} is the bargained wage if the job is dirty. Here, too, outcomes will depend on whether ϵ takes values above or below the break-even levels in the two systems. Since $T(w)$ is a function only of w (the government cannot see and does not base taxes and transfers on τ), dirty jobs are more likely to pay a wage that is above break-even levels than clean jobs.

To understand the impact of switching from an IA system to a BI system, I will consider clean jobs first. The key elements of wage setting are the value functions for workers when unemployed and employed and for firms when they have an unfilled and a filled vacancy. The value function for unemployed workers is given by:

$$\rho U^u = b + d + \psi[\sum_{j \in C} \eta_j U^e_{jc} + \sum_{j \in D} \eta_j U^e_{jd}] \tag{2}$$

where, ρ is the discount rate common to all agents, b is the value of the benefit paid when not working, d is the value of home production, U^e_{jc} is the value to the worker of being employed at firm j if it operates a clean job, U^e_{jd} is the value of working at firm j if it operates a dirty job, ψ is the probability the worker meets any job, and η_j is the proportion of jobs (and vacancies) that are of type j. Because search is random, the likelihood of meeting any particular type of job opening is given by the proportion of jobs of that type.

For employed workers, the value depends on the firm they meet, which determines whether the offered job is clean or dirty and the wage it pays. So, the value of a clean job is determined by,

$$\rho U^e_{jc} = w_{jc} + \delta(U^u - U^e_{jc}) \tag{3}$$

and the value of a dirty job is determined by,

$$\rho U^e_{jd} = w_{jd} - \tau + \delta(U^u - U^e_{jd}) \tag{4}$$

where, δ is the exogenous probability that a job ends.

Firms decide whether to post clean or dirty vacancies and continue to operate the job in the selected mode after filling the vacancy. The value of a filled, clean vacancy is determined by,

$$\rho V^f_{cj} = R_c \epsilon_j - w_{jc} - c_j + \delta(V^v_{cj} - V^f_{cj}) \tag{5}$$

where, R_c is a parameter determining whether running a clean job increases worker productivity ($R_c > 1$, as might happen in an efficiency wage type story in which workers are willing to provide more effort when the work place is agreeable); decreases productivity ($R_c < 1$, as would happen if costs of maintaining a clean workplace have both a direct component, c_j, and a component that is proportional to productivity); or is neutral with respect to productivity($R_c = 1$). For now, I will set $R_c = 1$. The bargained wage if the firm operates a clean shop is w_{jc}, and V^v_{cj} is the value of an unfilled clean job vacancy. Similarly, the value of a filled, dirty vacancy is determined by (noting that there are no clean-up costs or potential productivity advantages to a dirty job),

$$\rho V^f_{dj} = \epsilon_j - w_{jd} + \delta(V^v_{dj} - V^f_{dj}) \tag{6}$$

The value of opening a clean vacancy is given by,

$$\rho V^v_{cj} = -h + \phi(V^f_{cj} - V^v_{cj}) \tag{7}$$

where h is the flow cost of an open vacancy (which we will normalize to 0 in order to simplify the exposition) and ϕ is the probability an open vacancy meets an unemployed worker. V^v_{dj} is defined analogously.

Firms open a vacancy as a clean job if $V^v_{cj} > V^v_{dj}$. Firms of each type also need to decide whether to open a vacancy at all, which is determined by the values of vacancies being non-negative. As we will see, firm types with low productivity and high costs of clean up will tend not to open vacancies. We can make the distribution of firms across types, j, endogenous, allowing firms to select their type subject to a cost of beginning operations,

$$k_j = \frac{n_j}{\Omega_j} \tag{8}$$

where n_j is the number of workers in jobs of type, j, and Ω_j is a type specific cost advantage. Thus, the costs of opening will rise with the number of firms already opening jobs of a specific type, preventing all firms from entering just one type. The ultimate distribution of firms across types will depend on the distribution of Ω_j. Our main conclusions don't vary with whether the distribution of firms by type is exogenous or endogenous so I will maintain the assumption of exogeneity to simplify the exposition.

