The status of workers and platforms in the sharing economy

Andrei Hagiu* and Julian Wright†

June 20, 2018

Abstract

We consider whether workers who provide their services through online platforms like Handy and Uber should be classified as independent contractors or employees. Using a formal economic model, we show how being too strict or too liberal in classifying workers as independent contractors (relative to the actual degree of control workers have) can be detrimental, not just to firms and welfare, but sometimes to the workers themselves. We also use the model to explore the extent to which an intermediate classification of workers between employees and independent contractors may lead to better outcomes. The intermediate classification is meant to apply to firms that retain control of some actions while their workers control others, as is the case for many online platforms.

1 Introduction

As the sharing (or gig) economy grows in importance, policymakers have been grappling with what regulations and laws should apply. A key area of concern regards the proper legal status of workers who provide their services to consumers through online platforms such as Amazon Mechanical Turk, Freelancer, GrubHub, Handy, Instacart, Postmates, Task Rabbit, Uber, and Upwork. Should Handy cleaners, Uber drivers and other "gig workers" be considered employees or independent contractors? On the one hand, online platforms typically give their workers more control over when they work, how they work, and for whom they work relative to traditional employers, which suggests the workers should be thought of as independent contractors. On the other hand, the workers are oftentimes entirely reliant on the associated platform for their income, have little or no bargaining power, and some important decisions are controlled by the platform (e.g., the price charged to riders is controlled by Uber, not by drivers, and Handy cleaners must follow restrictive guidelines for how to interact with their customers).

*MIT Sloan School of Management, Cambridge, MA 02142, E-mail: ahagiu@mit.edu
†Department of Economics, National University of Singapore, Singapore 117570, E-mail: jwright@nus.edu.sg

There is little agreement on how to classify workers in such settings. The platforms involved have invariably classified workers as independent contractors, making them ineligible for most employee benefits and labor protections.[1] Indeed, Handy has proposed bills in eight state legislatures in the U.S. to permanently classify most gig workers as independent contractors, with the bills passing the house or senate in each state, and so likely to become law soon. At the same time, several courts have ruled that Uber drivers should be classified as employees, and eligible for the associated benefits and rights.[2] The regulators' concern is that some platforms are attempting to have it both ways: avoid paying payroll taxes and other employment benefits, while at the same time reaping the benefits of significant control over how their workers perform services for their customers.

In this paper we develop a simple model which allows us to analyze the implications of classifying workers as employees rather than independent contractors, even when workers have a significant degree of control over how they perform their services. In particular, we highlight an important inefficiency that arises when policymakers take an over-inclusive view of employment. The resulting cost burden imposed on platforms can lead them to make inefficient choices regarding the degree of control and the revenue shares given to their workers. These inefficient choices can sometimes be detrimental not only to the platforms themselves, but also to their workers.

In our benchmark model, in which there is a single costly revenue-enhancing action that can be controlled either by the firm or a worker (e.g., expenditure on the quality and maintenance of the relevant equipment, or hours of work), the key distinguishing feature of a platform relative to a traditional employer is that workers have control over this costly action. Thus, when workers control the costly action, the firm (i.e., the platform) must give them a larger share of variable revenues in order to incentivize them to choose high levels of the action. Indeed, real-world platforms typically let workers keep a large share of the revenues generated, but at the same time provide no fixed payments to the workers by way of salaries or other fixed benefits. In contrast, traditional firms which tend to take control over revenue-enhancing costly actions prefer to attract workers mostly through fixed salaries and benefits, since that leaves the firm with a greater share of variable revenue, thereby minimizing the distortions in its own choices of these actions. In this context, if the law requires firms that give workers a lot of control to provide the same amount of fixed benefits as traditional employers, firms may reoptimize how much control they grant workers and how much revenue to share. In particular, in our model whenever workers are liquidity constrained (so firms cannot simply offset the fixed benefits by charging workers upfront fees), misclassifying workers that retain control over the costly action as employees can lead the firm to take back control from workers and/or to give workers a lower share of revenue, both of which can be inefficient.

The effect of worker misclassification on worker welfare can go in either direction in our model. When firms reoptimize by taking back control over the costly action and shifting to a fixed salary model, workers are never better off since their salary is now pinned down by their outside option.

[1] For example, in the United States, independent contractors are not protected under the Fair Labor Standards Act. Thus, they are not guaranteed a minimum wage, do not receive workers' benefits such as holidays, insurance, maternity leave or paid sick leave, and are responsible for the full cost of social security contributions.

[2] See Artecona and Chau (2017) for a survey of the issues and a discussion of some of the cases.

They will in fact be worse off under worker misclassification if the platform prefers to leave them with some positive surplus under the correct classification in order to better incentivize their choice of the revenue-enhancing costly action. On the other hand, when firms continue to act as platforms despite worker misclassification, leaving workers in control of the costly action, and only reoptimizing the revenue share that workers receive, workers can either be better off or unaffected.

We extend our model to allow for multiple (two) costly actions, such that the efficient allocation of control rights requires the firm to control one action and the worker to control the other. We use this extension to show the benefits of introducing a third work classification, intermediate between employees and independent contractors. When firms choose this third category, which we refer to as *dependent contractors* (also known as independent workers), they grant some but not all control rights to their workers, and are only required to provide their workers with some but not all of the benefits normally required for employees. More generally, we think firms should be viewed along a continuum determined by the extent of control that they give their workers over various aspects of the services provided. This continuum makes it clear that platforms that give workers a lot of control are substantially different from traditional employers, and not just service firms that use technology as an excuse to skirt regulations. Consequently, we think it makes sense to treat them differently from traditional employers or service firms.

