Macroeconomic Implications of Executive Pay Caps

Francesco Caselli*

June 2025

Abstract

I study the introduction of statutory caps on excutive pay ratios (pay caps) in a span-of-control economy with heterogenous managerial talent. Managers hire fewer workers than they would in the absence of pay caps, and the marginal product of labor is not equalized among firms. The ability of the marginal manager declines. Inequality between managers and workers, and among managers, declines, but inequality among workers increases. Quantitatively, overall inequality decreases substantially, especially via a reduction in top income shares, but the GDP losses are also significant. In addition, the redistribution favours the "middle" deciles between 50 and 90, rather than the bottom 50% of the income distribution. Future work will build on this "first pass" to analyze the effect of pay caps in economies distorted by monopsony power, non-pecuniary benefits of control, and managerial rent extraction.

*London School of Economics, Centre for Macroeconomics (CFM), and Centre for Economic Policy Research (CEPR). Thanks: Thomas Brzustowski, Russell Cooper, Wouter Den Haan, Pedro Gomes, Johannes Matt (reasearch assistance), Chris Pissarides, Silvana Tenreyro, Martin Weale, and Shengxing Zhang. Also seminar participants at the LSE, Bocconi, ETH Zürich, Singapore Management University. I acknowledge financial support from CFM and declare that I have no financial relationships or other potential conflicts of interest in undertaking or circulating this research.

1 Introduction

In recent years and in multiple countries there have been calls for the introduction of statutory limits on the pay of top corporate officers, relative to the employess who work under them. As recently as 2023 Congresspersons in the United States have submitted legislation limiting the ratio of CEO pay to median pay in the company to 50 - with progressive punishing taxation for firms violating the cap. The British Labour Party's 2017 manifesto called for a similar cap set at 20 for companies bidding for government contracts. The European Parliament has at various times debated mandatory pay ratios ranging between 20:1 to 50:1. Executive pay ratios ranging from 10 to 50 have been advocated by political forces in France, Spain, Canada, Australia, South Africa, and several other countries.

A partial cap is already in place in Israel's financial system, as well as in Portland, Oregon, and San Francisco, California. There are also many jurisdictions (including the USA, the UK, France, Spain, and India) which mandate public companies to disclose their executive pay ratios. To the extent that these disclosures interact with reputational concerns to deter firms from granting their top executives salaries in even higher multiples of wages, these countries can be thought of as already having a "soft" cap in place.

Proponents of executive pay caps have a variety of objectives, but certainly some of the goals relate to the dramatic changes in the income distribution observed in most market economies over the last quarter century. Two developments are of particular concern for some pay-cap advocates. The first is the remarkable increase in the concentration of income at the top of the income distribution, which is deemed by many to be deleterious for social cohesion, and to have led to an ever-more disproportionate influence of super-wealthy individuals in politics – a rise in influence which many consider undemocratic, if not anti-democratic. The second is the extremely slow growth in the living standards of households in (roughly) the bottom 50% of the income distribution. Proponents hope that pay caps may help ameliorate the system's failure to allow this large constituency to partake in the benefits of growth.

Despite the proliferation of pay-cap proposals, there is little economic analysis of their likely effectiveness in achieving their goals, or of their consequences for economic efficiency. This paper takes a first step in filling this gap, by analyzing the impact of pay-cap policies on inequality and GDP in an otherwise frictionless economy. The frictionless environment is helpful in developing transparent novel insights on the impact of pay caps. As discussed in the Conclusions, however, a full quantitative evaluation of pay cap proposals will require a more realistic description of the economy pre-pay caps - for example by including employer monopsony power.

The framework for this analysis is a span-of-control economy à la Lucas (1978), in which agents have heterogeneous managerial abilities. The model's focus on labour income and its stark manager-worker dichotomy makes it ideal for assessing policies that constrain managerial income relative to worker wages, without having to track the ways in which labour-income inequality interacts with capital income inequality.

The key observation in assessing the impact of statutory pay ratios is that these provisions tie a manager's remuneration to the wage received by the workers under him: the workers need to be paid well if the manager is to be paid well. Given this, the manager's objective becomes very close to maximizing output per worker (and exactly that if the statutory pay ratio is equal to 1). Decreasing marginal productivity of labour pins down the firm size which maximizes managerial income. This firm size is independent of managerial ability, and so in the Pay-Cap economy all managers for whom the pay-cap is binding take on the same number of workers. This contrasts sharply with the Laissez-Faire economy, in which firm size is increasing in managerial ability. Intuitively, in the Laissez-Faire economy the cost of adding an extra worker is set by the market, and so managers are happy to add workers up to the point at which the marginal product of labour equals the wage. With pay caps, instead, each decline in the marginal productivity of labour directly impacts the *per worker* resources available for worker pay. This induces managers to stop hiring earlier that in the Laissez-Faire case.