When workers and firms meet, there is a match specific surplus (which must exist for firms to be willing to open a vacancy that is not instantly filled) and they bargain a wage to divide that surplus according to Nash bargaining. Note that the surplus is diminished by the cost of keeping the job clean, c_j, in the case of a clean job and by the utility cost, τ, to the worker if the job is operated dirty. That means that the bargained wage in each mode will reflect those parameters. The bargained wage in a clean job is given by:

$$w_{jc} = \gamma_1(T(w) + b) + \gamma_2(R_c\epsilon_j - c_j) + \gamma_3 \sum_{k \in C} \eta_k w_{jc} + \gamma_4 \sum_{l \in D} \eta_l (w_{ld} - \tau) \tag{9}$$

The γ's are functions of the key parameters in the model. In particular, they are functions of the probabilities that workers meet vacancies (ψ) and vacancies meet workers (ϕ) as well as the marginal tax effective rates $\frac{\partial T(w)}{\partial w}$. Thus, the wage represents a mid-point between the worker's outside option (which includes both b and the fact that, when unemployed, a worker would access job options according to the distribution of job types) and the productivity of the match, $(R_c\epsilon_j - c_j)$, which reflects the cost of keeping the job clean. Following Hungerbuhler et al(2006)'s analysis of optimal taxation in a similar search setting without clean and dirty jobs, the bargained wage is increasing in total taxes $T(w)$ and decreasing in the marginal effective tax rate - something that is reflected in the γ parameters.

The bargained wage equation if the job is dirty is given by:

$$w_{jd} = \gamma_1(T(w) + b) + \gamma_2\epsilon_j + \gamma_3 \sum_{k \in C} \eta_k w_{jc} + \gamma_4 \sum_{l \in D} \eta_l (w_{ld} - \tau) + \gamma_5\tau \tag{10}$$

Thus, it is similar to the wage on a clean job except that there is no clean up cost to subtract from the surplus but there is a utility cost, τ, that must be covered.

The effect of the BI system on wages reflects offsetting forces. On one side, the high b value implies wage increases. But for job types with wages that would be below the break-even point in a BI system, T(w) would be negative in the BI system and that would imply lower wages. In essence, the workers' surplus from working is bigger under a BI for these jobs because the person gets an income greater than just the wage (since they keep some of

b when working). In Nash bargaining, that higher surplus for workers must be balanced by a higher surplus for firms and that is accomplished through a lower wage. Put another way, firms are able to capture some of the BI benefits that workers get while working.

We can use this model, in part, to consider difference in wage outcomes under a BI versus an IA system. The effect of the tax and transfer system on wages is given by $T(w) + b$. In the BI system, $T(w) = -b + \theta w$, so $(T(w) + b) = \theta w$. This is necessarily less than b but greater than zero for wages below the break even level. Whether this represents an increase or decrease relative to an IA system depends on the relative sizes of the b in each system and the size of θ. It will correspond to a negative effect relative to the IA system when wages are low (so that θw is small) and then rise with w up to the BI system b value. Since the BI system b value is larger than the IA system value, there will be a crossing point at which the wage impact will turn positive, with that crossing point depending on the values of b in the two systems and the value of θ. Thus, while going from no transfer system to a BI would increase wages, a shift from an IA system to a BI system would have negative effects on the lowest wages offset by positive effects at somewhat higher wages. This stands in contrast to blanket statements by BI proponents that a shift to a BI system would allow all workers to bargain higher wages.

What does empirical work tell us about the likely size and sign of these bargaining effects? In the derived wage equations, the effect depends both on the size of the changes in the taxes and transfers and on the γ_1 parameter that translates those changes into wage changes. Beaudry et al(2012) examine the impact of changes in outside options related to industrial composition and find that the effects are positive and sizeable. In contrast, **?** find no effect of UI benefit changes on wage setting in Austria. It is possible that outside options have asymmetric effects: increased access to higher options can be used to bargain higher wages while reduced downside risk is not as useful in bargaining. It is worth noting that standard BI proposals put the benefit level at about the poverty line. It seems possible that workers cannot credibly threaten to quit a lower paying job for a benefit level set this low. In any case, whether a BI would lead to higher wages remains an open question with the theoretical results here suggesting the answer may be mixed.