We are not aware of any existing formal economic analysis that studies the welfare implications of how workers are classified. The only exception is our own recent work (Hagiu and Wright, 2018) in which we provide a theory of when transferable control rights should be allocated to the firm or to the worker. In that paper, an important feature of the model is double-sided moral hazard, which complicates the analysis and is why we rely on numerical examples to examine what happens if workers that have full control are misclassified as employees. The current paper provides a much more general analysis by abstracting from double-sided moral hazard, and instead incorporating efficiency differences between the worker and the firm at making the costly decision. This enables us to enrich the model and capture some key features relevant to the issue at hand. In particular, we allow for the costs incurred by the firm when providing fixed worker benefits to be different than the value of the benefits to the workers, and we also consider the possibility of introducing the intermediate work classification (dependent contractor). Neither of these features were present in the numerical analysis performed in Hagiu and Wright (2018).

Katz and Krueger (2016) document the significant rise of alternative work arrangements in the U.S. economy from 2005 to 2015. There is also a burgeoning academic literature that discusses whether sharing economy workers should be viewed as employees or independent contractors, and that reviews some of the associated court cases. For example, Means and Seiner (2016) argue that the defining question which determines how workers should be classified is whether the workers are "truly free to choose the time, place, price, frequency, and manner of the work". Bales and Woo (2017) review Uber court cases, arguing that the standard control test and the economic realities test for worker classification provide mixed or indeterminate results, which may explain why Uber court cases have resulted in different outcomes in different jurisdictions. Recognizing that gig economy

work relationships do not fit easily into the existing legal categories of "employee" and "independent contractor", Harris and Krueger (2015) propose a new hybrid category labeled "independent workers," in which workers have the ability to choose when to work, with whom to work, and whether to work at all, while the firm retains significant control over the way independent workers perform their jobs. Independent workers would qualify for some of the benefits and protections that employees receive, such as tax withholding, and employer contributions for payroll taxes, but not for others, such as hours-based benefits like overtime and minimum wages. Their proposed hybrid solution is an example of the dependent contractor classification we explore. Cherry and Aloisi (2017) provide a comparative analysis of using such a hybrid approach based on the experiences from various countries. As noted above, none of this literature is based on formal economic modeling.

A possible unintended consequence of requiring firms to provide fixed benefits for workers (even when they remain free to work for multiple platforms) is that an individual platform may impose exclusivity on its workers in order to prevent rival platforms from free-riding on the fixed benefits it provides. Bryan and Gans (2018) analyze the implications of such exclusivity in the context of ride-sharing platforms, showing it has non-trivial consumer market implications in terms of prices and wait time, thereby suggesting another channel by which worker misclassification could reduce total surplus.

At a high level, our approach of defining worker classification based on the allocation of control rights relates our paper to the older literature on the theory of the firm (Williamson, 1975, Grossman and Hart, 1986, Hart and Moore, 1990 and Holmstrom and Milgrom, 1994). One similarity is the prediction that high-powered incentives (e.g. larger revenue share) should go hand-in-hand with low-powered incentives (control over revenue-enhancing actions or ownership of assets). However, a key difference is that in the theory of the firm literature based on property rights and incentive systems, the key instrument determining the choice of organizational mode is the split of asset ownership (the classic make vs. buy decision). This determines the ex-post payoffs earned by the various parties from their respective outside options. Thus, different configurations of asset ownership must lead to different relative configurations of outside options in order for a tradeoff to exist between make and buy. By contrast, in our framework, the key instrument is the allocation of control rights over non-contractible decisions that are chosen ex-post and affect joint payoffs (as in Hagiu and Wright, 2018).

2 Model

Consider a stylized model of a service firm that requires one worker to perform some costly action to generate revenue. The level of this costly action is denoted x, with $x \geq 0$. It can either be controlled by the firm, which we will call the *employment mode*, or by the worker directly, which we will call the *agency mode*. The agency mode captures the case of a platform that hosts the worker, where the worker controls the provision of services to the consumers (i.e., by choosing the level of x).

Suppose the revenue generated from the costly action is $R = R_k(x)$ where $k = f$ when x is chosen

by the firm and $k = w$ when x is chosen by the worker, provided both the firm and the worker are willing to participate. Whichever party chooses the level of x incurs the associated cost $C(x)$, which is the same for the firm and the worker. We assume $R_f(.)$ and $R_w(.)$ are strictly increasing, concave and twice-differentiable, while $C(.)$ is strictly increasing, convex and twice differentiable. Furthermore, for $k \in \{f, w\}$, we assume $C(0) = 0 \leq R_k(0)$, $R_k'(0) > C'(0)$ and $R_k(x) < C(x)$ for some sufficiently high finite x.

Let

$$x_k \equiv \arg\max_x \{R_k(x) - C(x)\} > 0$$

and the associated profit

$$\pi_k \equiv R_f(x_k) - C(x_k).$$

For example, if $R_k = \beta_k x$ with $\beta_k > 0$ and $C(x) = \frac{1}{2}x^2$, then $x_k = \beta_k$ and $\pi_k = \frac{\beta_k^2}{2}$. We will refer to this as the linear-quadratic example.