The fact that firm size is independent of managerial ability implies that, unlike in the Laissez-Faire economy, the marginal product of labour is not equalized across firms. Hence, the Pay-Cap economy is inefficient. As for inequality, the effects of Pay-Caps are ambiguous. Obviously inequality *within firms*, i.e. between managers and workers, declines relative to the Laissez-Faire case. And since managers of different ability now operate firms of the same size, inequality among managers also falls. On the other hand, inequality among workers will tend to increase. The reason is that the (lucky) workers employed by high-ability managers earn more than the ones employed by lower-ability managers.

I calibrate the model so that its Laissez-Faire version matches moments of the US firm size distribution. I solve for key equilibrium macroeconomic aggregates both under laissez faire and under a wide variety of pay caps. The results confront us with a classic efficiency-inequality trade off. Output losses are steeply increasing with the tightness of the pay cap, while the share of income accruing to the top one percent is steeply decreasing. Focusing here on a pay cap of 50, as in the US proposal, the policy causes an output loss of 12 percent, but also reduces the share of the top 1 percent by one third. The impact on the share going to the top 0.1 percent, another frequent focus of recent debates on inequality, falls even more dramatically – by two-thirds. This is a trade-off that may perhaps be acceptable to those who are particularly concerned about the concentration of income and wealth at the very top. Having said that, it is also important to stress that the beneficiaries of this redistribution are not the workers in the bottom fifty percent of the income distribution, whose share of income is essentially insensitive to the pay cap, but those between the 40_{th} and 90_{th} percentiles. The main reason for this is that the fall in the demand for labor by the super-productive managers at the top of the ability distribution allows the "middling" managers below them to hire labour at depressed prices.

A methodological contribution of the paper is to propose an equilibrium concept for a Pay-Cap economy. Under Laissez-Faire, all workers earn the same wage. The reason is that, if worker i is paid less by manager j than worker i' is paid by manager j', worker i can successfully approach manager j' with a proposal which undercuts worker i'. In a Pay Cap economy no such undercutting offer will be accepted. Because the manager's pay is tied to the pay of workers, a manager who takes on a worker at a lower wage than what he currently pays is cutting his own income. The fact that there is no market-clearing wage means that we need to come up with an alternative definition of equilibrium in order to determine who will be a manager and how workers are assigned to managers.

Section 2 describes the basic physical set up of the economy. Section 3 briefly reviews the (well-know) equilibrium solution for the Laissez-Faire economy. Section 4 is where the paper's novel contribution begins. In particular, it proposes an equilibrium concept for a Pay-Cap economy with a unitary pay ratio, and introduces an algorithm to solve for it. The subsequent section extends the model (and the equilibrium concept) to a generic pay-ratio. Section 6 calibrates the model and presents the quantitative comparison between the Laissez-Faire economy and Pay-Cap economies for a broad range of values for the pay cap.

2 Environment

The economy is static with a population of mass 1. Each agent can be either a manager or a worker, and has a managerial ability z, where the CDF of z is $G(z)$ and is defined on $[\underline{z}, \infty)$. Everyone has the same ability when employed as a worker. A production unit can only function if it has a manager. Conditional on having a manager of ability z, output is $zF(l)$, where l is the mass of workers. The manager has a limited span of control, and therefore $F'(l) > 0$ and $F''(l) < 0$. The manager income is $zF(l)$ less the wage bill paid to the workers.

Some words on interpretation are in order. Since the manager is the residual claimant after worker wages, a literal interpretation is that managers in the model are *owner managers*, and thus their income is a mix of labour income and proprietor income. Under this interpretation, a pay cap is a cap on overall labour-plus-proprietor income. Given the arbitrariness of the distinction between these two forms of income, this would probably be the only way of imposing a statutory pay ratio on this class of managers anyway.