The other main question that I want to address with this framework is whether a BI would lead to fewer dirty jobs. This involves the interaction of two firm decisions. The first is whether to open a vacancy as a clean or a dirty job, which is determined by the difference in the values of unfilled vacancies of each type. That is, a firm opens a clean vacancy iff:

$$V_{cj}^{v} > V_{dj}^{v} \tag{11}$$

Using the definitions of the values of vacancies and the equations for bargained wages, this is equivalent to:

$$(1-\gamma_2)[(R_c-1)\epsilon_j - c_j] > -\gamma_5\tau \tag{12}$$

If $R_c = 1$, this condition reduces to $c_j < \tau$ (i.e., whether the cost of cleaning up the job is less than the cost to the worker of working in a dirty job), but whether or not $R_c = 1$, the condition determining whether a given firm opens a clean vacancy is not a function of $T(w)$. That is, whether the environment includes a BI instead of IA will not affect a firm's decision. In this dimension, claims that a BI could support workers in a way that would cause firms to shift to offering clean ('good') jobs are not correct.[8]

But a BI could affect whether a firm operates at all, which, of course, is determined by the non-negative profit conditions, $V^v_{cj} \geq 0$ and/or $V^v_{dj} \geq 0$ (depending on which job mode is more profitable for the given firm type). We can represent these conditions in c_j, ϵ_j space. In figure 3, I plot the line representing whether a firm opens a clean or a dirty vacancy which, under an assumption that $R_c = 1$ is a horizontal line with height τ. I also plot the $V^v_{cj} \geq 0$ line determining whether a clean firm would operate at all. That line is upward sloping because the more productive is the match in terms of ϵ_j, the higher is the cost of cleaning that will still allow for a positive match surplus. Job types with combinations of c_j and ϵ_j that are below this line (and above the horizontal axis since the costs of cleaning cannot be negative) will operate. The intersection of this line with the horizontal axis is

$$\bar{\epsilon} = [\frac{1}{(1-\gamma_2)}[\gamma_1(b+d+T(w_{jc})) + \gamma_3\sum_{k\in C}\eta_k w_{jc} + \gamma_4\sum_{l\in D}\eta_l(w_{ld}-\tau)]] \tag{13}$$

and only (clean) firms with $\epsilon_j > \bar{\epsilon}$ will operate. In essence, this condition says that the productivity of a firm must cover the value of a worker's outside option plus their tax obligations. Under an IA system, the implied wage from the bargain must also be above the break-even level or workers won't participate. In a BI system, the ϵ corresponding to the break-even wage when $c_j = 0$ could be above or below $\tilde{\epsilon}$ depending on the values of b, the tax back rate, and the relative bargaining power of workers and firms. If it is below $\tilde{\epsilon}$ then the presence of the added benefit to workers while working pushes down the productivity level needed to operate relative to what it would be if $\tilde{\epsilon}$ were above the break-even level. In that case, moving to the BI system will increase the number of operating clean firms. There is also an operating condition for dirty firms. If $R_c = 1$, that amounts to a vertical line - a

[8]The only qualification to this is that $w_{dj} > w_{cj}$ and so the two wages could be on different sides of the BI break-even point for some job types, which would result in a function of $T(w)$ entering the clean job decision. The number of job types in this situation is likely to be small, though.

cut-off determined entirely by the ϵ_j value since dirty firms do not engage in clean up and, so, the cost of clean up c_j does not enter the operating decision.

Returning to figure 3, the upper left quadrant corresponds to firms that would choose to operate dirty jobs but with ϵ_j values that are too low for it to be profitable for them to open such vacancies, so this quadrant is empty. The lower left quadrant corresponds to clean operating firms with lower productivity and only the firms in the area below the $V_c^v = 0$ line (area A) will operate. The $V_c^v = 0$ line intercepts the horizontal line of height τ at an ϵ value corresponding to the zero profit line for dirty firms. Thus, to the right of this line, firms with $c_j < \tau$ will operate and will open clean jobs (area B). Firms with higher clean-up costs will open dirty job vacancies (area C).