Implicit in our setup is that consumer surplus is subsumed into revenues, which is why we measure welfare by the sum of the firm's profit and the worker's surplus. This could reflect, for example, that consumer surplus is always fully extracted. We assume the worker has an outside option worth $u_o > 0$, and the firm makes a take-it-or-leave-it offer to the worker.[3] We focus on a simple type of contract, which is widely used in practice. Specifically, we assume the firm can pay the worker a fixed wage W and some share of the revenue R that is generated (i.e., a bonus or commission). Let $t \in [0, 1]$ be the share of revenue R that the firm takes, so the worker receives $(1 - t)R + W$. The terms of the contract (W, t) are specified by the firm at the contracting stage. Importantly, we assume that the fixed wage W must be non-negative, so that the firm cannot charge the worker a fixed fee. This could be because the worker faces a liquidity or credit constraint. In a context with multiple workers and network effects, this could also reflect a coordination problem: charging fixed fees creates an equilibrium in which no worker joins and so there is a failure to launch. It is also worth noting that in practice firms do not normally charge any fixed fees to workers: this is true both for traditional firms hiring employees and for platforms enabling workers to provide services directly to customers.

We assume that the law requires the firm to provide the worker certain fixed benefits if the relevant authorities deem that it operates in employment mode (so it is considered to be employing the worker). Suppose that these benefits are valued by the worker at B and they cost the firm F to provide. We assume $B \leq F$ throughout in order to capture that the firm would not offer these benefits unless it is required to do so.[4] We also impose that $B < u_o$, so that the worker would not work just to get B. If instead the firm is deemed to operate in the agency mode (so the worker is considered to be in control of x), then the firm does not have to offer B to the worker. We assume that the worker is willing to participate if he gets to control x and keep all the revenue, which is true provided $\pi_w > u_o$, and that

[3] This assumes that the firm has all the bargaining power. An important implicit assumption is that the worker's bargaining power remains unaffected by the worker's classification.

[4] We have also analyzed what happens if $B > F$. In brief, either we find (i) misclassification is irrelevant in this case, or (ii) the welfare results are qualitatively the same as those with $B \leq F$ because B is sufficiently close to F. Details are available from the authors upon request.

the firm is willing to participate even when it is required to offer the worker B, which as we will show is true provided $\pi_f - F > u_o - B$.

We will consider three possible regulatory regimes, depending on whether the firm is correctly classified by regulators or not:

1. Correct classification: the firm is deemed to be in employment mode (and therefore obligated to offer benefits B) whenever it controls the level of x, and in agency mode whenever x is chosen by the worker.

2. Employment misclassification: the firm is always deemed to be in employment mode (and therefore always obligated to offer benefits B), even if it is actually in agency mode.

3. Agency misclassification: the firm is always deemed to be in agency mode, even if it is actually in employment mode.

For any given regulatory regime, the timing of the game is as follows:

1. The firm determines the contract, which involves choosing who has control over x and setting (W, t).

2. The worker decides whether or not to accept the contract.

3. Provided the worker accepts the contract, the party given control over x chooses the level of x and incurs the associated cost. If the firm is deemed (correctly or incorrectly) to be in employment mode, it provides benefit B to the worker at cost F. Revenues are distributed according to the contract.

3 Results

In order to determine the equilibrium outcome under the three different regulatory regimes, we first determine the equilibrium outcomes under employment mode and agency mode, assuming in each case the firm offers B at the cost F. The case in which B is not offered can then be calculated simply by setting $B = F = 0$. This will allow us to work out the optimal choices of the firm when it is correctly classified, determine whether it wants to offer B, and explore what happens under the two types of misclassification.

3.1 Employment mode

If the firm chooses x, then it has no reason to pay any share of the revenue to the worker, so it will set $t = 1$ and $x = x_f$. If it provides the benefit B, it will pay the worker the fixed wage $W = u_o - B$ to ensure the worker's participation. The firm's resulting profit is $\pi_f - u_o + B - F$ and the worker obtains a payoff of u_o. If the firm does not provide the benefit B, the firm's profit is just $\pi_f - u_o$.

Thus, in employment mode, the firm extracts all the variable profit and just compensates the worker with a fixed wage. Any fixed benefits that the firm is required to give the worker just offset the wages it would otherwise pay. Since $B \leq F$, the regulatory requirement to provide the fixed benefits lowers the firm's profit in employment mode.

3.2 Agency mode

If the worker sets x and retains a share $1-t<1$ of the revenue generated, then he will set x to maximize $(1-t) R_w(x) - C(x)$, provided this together with any fixed wage leaves him with sufficient surplus that he wants to participate. Denote the resulting choice of x by $x_w(t)$. Given that R_w is concave and C is strictly convex, x_w must be strictly decreasing in t from the implicit function theorem. Thus, $x_w(t) < x_w$.

Denote the associated worker surplus by the function $u_w(t) \equiv (1-t) R_w(x_w(t)) - C(x_w(t))$. Assuming the worker receives the fixed benefit B, the worker obtains $B + u_w(t)$. It is straightforward to verify that the function $u_w(t)$ is strictly decreasing in t given our assumptions. The firm solves

$$\begin{aligned} & \max_{t,W} \{tR_w(x_w(t)) - W - F\} \\ \text{s.t.} \quad & u_w(t) + B + W \geq u_o \\ & W \geq 0. \end{aligned}$$

Thus, the worker must be paid a fixed wage of $W = \max\{u_o - B - u_w(t), 0\}$ to be willing to participate. Recall that $W \geq 0$ reflects our assumption that the worker cannot be charged a fixed fee. It is easily seen that, since $B \leq F$, the firm does not want to offer the benefit B (and therefore incur the cost F) if it has a choice.[5] However, the firm may have to offer the benefit if it is misclassified.

Note then that the constraint $W \geq 0$ must be binding, so that $W = 0$. Indeed, if $W > 0$, then the optimum t would have to be $t^* = \arg\max_t \{tR_w(x_w(t)) + u_w(t)\} = 0$. But then $u_w(0) = \pi_w > u_o$, so neither constraint is binding, which means the firm could lower W. The firm's program simplifies then to

$$\begin{aligned} & \max_{t} \{tR_w(x_w(t)) - F\} \\ \text{s.t.} \quad & u_w(t) + B \geq u_o. \end{aligned}$$

There are two cases depending on whether the remaining constraint is binding or not.