A loser, but perhaps more relevant, interpretation is that managers are professional (hired) managers, and that $zF(l)$ is only the *labour share* of the value added produced by the firm. In a fully-articulated version of this interpretation, there would be another category of agents (shareholders, capital suppliers, banks, etc.) entitled to their own share of value added. To be sure, this broader model would be formally identical to the current one only if these external agents provided a fixed input at an exogenous rental rate. In practice, it is likely that managers' demand for additional inputs, as well as these inputs' market prices, would be affected by the introduction of pay caps. But it is reasonable to presume that such effects would be second-order compared to those that the "labor-income only" model allows me to focus on.[1]

The assumption that agents are homogeneous as workers is not without loss of generality. Under worker heteorgeneity, the imposition of a statutory pay ratio generates an incentive for the manager to hire high-ability workers - similar to the incentive for segregation by ability in equalitarian teams [Farrell and Scotchmer (1988)]. This effect is absent in this paper. Having said this, standard supermodularity assumptions tend to produce assortative matching even without pay caps. If managers and workers are already matched by ability in the Laissez-Faire case, the additional incentive for assortative matching coming from the pay cap should have little effect on the comparison between the Laissez-Faire and the Pay-Cap economy. For this reason, in this first-pass I prefer to focus on an effect that is unquestionably only present in the Pay Cap economy.

3 Laissez-Faire Economy

Before studying the consequences of a pay cap, it is important to be reminded of the properties of the span-of-control economy in the absence of constraints on executive pay.

The income of a manager of ability z is

$$zF(l) - wl,$$

[1]The simplest way to allow for a distinction between labour income and other income sources is to add physical capital, and a competitive rental market. In such a model a likely response by managers to a statutory pay ratio would be to reduce the demand for labour (as in my labour-only benchmark) and increase the demand for capital. If the resulting increase in rental rates elicits an increase in the supply of capital, this might reduce the output loss from pay caps - relative to the benchmark in this paper. A more ambitious extension would model explicitly the agency problem between shareholders and managers. I return on this in the Conclusions.

where w is the market wage. An equilibrium is defined by the following requirements.

- Managers maximize their own income taking the wage w as given. This implies that, for every z such that some agents of type z are managers,

$$zF'(l) = w. \tag{1}$$

  We define $l_z(w)$ the solution to this equation for given z, and $\pi_z(w)$ the corresponding maximized values of managerial income. Both $l_z(w)$ and $\pi_z(w)$ are increasing in z.

- Define $\tilde{z}$ by

$$\tilde{z}F(l_{\tilde{z}}(w)) - wl_{\tilde{z}}(w) = w. \tag{2}$$

  Then, all agents with ability $z > \tilde{z}$ ($z < \tilde{z}$) are managers (workers).

- The labor market must clear. The market-clearing condition is

$$\int_{\tilde{z}}^{\infty} l_z(w)dG(z) = G(\tilde{z}) \tag{3}$$

Equations (2) and (3) pin down the equilibrium wage w and the cutoff $\tilde{z}$.

4 Economy with Unitary Pay Cap

We now turn to what happens if the government imposes a rule that the income that managers can take from their firms must equal the minimum salary that pay to their workers, i.e. the statutory pay ratio is 1. In the next section I generalize the analysis to a generic pay cap.

Because of the constraint that everyone earns the same, the income of a manager with ability z who hires l workers is

$$y(z, l) = \frac{zF(l)}{1 + l}.$$

The numerator is the total revenue generated by the firm. The denominator is the number of people sharing (equally) this pie: the manager and the l workers.

Just as in the LF economy, managers choose the number of workers which maximizes their income. However, this translates into very different labour-demand behaviour. In particular, the first order condition for maximization of $y(z, l)$ with respect to l is

$$\frac{zF(l)}{1+l} = zF'(l), \tag{4}$$

which says that average product must be equal to marginal product. The intuition for this optimality condition is well understood from the literature on equalitarian cooperatives [Ward (1958), Vanek (1970)]. If marginal revenue exceeds average revenue the manager can enhance her objective (which is to maximize average revenue, just as in a coop) by hiring more workers - and viceversa.

An immediate consequence of (4) is that firm size in the PC economy is invariant with the ability of the manager, z. We denote this common firm size l^*. Intuitively, an increase in z increases average and marginal income by the same proportion, and thus cancels out in the choice of the optimal employment level. It is easy to foresee that this insensitivity of firm size to managerial skill will be one important source of inefficiency in the PC economy. In particular, it means that the marginal product of labour is not equalized across firms.

To further compare firm size in the LF and PC economies, we can combine (1) and (2) to find that, in the former, the employment level chosen by the marginal manager (whose ability we denoted $\tilde{z}$) also satisfies equation (4) - and hence the marginal firm employs l^* workers. This is a reflection of the LF economy's efficiency (and indeed it is the core observation to prove it, which is easily done). The marginal manager is one who is equally useful as a worker or as an employer, and the two sides of the equation represent her contribution to output under the two choices. The difference, however, is that in the LF economy all non-marginal firms will have employment l larger than l^*, while in the PC economy all firms employ l^* workers. In other words, in the PC economy all firms have the size of the smallest firm in the LF economy.