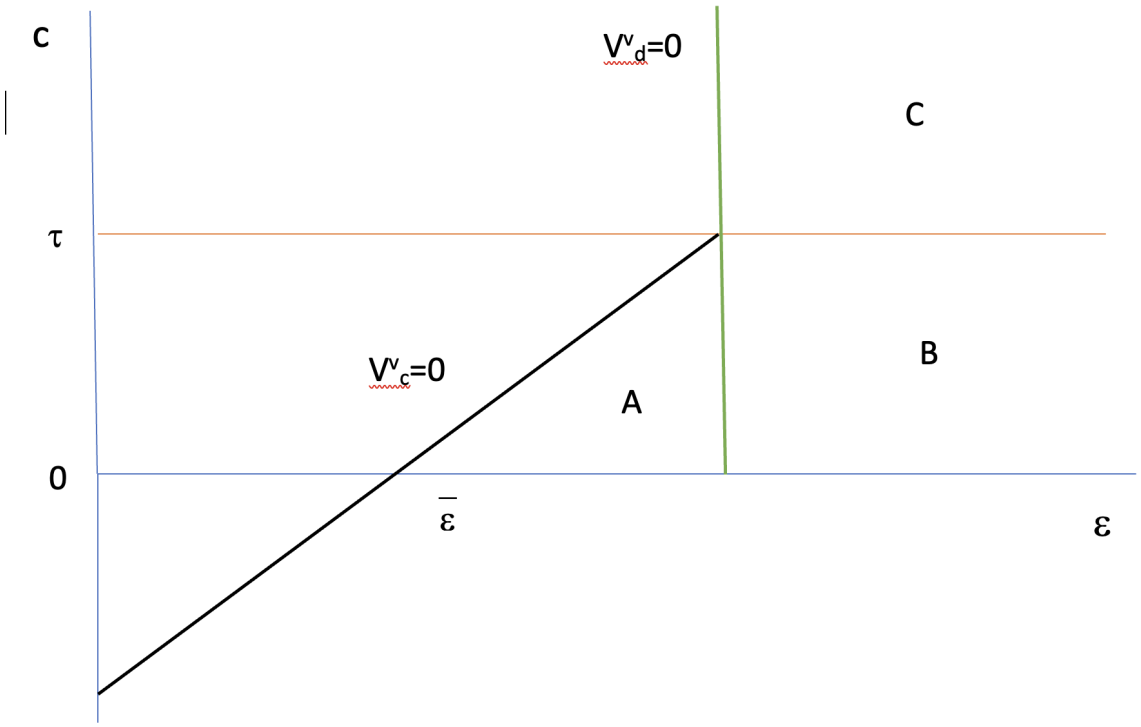

Figure 3: Firm Operation and Choice of Type

If we switch from an IA system to a BI system, the impacts on the proportion of firms that are dirty is ambiguous for a few reasons. First, under the same reasoning we saw with wage setting, the higher level of b, the non-work transfer, pushes in one direction while the fact that benefits are delivered to workers with wages below the break-even level pushes in the opposite direction. The values of ϵ at which a firm would be willing to operate have to cover the outside option of the worker and, so, are higher when b is higher, but don't

have to do this to the same extent when workers receive benefits while working. Suppose, though, that the effect of b dominates. In that case, shifting to a BI system will shift both of the zero profit lines to the right. This is shown in figure 6. The shaded area between the diagonal solid and dashed lines represents the set of clean firms that will no longer operate. The shaded area between the vertical solid and dashed lines represents the set of dirty firms that will no longer operate. This brings us to our other source of uncertainty: whether these movements amount to an increase or decrease in the proportion of jobs that are dirty will depend on the distribution of firms over (c_j,ϵ). The impact of switching to a BI on the extent of dirty jobs is an empirical question, but there is no theoretical reason to expect that it will shift toward fewer dirty jobs because it will lead to the shut-down of both types of jobs.[9]