(i) **Low** u_o: If the constraint $u_w(t) + B \geq u_o$ is not binding, then the firm's profit is $\max_t \{tR_w(x_w(t))\} - F$. Note the revenue the firm extracts $tR_w(x_w(t))$ is zero when $t = 0$ and when $t = 1$, and we assume R and C are such that $tR_w(x_w(t))$ is single peaked in t over $[0,1]$,[6] with $t_w \in (0,1)$

[5] Denote by $(W(B), t(B))$ the optimal solution when the benefit is offered. Then, when the benefit is not offered, the firm can choose $(W = W(B) + B, t = t(B))$, which clearly satisfies the constraints and yields higher profits than the optimal solution when B is offered (strictly higher if $F > B$).

[6] The firm's profit is single peaked in the linear-quadratic example. More generally, it can be shown using the implicit

defining the unique maximizing level of t. The maximizer t_w represents the unconstrained revenue share that the firm takes in agency mode. Denote the profit the firm extracts in this case by $\pi_w(t_w) \equiv t_w R_w(x_w(t_w))$, so the firm's profit is $\pi_w(t_w) - F$. This case arises when $u_o \leq u_w(t_w) + B$: the level of t maximizing the revenue extracted by the firm is sufficient to make the worker willing to participate. The worker's payoff is $u = u_w(t_w) + B$. To the extent that $u_o < u_w(t_w) + B$, the worker is left with more surplus than his outside option—the firm cannot extract more of the worker's surplus because it cannot charge the worker a fixed fee (by assumption) and because increasing t would decrease the firm's revenue.

(ii) **High** u_o: If instead, $u_o > u_w(t_w) + B$, then the constraint $u_w(t) + B \geq u_o$ is binding and the firm will have to let the worker keep more than the fraction $1 - t_w$ of the revenue it generates in order to get him to participate. The firm will therefore optimally lower t below t_w to the point where the worker is just willing to participate. Let this constrained level of t be written as $t_w(B)$, where $t_w(B) \equiv u_w^{-1}(u_o - B) < t_w$. Note $t_w(B)$ is increasing in B. This gives the firm a profit of $\pi_w(t_w(B)) - F$ and the worker a payoff of u_o.

If the firm does not offer the fixed benefit in agency mode, then all the results above still hold provided we set $B = F = 0$ everywhere.

To understand these results, note that without any liquidity or credit constraint, the firm would do best by charging an upfront fixed fee to the worker but then let the worker control x and keep all the revenue generated (i.e., $t = 0$). Without the ability to charge workers such a fixed fee, the firm determines t by trading off the additional share of revenue it gets with a higher t with the distortion in the choice of x by the worker. The unconstrained solution to this problem arises when the resulting surplus for the worker is sufficient to cover his participation constraint. This is case (i) above, which arises with low u_o. However, if the unconstrained solution does not leave enough revenue for the worker to be willing to participate, then the firm will have to increase the share of revenue it leaves with the worker to ensure the worker's participation. This is case (ii) above, with high u_o. Notice the firm never pays the worker any fixed payment (i.e., a wage) in agency mode, which is consistent with what we observe in most platform settings. A positive fixed wage would be suboptimal because it is always better for the firm to compensate the worker with a larger share of variable revenues, thereby reducing the distortion associated with the worker's choice of the costly action x.

3.3 Optimal choices

To work out which mode the firm wants to choose, it is helpful to summarize the firm's profits for different choices, which we do in Table 1. Note the table shows the payoffs when B is provided and when it is not. This will allow us to determine if the firm would actually want to offer B, and also what happens under the two different misclassification regimes, which we consider subsequently.

function theorem that sufficient conditions for $tR_w(x_w(t))$ to be concave are either that $R_w'''(x) \leq 0 \leq C'''(x)$ or $0 \leq R_w'''(x) \leq C'''(x)$ for all x.

Table 1

	Fixed benefits not provided		Fixed benefits provided	
	$u_o < u_w(t_w)$	$u_o \geq u_w(t_w)$	$u_o < u_w(t_w) + B$	$u_o \geq u_w(t_w) + B$
Agency mode	[1] $\pi = \pi_w(t_w)$	[2] $\pi = \pi_w(t_w(0))$	[3] $\pi = \pi_w(t_w) - F$	[4] $\pi = \pi_w(t_w(B)) - F$
Employment mode	[5] $\pi = \pi_f - u_o$		[6] $\pi = \pi_f - u_o + B - F$	

Note the worker's payoff is u_o in all cases except in the case corresponding to cell (1) where it is $u_w(t_w)$ and in the case corresponding to cell (3) where it is $u_w(t_w) + B$

Under the correct classification, we have already explained above that the firm never wants to provide the fixed benefit B, but it is obligated to provide it in employment mode. Thus, directly comparing the firm's profit in cells (1) and (2) in Table 1 with the firm's profit in cell (6) we obtain the following results.

Proposition 1. (Correct classification) *The firm never offers any fixed benefits in agency mode.*
(i) If $u_o < u_w(t_w)$, the firm prefers agency mode iff

$$\pi_w(t_w) > \pi_f - u_o + B - F. \tag{1}$$

(ii) If $u_o \geq u_w(t_w)$, the firm prefers agency mode iff

$$\pi_w(t_w(0)) > \pi_f - u_o + B - F. \tag{2}$$

These results show that under the correct classification, the agency mode is preferred if R_w is sufficiently large relative to R_f, u_o is large (subject to the constraint on u_o for each case), and $F - B$ is large, so fixed benefits are particularly inefficient to provide.