Since $y(z, l^*)$ is increasing in z, and workers enjoy the same income as their managers, the results derived so far imply that workers assigned to better managers earn more than workers assigned to worse ones. The standard argument for workers receiving a common wage is that a worker receiving a lower wage than another worker can always offer to replace the higher-wage worker for a slightly lower salary. No such undercutting needs exist in a pay-ratio economy. A worker approaching a manager offering to work for less than $y(z, l^*)$ will be rebuffed, because the manager would then have to pay herself this lower income as well.

The absence of an equilibrating wage mechanism need not prevent us from coming up with a plausible equilibrium concept that satisfies sensible rationality criteria. Two observations are pertinent. First, everyone would prefer to work for a manager with higher ability than themselves, rather than being a manager. Second, everyone would prefer to be a manager, rather than working for someone with lower ability than themselves. My proposal for an equilibrium concept, therefore, is to allocate to every agent of type z' who manages a firm l^* workers chosen *at random* from the population with ability $z < z'$, and to assign the role of manager to all the agents who fail to be so allocated.

Heuristically, we could loosely interpret this idea as the result of a costless search process. All agents in the model apply for jobs with all agents with abilities higher than themselves. If they receive one or more job offers, they accept the one from the highest-ability manager among these. If they do not receive any offer, they (reluctantly) become managers themselves, and hire a random sample l^* of agents from the pool who have applied to work for them.

The next definition formalizes this discussion.

Definition 1. *Define*

$$y_z = \max_l \frac{zF(l)}{1+l} = \frac{zF(l^*)}{1+l^*}.$$

An equilibrium for the PC Economy with unitary pay cap is a function $\mu(z)$ which represents the measure of agents of type z who become managers, and a function $\omega(z, z')$ which represents the *fraction of agents of type z who work for a manager of type z', where*

1. *Every agent of type z who is a manager hires l^* workers and earns y_z*

2. *Workers who work for a manager of type z earn y_z*

3. *An agent of type z is either a manager or works for a manager of type $z' \geq z$*

4. *If $A_{z'} = \{z; z' > z\}$, then $\omega(z, z') = \omega(z')$ for all $z \in A_{z'}$, and $\omega(z, z') = \emptyset$ if $z \notin A_{z'}$*

5. *The labor market clears, or*

$$(1 + l^*) \int_{\underline{z}}^{\infty} \mu(z) dz = 1$$

Equilibrium conditions (1-3) and (5) are straightforward and follow immediately from the discussion in the preceding paragraphs without the need for further commentary. Condition (4) is the formalization of the randomization criterion for hiring. The set $A_{z'}$ can heuristically be interpreted as the set of applicants for jobs with managers of ability z'. The condition then simply says that all applicants have equal probability to be offered a job.

The key equilibrium objects are the functions $\mu(z)$ and $\omega(z)$. The definition of the equilibrium implies that these are determined as

$$\mu(z) = g(z) \left(1 - \int_{z}^{\infty} \omega(z') dz' \right)$$

and

$$\omega(z') = \frac{\mu(z') l^*}{G(z')}.$$

The first equation states that managers of type z are the ones "left over" after managers of type $z' > z$ have completed their hiring. The second equation reflects the fact that everyone with ability less than z' has equal probability of being hired by a manager with ability z'. Plugging the latter in the former and solving (calculations in unpublished appendix) we find

$$\mu(z) = g(z) G(z)^{l^*}.$$

This result is very intuitive: the larger l^*, the more workers of type z will have been hired by managers of greater ability - and so the smaller the proportion left to become managers. By the same token, the larger is $G(z)$, the fewer the managers with ability greater than z who can hire agents of ability z. With $\mu(z)$ at hand it is easy to construct all the economy's macroeconomic aggregates (e.g. GDP), inequality statistics, etc.

5 Economy with Generic Pay Cap

In this section we examine the general case, in which a pay-cap $\xi \geq 1$ is imposed on managers. In particular, if a manager pays wage w to her workers, her income is not allowed to exceed ξw.

With generic ξ, the pay-cap may not bind for all managers, requiring analysis of both constrained and unconstrained cases.