In **?**, Daron Acemoglu presents a DMP model in which firms can create more or less capital intensive jobs. The more capital intensive jobs can be held up for higher wages and, so, are considered 'good' jobs by workers. Acemoglu shows that firms do not take account of the worker preferences and end up under-producing good jobs. Could something similar happen here and, if so, would a BI be a policy response? In the model presented here, a worker prefers that the job is clean if:

$$w_{jc} > w_{jd} - \tau \tag{14}$$

It is straightforward to show that when $R_c = 1$, this condition collapses to $c_j > \tau$. In other words, workers would use the same criterion as firms if they were in charge of deciding whether to make a job clean. This happens because the costs of clean up and the disamenity value from a job being dirty happen as flows, affecting the common surplus which the worker and firm divide up and have the same incentive to maximize. In contrast, in **?**, the cost of creating a good job is a sunk cost up front and the bargained wage reflects the worker holding up the firm for some of that cost. In that situation, they are not jointly optimizing in the same way and the firm does not take account of the worker's desire for higher wages when it is deciding on which type of job to create.

In the world of the current model, there is no reason for regulation. The workers and firms have come to an agreement that neither would alter if given the option. A more extreme version of this is wage posting models where workers are indifferent about wage, amenity

[9]Michael Veall pointed out to me that a stronger statement might be made if workers have different skills and if lower skilled workers are more likely to face dirty jobs (which seems plausible, empirically). In that case, even an equal shutting down of dirty and clean firms for lower skilled workers due to a shift to a BI will lower the proportion of dirty jobs in the economy since it won't affect firm decisions related to high skilled workers whose incomes are too high for them to receive the benefit. I believe, though, that the conclusion that moving to a BI will not cause any firms to move from being a dirty to a clean firm remains.

combinations at the margin. Allowing workers to search longer by giving them a better benefit will, in principle, allow them to buy more of what they like (clean jobs) but since there is already a compensating differential for dirty jobs (a higher wage to compensate for having to work with dirt), what is being generated is a general income effect not a correcting of a problematic distribution of job types.

The philosopher Elisabeth Anderson in her 2017 book, *Private Government: How Employers Rule Our Lives (and Why We Don't Talk about It)*, examines work relationships in the context of the standard neoclassical model. Models of this type, including the new classical monopsony model (e.g., **?**), treat the work relationship as one of relatively equal power in which both sides can walk away when the arrangement is not to their liking. But, as Anderson points out, a market exchange of work for pay is not the same as a market exchange of money for carrots. In the latter situation, both parties walk away after the exchange and the types of social interactions that are key elements in justice need not play a substantial role (although the work on fairness and exchange prices for goods and labour in **?** indicate they are often still present). But in a labour exchange, the worker is granting control over themselves to the employer. That is, the worker gives up at least some autonomy. This is necessary in order to obtain efficient production (which would not occur if a worker were hired but was then completely free to disregard orders from a boss), implying a necessary trade-off. Anderson's contention is that the pendulum in that trade-off has swung too far in one direction (especially in the US and especially with the decline of the labour movement). But regardless of the existing balance, the surrender of at least some autonomy brings challenges to self-perceptions and respect. The central question is how to create a labour market conducive to a fair society with equality in respect and social standing. A labour market which involves too much surrender of autonomy seems to fail in this sense practically by definition.