3.4 Misclassification

In this section we explore the implications of the two misclassification regimes. Recall that misclassification arises either when the authorities always interpret the firm as hiring independent contractors (agency misclassification), or when the authorities always interpret the firm as hiring employees (employment misclassification). We will consider the implications of each regulatory regime in turn, relative to what happens with correct classification. (The proofs of propositions are given in the Appendix.)

Proposition 2. (Agency misclassification)
(i) If the firm prefers employment mode under correct classification, agency misclassification does not change the firm's choice of mode, it increases the firm's profit, it leaves the worker unaffected, and it increases welfare.
(ii) If the firm prefers agency mode under correct classification, $u_o < u_w(t_w)$ and $\pi_f - u_o > \pi_w(t_w)$, agency misclassification causes the firm to switch to employment mode, it increases the firm's profit, it decreases the worker's payoff, and it has an ambiguous effect on welfare.

(iii) If the firm prefers agency mode under correct classification, $u_o \geq u_w(t_w)$ and $\pi_f - u_o > \pi_w(t_w(0))$, agency misclassification causes the firm to switch to employment mode, it increases the firm's profit, leaves the worker unaffected, and it increases welfare.
(iv) In all other cases, the firm's and worker's payoffs are unaffected by agency misclassification.

From the perspective of worker welfare, the proposition captures one potential downside of being too liberal in classifying workers as independent contractors. Without having to provide fixed benefits to workers, agency misclassification sometimes leads the firm to prefer taking control of the costly action (employment mode), and so extracting all the surplus from the worker until his participation constraint binds. This is in contrast to the agency mode, where the worker is sometimes left with a positive surplus since the firm needs to incentivize the worker's choice of costly action (case (ii) in Proposition 2). Moreover, because the firm extracts more of the worker's surplus when it switches to employment mode, this raises the possibility that the firm switches to employment mode even when this lowers total welfare. However, in all other cases, agency misclassification either has no effect or it leads to a Pareto improvement. It leads to a Pareto improvement if it allows a more efficient choice to be made (either in terms of changing to a more efficient mode or avoiding an inefficient fixed benefit) without affecting the worker's payoff when the worker's participation constraint is binding both before and after the misclassification.

Now consider the effects of employment misclassification relative to what happens with correct classification.

Proposition 3. (Employment misclassification)
(i) If the firm prefers employment mode under correct classification, the firm's and worker's payoffs are unaffected by employment misclassification.
(ii) If the firm prefers agency mode under correct classification, $u_o < u_w(t_w) + B$ and $\pi_w(t_w) \geq \pi_f - u_o + B$, employment misclassification does not change the firm's choice of mode, it decreases the firm's profit, it increases the worker's payoff, and decreases welfare.
(iii) If the firm prefers agency mode under correct classification, $u_o < u_w(t_w)$ and $\pi_w(t_w) < \pi_f - u_o + B$, employment misclassification causes the firm to switch to employment mode, decreases the firm's profit, decreases the worker's payoff, and decreases welfare.
(iv) If the firm prefers agency mode under correct classification, and either (a) $u_w(t_w) \leq u_o < u_w(t_w) + B$ and $\pi_w(t_w) < \pi_f - u_o + B$ or (b) $u_o \geq u_w(t_w) + B$, employment misclassification decreases the firm's profit, leaves the worker unaffected, and decreases welfare.

The proposition shows that taking an over-inclusive view of employment is generally bad for welfare. To understand this result note that when the firm prefers agency mode under correct classification, the agency mode also leads to higher welfare. This reflects that the firm's profits are aligned with total welfare given the worker's participation constraint is always binding, with the only exception being that under agency mode the worker also gets an additional benefit when u_o is low. Thus, whenever

requiring the firm to offer the additional B pushes the firm to instead choose the employment mode, this decreases the firm's profit, (weakly) decreases the worker's surplus, and decreases welfare.

If instead, the firm wants to stay in the agency mode, having to provide the additional fixed benefit B will lead to one of two effects. In case (ii) of Proposition 3, the firm would not want to adjust t since it is already extracting the maximum commission revenue, in which case employment misclassification results in a transfer from the firm to the worker. This provides some support for the view that platforms should be forced to offer various standard benefits to help better compensate workers, although note that $F \geq B$ implies this transfer leads to (weakly) lower welfare. The other possibility is that the firm will increase t to offset the fact that the provision of B relaxes the worker's participation constraint. Since a higher t leads to a greater distortion in the worker's choice of the costly action in agency mode, this decreases the firm's profit, and decreases welfare. Workers are worse off under employment misclassification when the firm previously left them with some positive surplus in order to better incentivize their choice of x. Thus, it is also quite possible that forcing platforms to provide fixed benefits to workers that are actually independent contractors is not only bad for firms and welfare, but also for the workers that authorities are trying to protect.

4 Hybrid case

We now extend the previous model by allowing for two actions, x_1 and x_2. We make the simplifying assumption that x_1 only generates revenue $R(x_1)$ when chosen by the worker at cost $C(x_1)$, while x_2 only generates revenue $R(x_2)$ when chosen by the firm at cost $C(x_2)$. In other words, the firm has no competency at choosing x_1 and the worker has no competency at choosing x_2. Thus, the most efficient outcome involves the worker controlling x_1 and the firm controlling x_2, which we call the hybrid mode. However, this may not be the chosen outcome if the firm is required to provide fixed benefits when it controls x_2. In this case, it may instead prefer to give up control over both actions to the worker (agency mode).