If the pay cap is binding the following resource constraint governs the distribution of income within the firm.

$$(\xi + l)w = zF(l).$$

The right hand side is firm revenue. The left hand side is the sum of the payments to workers and to the manager. Using the resource constraint we can solve for the wage and managerial income in the pay-constrained firm:

$$w_z^c = \frac{zF(l)}{\xi + l},$$

$$\pi_z^c = \frac{\xi zF(l)}{\xi + l}.$$

Note that π_z^c is the objective function of the pay-constrained manager and hence employment in her firm must be chosen to maximize it. The first order condition for l implicitly defines the labor demand of a pay-constrained firm, l^c, as

$$\frac{F(l^c)}{\xi + l^c} = F'(l^c).$$

Note that, just like in the special case $\xi = 1$, firm size (in firms managed by pay-constrained managers) is invariant to z. However, $l^c \geq l^*$ and increasing in ξ. Note also that, given l^c constant, both w^c_z and π^c_z are increasing in z.

An unconstrained manager is one whose income $zF(l) - wl$ is less than ξw. An immediate consequence of this is that all unconstrained managers must pay the same wage w^u. The reasoning is the familiar one from traditional labor-market modelling: any unconstrained manager paying higher wages than another manager would be very happy to entertain an undercutting offer from the workers of the lower-paying manager. We will denote w^u the common wage paid by unconstrained managers.

Given that unconstrained managers all face the same wage w^u, their labor demand behaviour $l^u_z(w^u)$ and income $\pi^u_z(w^u)$ as a function of the wage are identical to those in the LF economy. We therefore know that $\pi^u_z(w^u)$ is increasing in z, and hence so is $\pi^u_z(w^u)/w^u$. This therefore means that there exists a threshold $\tilde{z}^c$ such that agents who have ability $z < \tilde{z}^c$ and who manage firms are unconstrained, while managers with ability $z \geq \tilde{z}^c$ are pay-constrained. The threshold $\tilde{z}^c$ is implicitly defined by the equation

$$\frac{\pi^u_{\tilde{z}^c}(w^u)}{w^u} = \xi. \tag{5}$$

Agents with managerial skill z less than $\tilde{z}^c$ who are not employed by managers with ability greater than $\tilde{z}^c$ can either manage firms as unconstrained managers or be employed at wage w^u. They are therefore in the same situation as agents in the LF economy. Accordingly, there is a threshold $\tilde{z}^u \leq \tilde{z}^c$ such that agents with $z < \tilde{z}^u$ ($z > \tilde{z}^u$) strictly prefer to work (manage) at wage w^u. $\tilde{z}^u$ is defined by

$$\frac{\pi^u_{\tilde{z}^u}(w^u)}{w^u} = 1. \tag{6}$$

We can summarize the discussion so far as follows:

- For $z < \tilde{z}^u$, nobody wants to be a manager. Instead, agents seek to maximize the z of the manager they work for - except that they are indifferent among managers with ability $z \leq \tilde{z}^c$

- For $\tilde{z}^u \leq z < \tilde{z}^c$, agents would prefer to work for a constrained manager with ability z' if $\pi_z^u < w_{z'}^c$, and manage otherwise. This condition can be rewritten as $\xi \pi_z^u < \pi_{z'}^c$.

- For $\tilde{z}^c \leq z$, agents would prefer to work for a constrained manager with ability z' if $\pi_z^c < w_{z'}^c$, and manage otherwise. This condition can be rewritten as $\xi z < z'$.

Essentially, the economy with a generic pay cap resembles a mixture of the unitary cap economy and the LF economy. The labour market is segmented. Some lucky workers are hired by high-ability managers who are pay constrained, and who hire l^c employees each. The unlucky ones seek work in the "residual" labour market, where manages of middling ability offer employment at the common wage w^u.

Having clarified how agents' incomes depend on their status and their ability, we are ready to generalize the definition of equilibrium in the PC economy.

Definition 2. *An equilibrium for the PC Economy with pay cap ξ is a function $\mu(z)$ which represents the measure of agents of type z who become managers, a function $\omega(z, z')$ which represents the fraction of agents of type z who work for a manager of type z', thresholds $\tilde{z}^u$ and $\tilde{z}^c$, and a wage for the residual labor market w^u such that*

1. *Every agent of type $z < \tilde{z}^c$ who is a manager earns $\pi_z^u(w_u)$ and hires $l_z(w^u)$ workers. Every agent of type $z \geq \tilde{z}^c$ who is a manager earns π_z^c and hires l^c workers.*

2. *Workers who work for a manager of type z earn w^u if $z < \tilde{z}^c$ and w_z^c if $z \geq \tilde{z}^c$*

3. *An agent of type z is either a manager or works for a manager of type $z' \geq z$*

4. *If $A_{z'} = \{z; max[w^u, w_{z'}^c] > max[\pi_z^u(w^u), \pi_z^c]\}$, then $\omega(z, z') = \omega(z')$ for all $z \in A_{z'}$, and $\omega(z, z') = \emptyset$ if $z \notin A_{z'}$*