In the standard neoclassical model, workers can decide how much autonomy they want to surrender, shopping around among employers and demanding a wage premium for jobs in which they are given less control or respect. Anderson argues that what happens inside the firm is, to some extent, walled off from the operations of the outside market. Workers are tied to jobs through social ties or forming a self-image that is tied to a particular job or the threat of employers providing a bad reference for them. It is then a matter of luck whether a worker ends up in a bad or a good job. This has parallels in the search and bargaining model in which workers meet jobs at random and then stay with them even though there are other jobs that pay more or have better working conditions (e.g.,**?**). And there are reasons to think that a perspective in which we don't have a compensating differentials equilibrium is at least partly correct. The economist Walter Block, taking an extreme efficiency view of markets,

argued that there is no reason for there to be regulation about workplace harassment because if such harassment is going on then the worker is being paid a compensating differential for it and disrupting that market exchange with regulation makes both parties worse off.(**?**) The fact that such a position seems extreme and unreasonable highlights that the work arrangement is not an efficient market exchange.

To the extent that the fully efficient market model doesn't seem to capture reality, our natural inclination as economists is to look for market failures of various kinds. In the search and bargaining model, they arise because workers can't direct their search. As we have seen, in **?**, firms create too few 'good' jobs. This doesn't happen in model developed here because the cost is an ongoing flow that becomes part of the surplus that workers and firms bargain over. In **?**, workers cannot offer to pay firms to create more good jobs because the job creation happens before they meet. We could create that situation in this context by assuming that firms have to sink costs of cleaning up a job site before meeting a worker. Interestingly, as **?** shows, in the context of that sunk cost model, an increase in b, the non-working transfer, essentially allows workers to search longer and wait for good jobs. This leads to a shift in composition toward the efficient distribution of job types. However, regulation, which directly raises the costs of bad jobs, is a more effective way to alter the job type distribution. Again, a BI is a possible solution but more direct and effective approaches exist.

But is this approach of identifying a market failure and then asking which policy tool will best correct it, the best way to examine the problem of creating more just labour exchanges? I think the answer is no for two reasons. The first is an ethical argument. If a key problem with bad jobs is that they damage worker's self-respect and rob them of autonomy then a cash payment is not the correct response. A key tenant of justice is that the remedy should operate in the same realm as the problem. Paying someone cash to compensate them for an insult may, for example, heighten the feeling of being insulted rather than remedy the problem. The right realm would be to alter the work arrangements to remove the direct insult to dignity. I think this is a key problem with economic models when thinking about the justice of worker-firm relationships since we write our models in terms of individual utility and, ultimately, monetary equivalents, thus taking us away from a consideration of different realms of exchange. Put another way, when we broaden our lens to consider outcomes in terms of justice then we are led to consider jobs as locations of self-respect and self-image rather than simply as vehicles of increased income and reduced leisure. In that context, it becomes harder to think about reducing everything to monetary equivalents.

The second reason is that our models miss the nature of the problem at hand. Changing workplace conditions in a situation in which individuals do not freely move among jobs

requires the workers to act collectively. Indeed, I believe that an accurate model of the creation of workplace amenities is one that involves both the employer and the set of employees, with a critical mass of the latter needing to act in order for changes to occur. Put another way, the set of employees at a workplace is a community and, in fact, is a key community for a person's notions of self-respect. One individual in that community deciding to walk away from the job because she doesn't like the level of amenities won't be sufficient to alter the work conditions because frictions in the labour market prevent market discipline of the kind that happens in a simple neoclassical model. Using a BI as a response in this situation involves hoping that each of the individuals takes their personal backstop as a means to engage in a community action. This might happen but there is no clear reason why it would. Maybe it would allow them all to walk away from a bad job but, as we have seen, it is not clearly the case that this will lead to a reduction in the proportion of bad jobs since it funds walking away from good jobs (in the hopes of finding an even better job) as well.

A problem with a BI as a response to an over-abundance of bad jobs is that it ignores both of these issues. It acts as if money is the right realm and that backstopping individual effort is sufficient to correct the problem. But that would only be true in a neoclassical world and in that world, we are either in a compensating differential equilibrium - in which case, no response is actually needed - or the more directly effective response is regulation not cash. In a world where individual workers don't have sufficient agency through the market to bring about change, a BI is not the right response.