This setting is closer to the real-world scenarios of firms like Uber and other sharing economy firms, which have less control over worker decisions than traditional employers, but more than pure platforms. In particular, this allows us to examine the possibility that some scholars have advocated of having an intermediate classification for the hybrid mode (e.g., see Harris and Krueger, 2015), in which workers are eligible for some fraction of the benefits that are required in employment mode. To do so in our framework, we assume that when the firm chooses the hybrid mode and controls only one of the actions (x_2 here, given our assumptions), it has to pay some fraction αB of the employment benefits B, which costs the firm αF, where $0 \leq \alpha \leq 1$. We continue to assume $F \geq B$. We refer to the classification of workers associated with the hybrid mode as *dependent contracting*.

To keep the analysis tractable, we directly focus on a version of the linear-quadratic example. Specifically, assume revenue is $R(x) = \beta x$ and cost is $C(x) = \frac{1}{2}x^2$, where we maintain our assumptions from the previous section, namely $u_o < \frac{\beta^2}{2}$. Our goal is to explore what happens when the dependent contracting category is introduced ($0 < \alpha < 1$) relative to what happens when the worker is classified

as an employee if the firm chooses the hybrid mode ($\alpha = 1$).[7]

We first note that the firm never prefers the employment mode regardless of whether the dependent contracting category exists or not. Given the employment mode requires maximum fixed benefits be provided, which the firm wants to avoid, this provides no advantage relative to the hybrid mode. Indeed, in hybrid mode the firm can always choose the same contract as in the employment mode, resulting in the same outcome as the employment mode, but can generally do better by reoptimizing its contract to take into account that the worker can generate additional revenue by choosing a positive level of x_1 and by only offering a fraction α of the fixed benefits. Thus, we restrict attention to the agency mode and the hybrid mode in what follows.

Consider the agency mode. The firm solves the following program:

$$\begin{aligned} & \max_{t,W} \left\{ t\left(1-t\right)\beta^2 - \alpha F - W \right\} \\ \text{s.t.} \quad & \frac{\left(1-t\right)^2\beta^2}{2} + \alpha B + W \geq u_o \\ & W \geq 0. \end{aligned}$$

Since $B \leq F$, it is easily seen from the program above that the firm's profits are decreasing in α, so that the firm will never offer any benefits in agency mode.[8]

There are then two possibilities:

- if $u_o < \frac{\beta^2}{8}$, then $t^* = \frac{1}{2}$ and $W^* = 0$, so $u = \frac{\beta^2}{8}$ and $\pi = \frac{\beta^2}{4}$.
- if $u_o \geq \frac{\beta^2}{8}$, then $W^* = 0$ and $t^* = t_w\left(0\right) = 1 - \frac{\sqrt{2u_o}}{\beta} \leq \frac{1}{2}$. Substituting this solution into the firm's profit function gives the profit in Table 2 below.

Now consider the hybrid mode. The firm solves the following program:

$$\begin{aligned} & \max_{t,W} \left\{ \left(t\left(1-t\right) + \frac{1}{2}t^2 \right)\beta^2 - \alpha F - W \right\} \\ \text{s.t.} \quad & \left(\frac{1}{2}\left(1-t\right)^2 + \left(1-t\right)t \right)\beta^2 + \alpha B + W \geq u_o \\ & W \geq 0. \end{aligned}$$

Again, since $B \leq F$, the firm's profits are always decreasing in α, so the firm will offer the minimum benefit permitted by law in the hybrid case, which with a slight abuse of notation, we also denote as αB. Note the firm's profit is increasing in t and the worker's payoff is decreasing in t for $t \in [0,1]$. It

[7] The worker is classified correctly as an employee if the firm chooses employment mode and is classified correctly as an independent contractor if the firm chooses agency mode.

[8] Consider $\alpha > \alpha'$ and let $\left(W\left(\alpha\right), t\left(\alpha\right)\right)$ denote the optimal solution to the firm's program when it offers benefits αB. Suppose now the firm offers benefits $\alpha' B$. Then it can set $W' = W\left(\alpha\right) + \left(\alpha - \alpha'\right)B$ and $t' = t\left(\alpha\right)$, which clearly satisfy the two constraints in the optimization program, and yield higher profits than what the firm was obtaining with $\left(W\left(\alpha\right), t\left(\alpha\right)\right)$ when it was offering benefits αB (because $B < F$). This implies optimal profits are higher when the firm offers benefits $\alpha' B$.

is straightforward to see that at the optimum the first constraint must be binding, which means the agent always obtains u_o. There are then only two possibilities:

- If $u_o < \alpha B + \frac{3}{8}\beta^2$, then $W^* = 0$ and $t^* = t_h(\alpha B) = \frac{\sqrt{\beta^2 - 2(u_o - \alpha B)}}{\beta} \geq \frac{1}{2}$.
- If $u_o > \alpha B + \frac{3}{8}\beta^2$, then $t^* = \frac{1}{2}$ and $W^* = u_o - \alpha B - \frac{3}{8}\beta^2 > 0$.

Substituting these solutions into the firm's profit function gives the profit in Table 2 below. Note in case of high u_o, the firm maximizes the sum of its profit and the worker's surplus, which is why it sets $t = \frac{1}{2}$. Note that the firm always chooses a (weakly) higher t in hybrid mode than in agency mode. This reflects that in agency mode only the agent has control over costly actions and incentivizing the agent requires a low t, whereas in hybrid mode, the firm and the agent each have control over one action, so the optimal t trades off the agent's incentives (low t) and the firm's incentives (high t). Note also that the worker's participation constraint is always binding in hybrid mode, but is only binding for high u_o in agency mode. To understand this, note that if the participation constraint is not binding, then the optimal wage must be $W^* = 0$ (in either mode). The difference then is that in hybrid mode, the firm wants to set t as high as possible, which eventually makes the participation constraint binding because the worker's payoff is decreasing in t. However, in agency mode, the firm never wants to set t above 1/2, as explained above.