5. *The labor market clears, or*

$$\int_{\tilde{z}^u}^{\tilde{z}^c} (1 + l_z^u(w_u))\, \mu(z) dz + (1 + l^c) \int_{\tilde{z}^c}^{\infty} \mu(z) dz = 1 \tag{7}$$

The relationship which defines $\mu(z)$ in the generalized PC model is

$$\mu(z) = \begin{cases} g(z)\left(1 - \int_{\hat{z}(z)}^{\infty} \omega(z')dz'\right) & \text{if } z \geq \tilde{z}^u \\ 0 & \text{if } z < \tilde{z}^u, \end{cases} \tag{8}$$

where $\hat{z}(z)$ is the solution to

$$w^c_{\hat{z}(z)} = \max[\pi^u_z(w^u), \pi^c_z].$$

Compared to the unitary cap case, there are two differences in the definition of μ. First, there is now a marginal manager, $\tilde{z}^u$, who, like in the LF economy, is indifferent between being an unconstrained manager or a worker in the "residual" labor market. Second, the lower band in the integral is now $\hat{z}(z)$, instead of z. This accommodates the new definition of the set $A_{z'}$, i.e. the fact that now potential managers only want to be managed by agents who are sufficiently more productive than them.

Turning to the $\omega(z')$s, we have

$$\omega(z') = \begin{cases} \frac{\mu(z')l^c}{G(\hat{z}^{-1}(z'))} \text{ if } z' \geq \tilde{z}^c \\ \frac{\mu(z')l^u_z(w^u)}{G(\tilde{z}^u)} \text{ if } \tilde{z}^u \leq z' < \tilde{z}^c \\ 0 \text{ if } z' < \tilde{z}^u, \end{cases} \tag{9}$$

where $\hat{z}^{-1}(\cdot)$ is the functional inverse of $\hat{z}(\cdot)$. In other words, if $\hat{z}(z)$ is the marginal constrained manager for which agents with ability $z \geq \tilde{z}^u$ are willing to work, $\hat{z}^{-1}(z')$ is the marginal agent with ability $z \geq \tilde{z}^u$ willing to work for a constrained manager with ability z'.

Equations (5)-(9) jointly pin down $\mu(z)$, $\omega(z)$, $\tilde{z}^u$, $\tilde{z}^c$ and w^u. These equations are complicated and do not lend themselves to a closed-form solution for generic F and G. Nevertheless, a numerical solution is easy to obtain. One guesses a value of w^u, which pins down $\tilde{z}^c$ from (5), $\tilde{z}^u$ from (6), and the function $\hat{z}(z)$ (and hence its inverse). With these values, one substitutes from (9) into (8), and integrates numerically "backward" (from high values of z all the way down to

$\tilde{z}^u$). With $\mu(z)$ at hand one checks whether (7) holds for the initial guess of w^u, and updates the guess in the right direction.

6 Quantification

6.1 Calibration

I quantify this model under the assumption that $F(l)$ is Cobb-Douglas and $G(z)$ is Pareto:

$$F(l) = l^{\alpha}, \tag{10}$$

$$G(z) = 1 - \left(\frac{\underline{z}}{z}\right)^{\gamma} \text{ for } z \in [\underline{z}, \infty), \text{ and } \gamma > 1. \tag{11}$$

This minimal parametrization allows closed-form solutions for all the equilibrium variables of the LF economy, including GDP and inequality statistics (see online appendix).

The model has three parameters: α, γ, and $\underline{z}$. The last of these three has no bearing on ratios of aggregates from the two economies or on comparisons of inequality measures. Hence, we only have to calibrate α and γ.

We presume that the USA economy resembles the LF economy and hence we calibrate the parameters by picking two moments from the data and matching them to the LF version of the model.

Both of the parameters we need to pin down are "production side" coefficients which are closely linked to the firm-size distribution. α governs the speed at which decreasing returns set in, while γ governs the dispersion of productivity, which in turn has a first-order impact on the dispersion of firm sizes. Two available moments from the firm size distribution are average firm size, and the fraction of firms which have at least 10,000 employees. These will be our calibration targets.[2]

[2]Results from an alternative calibration strategy which uses a moment from the managerial income distribution are qualitatively and quantitatively similar.