6 Conclusion

My goal in this paper was to examine the relationship between a BI and the labour market in order both to try to make a contribution to the debate over a BI as a central policy tool and as a basis for considering deficiencies in our main models of the labour market. My focus was on three main claims about a BI: that it would affect paid labour supply; that it is necessary as a response to emerging technological upheaval in the labour market; and that it would provide a backdrop that would allow workers to force firms to provide jobs with better wages and working conditions. I argue, as do many proponents, that the paid labour supply effects of the BI proposals that have been put forward in the Canadian context are likely to be small. Importantly, I argue that these small estimated predictions arise not because people do not respond to incentives in choosing whether and how much to work but because different groups of people would respond in opposite ways to the introduction of a BI and because the BI systems being proposed generate incentives that are not dramatically

different from the system that already exists. That is, a BI system would not actually lower the 'welfare wall'.

On the claimed need for a BI stemming from technological change associated with AI and robotics, I point out that we are not actually in the midst of the end of work or even a trend rise in precarious forms of work. Instead, we are in a situation of 'persistent precarity' in which substantial proportions of workers face unstable working conditions with low wages and those proportions have been high for a long time. That points to a need to evaluate the third BI claim - that a BI can lead to an improvement in working conditions. It is also important to keep in mind that technological change is not an exogenous force beyond our control. It is an endogenous outcome of decisions in the economy that can be shaped by policies. That changes the question from, do we need a BI as a response to the unstoppable force of technological change, to, what role would a BI play in steering us toward one technological path or another?

In the final section of the paper, I examine the question of whether a BI could alter the quality of jobs and the level of wages by giving workers greater bargaining power. Working within a standard search and bargaining model, I show that there is no reason to think that it would have either of these effects. Theoretically, a BI could even lead to decreases in wages - particularly at the low end. But more importantly, the analysis points to deficiencies in our models that become particularly evident if we consider policies in light of a goal of creating a more just society. Once we look through that lens, the relationship of work to social and self respect becomes more salient. The correct response to denigrating and otherwise problematic work conditions are policies that focus on elements of respect in the workplace. Treating those as reducible to monetary equivalents, as we tend to do in our models, misses the extent to which the work relationship is special and different from other exchanges in the economy. In addition, a focus on self and social respect highlights people's need for community as the basis of that respect. Truly improving working conditions requires engaging the community of workers at a work site. Again, our models are not well suited to investigating how to do that. Improving our models would provide a better basis for discussing these key issues and their specific relationship to a BI. But on the face of it, a BI doesn't appear to be the best policy for improving job characteristics since it only delivers cash (which, again, is the wrong realm of consideration) and it is given to individuals in the hope that they might use it to backstop community building behaviours without any particular evidence that it would actually do so.

Overall, I would argue that labour market related claims in favour of a BI do not stand up under closer scrutiny. At the same time, there is little reason to fear that the BI proposals

put forward in Canada would generate large declines in labour supply. And there are reasons to conclude that we need to extend our models if we are going to be able to provide more definitive answers about the role of any policy in promoting a move toward a more just society.