The following Table summarizes the profit expressions derived from the above solutions, both for the agency mode, in which no fixed benefits are provided, and the hybrid mode in which αB fixed benefits are provided. The case in which the dependent contracting category is available corresponds to the case in which $\alpha < 1$, while the case without the dependent contracting category corresponds to $\alpha = 1$.

Table 2

	$u_o < \frac{1}{8}\beta^2$	$\frac{1}{8}\beta^2 \leq u_o \leq \frac{3}{8}\beta^2 + \alpha B$	$u_o > \frac{3}{8}\beta^2 + \alpha B$
Agency mode	$\pi = \frac{1}{4}\beta^2$	$\pi = \sqrt{2u_o}\left(\beta - \sqrt{2u_o}\right)$	
Hybrid mode	$\pi = \frac{1}{2}\sqrt{\beta^2 - 2(u_o - \alpha B)}\left(2\beta - \sqrt{\beta^2 - 2(u_o - \alpha B)}\right) - \alpha F$		$\pi = \frac{3}{4}\beta^2 - u_o + \alpha(B - F)$

Note the worker's payoff is u_o in all cases except under the agency mode when $u_o < \frac{\beta^2}{8}$ in which case $u = \frac{\beta^2}{8}$.

Using the profit expressions in Table 2, we can now determine the effect of introducing the dependent contracting category.

Proposition 4. (Introduction of dependent contracting category)
(i) If the firm prefers the agency mode with and without the dependent contracting category, then the introduction of the dependent contracting category has no effect on the firm's profit, the worker's payoff or welfare.
(ii) If the firm prefers the hybrid mode without the dependent contracting category, the introduction of the dependent contracting category increases the firm's profit, leaves the worker's payoff unchanged, and increases welfare.

(iii) Suppose the firm prefers the agency mode without the dependent contracting category and prefers the hybrid model with the dependent contracting category. In this case, if $u_o < \frac{\beta^2}{8}$*, then the introduction of the dependent contracting category increases the firm's profit, decreases the worker's payoff, and has an ambiguous effect on welfare. The welfare effect is more likely to be positive when* u_o *is larger,* α *is smaller,* B *is larger, and* F *is lower. If* $u_o \geq \frac{\beta^2}{8}$*, then the introduction of the dependent contracting category increases the firm's profit, does not change the worker's payoff, and increases welfare.*

In the first two cases in the proposition, the introduction of the dependent contracting category does not change the firm's preferred mode, so the welfare implications are straightforward. The interesting case is when the introduction of the dependent contracting category causes the firm to switch modes, which can only be from agency to hybrid because the dependent contracting category increases the attractiveness of the hybrid mode. In this case, the firm's profit clearly increases. The impact on the worker's payoff depends on the attractiveness of the outside option u_o. If the worker's outside option is large enough, the worker's participation constraint is binding in both modes, so the change in total welfare is equal to the change in the firm's profits, and therefore positive. If on the other hand the worker's outside option is small enough, then the worker's participation constraint is binding in hybrid mode, but not binding in agency mode (i.e., the worker gets more than u_o in agency mode which the firm cannot extract because it cannot charge the worker a fixed fee and because increasing t would give the firm less revenue). Consequently, in this case the introduction of the dependent contracting category leads to a lower payoff for the worker. Because the firm extracts more of the workers' surplus when it switches to the hybrid mode, this raises the possibility that the firm switches to the hybrid mode even when this lowers total welfare. This distortion decreases in u_o and α, so welfare will be higher when the firm switches to the hybrid mode provided u_o is not too low and α is not too high. Obviously, welfare is also more likely to increase with this switch when fixed benefits are not too inefficient (i.e., $F - B$ is lower).

5 Conclusion

We have provided the first formal analysis of how the (mis)classification of workers as independent contractors or employees can matter for firms' profits, workers, and overall welfare.

One limitation of our analysis is that in employment mode, workers never get more than their (exogenously given) outside option. There are potentially two problems with this. First, employees could enjoy greater bargaining power (e.g., through unions) than independent contractors, so their surplus exceeds their outside option in employment mode. This means there can be higher gains to workers when true independent contractors are misclassified as employees, and also greater losses to workers when true employees are misclassified as independent contractors. Second, we have not accounted for the possibility that forcing a large number of firms to provide certain benefits could change the outside option for workers, which we have treated as exogenous. Capturing either or both of these additional effects would enrich our analysis, and remains an interesting direction for future

research. Finally, our treatment of the case with multiple actions was limited to a special symmetric setting. Future research on the use of the dependent contracting category could provide a more complete analysis of this case.

References

[1] Artecona, Raquel, and Terence Chau (2017) "Labour issues in the digital economy," Working paper.

[2] Bales, Richard A., and Christian P. Woo (2016) "The Uber Million Dollar Question: Are Uber Drivers Employees or Independent Contractors," *Mercer L. Rev*, 68 (2016): 461.

[3] Bryan, Kevin, and Joshua Gans (2018) "A theory of multihoming in rideshare competition," Working paper.

[4] Cherry, Miriam A., and Antonio Aloisi (2016) "Dependent contractors in the gig economy: a comparative approach," *American University Law Reiew*, 66 (2016): 635.

[5] Hagiu, A, and Julian Wright (2018) "Controlling vs. enabling," *Management Science*, forthcoming.

[6] Harris, Seth D., and Alan B. Krueger (2015) "A proposal for modernizing labor laws for twenty-first-century work: The independent worker," Hamilton Project, Brookings, 2015.