Using results for $\tilde{z}$ derived in the unpublished appendix, average firm size in the LF economy can be written

$$s = \frac{1}{1 - G(\tilde{z})} = \left(\frac{\tilde{z}}{\underline{z}}\right)^{\gamma} = \frac{\gamma - 1}{(1-\alpha)\gamma - 1}. \tag{12}$$

This is increasing in α (the more slowly marginal productivity declines, the larger the span of control) and decreasing in γ (the thinner the tails of the productivity distribution, the smaller the average firm size).

As for the second calibration target, define λ_q as the share of aggregate employment in firms at least as large as q. In the unpublished appendix I show that this can be written

$$\lambda_q = \frac{\gamma(1-\alpha)q - 1}{\gamma - 1}\left(\frac{\alpha}{(1-\alpha)(q-1)}\right)^{\gamma(1-\alpha)}. \tag{13}$$

Equations (12) and (13) can be jointly solved for α and γ as functions of s and λ_q (for a given choice of q among those available from the data). More details on the joint solution of these equations in the unpublished appendix.

We measure s in the US economy as the total number of employees (including the self-employed), divided by the total number of firms in the economy. Using the 2019 US Census's Annual Business Survey this is 22.4. (The data do not include public sector workers, agriculture, religious services, or private household work). Note that $1/s$ in our model is the percentage of managers in the labour force, so this figure implies that less than 5% are managers. Using the Business Dynamics Statistics of the US Census Luttmer (2010) constructs λ_q for a large number of values of q - and shows that they are remarkably stable over time. For calibration purposes I (arbitrarily) pick $q = 10000$, for which $\lambda_q = 0.25$. The values of α and γ implied by the calibration described above are 0.75 and 4.67.

6.2 Quantitative Results

6.2.1 GDP

The first row of Table 1 shows computed values of GDP for selected values of ξ - all as ratios to GDP in the LF economy (which is the limit of the Pay Cap economy for $\xi \to \infty$). The efficiency loss is monotonic in ξ, reaching a maximum in the unitary pay-cap economy, where 45% of Laissez-Faire GDP is lost. A pay ratio of 50, which, as we have seen, often features in policy proposals, would cause a GDP loss of 12%. Even a pay-cap of 1000 continues to be significantly distortionary, with GDP being still only 93% of LF GDP.

ξ	1	10	50	100	1000	Inf
Macroeconomic Outcomes						
GDP	0.65	0.84	0.88	0.89	0.93	1.00
Avg Firm Size	0.18	0.45	0.58	0.63	0.73	1.00
Marg. Manager z	0.52	0.83	0.88	0.90	0.93	1.00
Inequality						
Top 10% Share	0.17	0.27	0.28	0.29	0.29	0.29
Top 1% Share	0.03	0.09	0.14	0.15	0.17	0.20
Top 0.1% Share	0.00	0.02	0.05	0.07	0.10	0.14
Bott. 50% Share	0.41	0.39	0.39	0.39	0.39	0.39
Between	1.00	3.01	3.98	4.32	5.11	7.00
90-10 Managers	1.82	6.86	6.74	6.76	6.70	6.58
90-10 Workers	1.82	1.20	1.07	1.02	1.00	1.00

Table 1: Aggregate Impact of Statutory Pay Ratios

A key reason for the output loss is that high-quality managers, who hire a disproportionate number of workers in the LF economy, now manage the same number of workers as less able managers. This "firm-size effect" is documented in the second row of Table 1. In the economy with unitary pay caps employment per firm is less than 20% of what it is in the LF economy. This relative firm size rises to 58% for a pay ratio of 50, and is still only just above 70% for $\xi = 1000$.

Since high-ability managers hire fewer workers, the economy ends up recruiting managers from further down the ability distribution. In the third row I report $\tilde{z}^u$ for variuous values of ξ. Recall that this variable captures the ability of the marginal manager, so this comparison tells how "low" does the economy ends up recruiting managers. In the unitary PC economy the least able agent who manages others is about 50% as capable as the least able manager in the LF economy. The ratio rises to about 80% with a pay cap equal to 50, and to 93% for $\xi = 1000$.

6.2.2 Inequality

The second Panel of Table 1 reports the implications of the model for inequality. Much of the discussion of top-executive pay has been in the context of the concentration of income at the top of the distribution. In the Table we can see that executive pay caps are quantitatively effective in reducing top-income shares. The impact is relative modest for the top decile, but quite significant for the top one percent and, even more so, for the top 0.1 percent. A statutory ratio of 50 reduces the income share accruing to the top one percent by more that a quarter, and the share going to the top 0.1 percent by almost two thirds. Even a pay cap of 1000 has a measurable impact on the very top income shares. Proponents of pay caps who are motivated by concerns with the adverse consequences of top-income concentration may regard these results as encouraging.