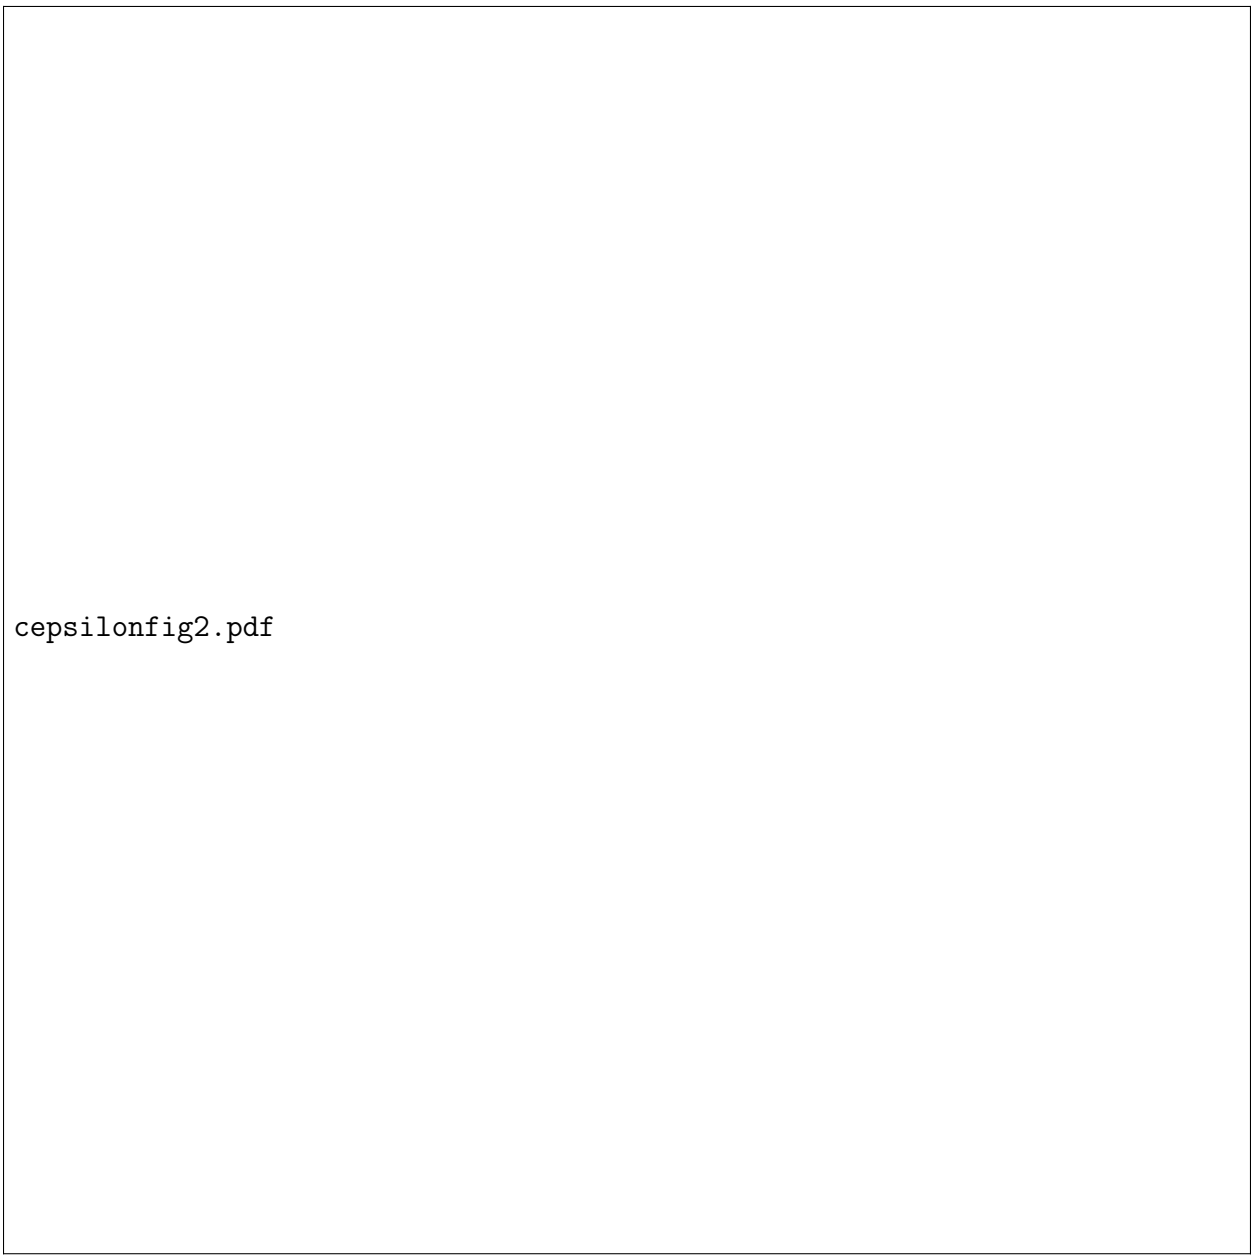

Figure 4: Firm Operation and Choice of Type, Impact of Shifting to a BI