[7] Katz, Lawrence F., and Alan B. Krueger (2016) "The rise and nature of alternative work arrangements in the United States, 1995-2015," National Bureau of Economic Research.

[8] Means, Benjamin, and Joseph A. Seiner (2015) "Navigating the Uber economy," *UCD Law Review*, 49: 1511.

6 Appendix

6.1 Proof of Proposition 2

Agency misclassification has no effect on the firm's profits in agency mode. It increases the firm's profit in employment mode from $\pi_f - u_o - (F - B)$ to $\pi_f - u_o$. If the firm prefers the employment mode under correct classification, then the firm will prefer to stay in employment mode and its profit will increase. The worker is unaffected as he still gets u_o.

Now suppose the firm prefers the agency mode under correct classification. If $u_o < u_w(t_w)$, comparing the firm's profit across cells (1) and (5) in Table 1 implies the results in case (ii) of Proposition 2: if the firm prefers to switch to employment mode, it must be better off by the difference $\pi_f - u_o - \pi_w(t_w)$, while the worker is worse off by the difference $u_w(t_w) - u_o$. The welfare change is $\pi_f - u_o - \pi_w(t_w)$, which can be positive or negative.

Alternatively, if $u_o > u_w(t_w)$, comparing the firm's profit across cells (2) and (5) in Table 1 implies the results in case (iii) of Proposition 2: if the firm prefers to switch to employment mode, it must be better off by the difference $\pi_f - u_o - \pi_w(t_w(0))$, while the worker's payoff is unchanged at u_o, meaning welfare is unambiguously higher.

6.2 Proof of Proposition 3

Employment misclassification has no effect on the firm's profit in employment mode. Thus, if the firm prefers the employment mode under correct classification, then making it provide the additional benefit B in agency mode will reduce the attractiveness of the agency mode and therefore the firm continues to prefer the employment mode and nothing changes. This is case (i) in the proposition. So the only case of interest is when the firm prefers the agency mode under correct classification. Consider each case in the proposition in turn.

(ii) If $u_o < u_w(t_w) + B$ and $\pi_w(t_w) \geq \pi_f - u_o + B$, comparing across cells (3) and (6) in Table 1, the firm will want to stay in agency mode. There are two subcases. Suppose $u_o < u_w(t_w)$. Then employment misclassification increases the worker's surplus by B and decreases the firm's profit by F, so the net effect on welfare is $B - F \leq 0$. If on the other hand $u_w(t_w) \leq u_o < u_w(t_w) + B$, then employment misclassification increases the worker's payoff from u_o to $u_w(t_w) + B$ and decreases the firm's profit from $\pi_w(t_w(0))$ to $\pi_w(t_w) - F$. The change in total welfare is

$$\pi_w(t_w) + u_w(t_w) - \pi_w(t_w(0)) - u_o + B - F = \pi_w(t_w) + u_w(t_w) - \pi_w(t_w(0)) - u_w(t_w(0)) + B - F,$$

which is negative because $\pi_w(t) + u_w(t)$ is decreasing in t, $t_w(0) < t_w$ and $B \leq F$. Since the change in the worker's payoff is positive, this also implies the change in the firm's profit must also be negative.

(iii) If $u_o < u_w(t_w)$ and $\pi_w(t_w) < \pi_f - u_o + B$, comparing across cells (3) and (6) in Table 1, employment misclassification leads the firm to switch to employment mode even though it was better off in agency mode under correct classification. Thus, the firm's profit decreases. The worker's payoff also decreases, from $u_w(t_w)$ to u_o. Thus, total welfare unambiguously decreases.

(iv) Consider the two cases in turn. If $u_w(t_w) \leq u_o < u_w(t_w) + B$ and $\pi_w(t_w) < \pi_f - u_o + B$, comparing across cells (3) and (6) in Table 1, the firm prefers to switch to employment mode, which decreases the firm's profit. The worker's payoff remains unchanged and equal to u_o, so total welfare decreases. If on the other hand, $u_o \geq u_w(t_w) + B$, then regardless of whether the firm switches modes, the worker obtains u_o. The firm's profit necessarily decreases and therefore so does total welfare.

6.3 Proof of Proposition 4

Since the firm's profit in hybrid mode is decreasing in α, if the firm prefers the agency mode with the dependent contracting category ($0 < \alpha < 1$), then it must have preferred the agency mode without the dependent contracting category ($\alpha = 1$), so the introduction of the dependent contracting category is irrelevant. On the other hand, if the firm preferred the hybrid mode without the dependent contracting category, then it must still prefer the hybrid mode with the dependent contracting category,

and indeed is better given this is equivalent to a decrease in α. Since workers still get $u = u_o$, total welfare must increase.

The remaining case is that the firm switches to the hybrid mode with the introduction of the dependent contracting category, due to the higher profit that is now available. There are two possibilities:

- if $u_o < \frac{\beta^2}{8}$, then the switch from agency to hybrid mode makes the worker worse off by $\frac{\beta^2}{8} - u_o > 0$, which implies the effect on total welfare is ambiguous. In particular, the change in total welfare is
$$\left(\frac{1}{2}\sqrt{\beta^2 - 2\left(u_o - \alpha B\right)}\left(2\beta - \sqrt{\beta^2 - 2\left(u_o - \alpha B\right)}\right) - \alpha F - \frac{\beta^2}{4}\right) - \left(\frac{\beta^2}{8} - u_o\right).$$
This expression is increasing in u_o, decreasing in α, increasing in B, and decreasing in F.
- if $u_o \geq \frac{\beta^2}{8}$, then the switch from agency to hybrid mode leaves the worker's payoff unchanged and equal to u_o, which implies the effect on total welfare is positive.