Some proponents, however, are more (or also) focused on the consequences for the bottom part of the income distribution, and on the potential of pay caps to ameliorate the stagnating living standards of the less well off. Unfortunately, the next row shows that the overall share of income going to the bottom 50% of the labor force is essentially insensitive to the pay cap. The reason for this result is that individuals in the bottom half of the income distribution are typically workers in the "residual" labor market or workers at firms run by relatively-low ability managers. The decline in labor demand by high-quality managers causes the wage in the unconstrained labour market, w^u, to decline, reducing the income share of the first group. The decline in the quality of managers at the bottom of the managerial skill distribution reduces the income share of the latter group.

The implication of decreasing top incomes and stagnant bottom-50 incomes implies, of course, that the distributional winners from pay cap policies are in the "middle" 50-90 income percentile bracket. This is made up of mostly unconstrained managers who benefit from the decline in the wage rate in the residualk labour market, w^u, as the pay cap tightens.

The next row of Table 1 reports the average ratio of managerial income to worker income, as a measure of "between" inequality (the alternative measure, average managerial income divided by the average wage, behaves in a quantitatively very similar way). Between inequality clearly increases in ξ, and again the impact of mandatory pay ratios is substantial even at high levels of the statutory cap. In the economy with unitary pay ratios, there is no between inequality because, for every managerial income, there is a constant measure of workers earning the same income.

The next two rows focus on "within" inequality. We capture inequality among managers and, respectively, among workers by the 90-10 percentile range of the respective income distributions. Not surprisingly, pay caps compress the managerial income distribution. This is because in the LF economy high-ability managers leverage their skills to operate larger firms, thereby amplifying their advantage over lower ability managers. In the PC economy all firms have the same size, so there is no amplification of the differences in ability. More interestingly, pay caps inject inequality in the distribution of income among workers. As already discussed, this follows from the fact that "lucky" workers assigned to better managers earn more. In the extreme unitary case inequality among workers is in fact by construction identical to inequality among managers.

6.2.3 Winners and Losers

In Table 2 I report more granular results on the expected incomes of individuals situated at specific percentiles of the ability distribution.

For all deciles up to the 8_{th}, expected income increases monotonically in the statutory pay ratio. For all these agents, the redistributive benefits of the pay cap are always dominated by the efficiency losses. A pay cap of 50 reduces the income of the bottom 9 deciles (relative to Laissez Faire) by 9%. The reason why the loss is uniform across these deciles is that – at this pay cap – the threshold $\tilde{z}^u$ is above the 90_{th} percentile of the distribution of z. This means that all these agents are strictly workers, and hence face the same expected income.

Perhaps surprisingly, agents above the 9_{th} decile tend to have finite "preferred" statutory pay ratios. This relates again to the fact that pay ratios redistribute the most from the managers at the very very top to those just below them. For example, with a pay cap of 50, even managers in the 99.9_{th} percentile of the distribution of z hire from the residual labor market, which means that they benefit from the decrease in the wage. This last finding implicitly underscores that the pay cap is only binding for a minuscule, albeit highly consequential, tail of the managerial ability distribution: $\tilde{z}^c$ is located at the 99.917_{th} percentile of it. I conjecture that, in some of the more realistic future extensions that I discuss in the next section, the pay cap will be binding on a larger share of managers.

ξ	1	10	50	100	1000	Inf
10_{th}	0.79	0.88	0.91	0.92	0.95	1.00
20_{th}	0.79	0.88	0.91	0.92	0.95	1.00
30_{th}	0.79	0.88	0.91	0.92	0.95	1.00
40_{th}	0.79	0.88	0.91	0.92	0.95	1.00
50_{th}	0.79	0.88	0.91	0.92	0.95	1.00
60_{th}	0.80	0.88	0.91	0.92	0.95	1.00
70_{th}	0.81	0.88	0.91	0.92	0.95	1.00
80_{th}	0.83	0.88	0.91	0.92	0.95	1.00
90_{th}	0.91	0.93	0.91	0.92	0.95	1.00
95_{th}	1.02	1.61	1.35	1.28	1.14	1.00
99_{th}	0.38	1.74	1.45	1.38	1.23	1.00
99.9_{th}	0.09	0.48	1.45	1.38	1.23	1.00

Table 2: Expected Income for Selected Percentiles of z

7 Conclusions

I have studied the introduction of caps on managerial pay ratios in a span-of-control economy with heterogeneous managerial talent. The introduction of pay caps transforms managers into something akin to equalitarian cooperators and leads them to hire fewer workers than they would in the absence of pay caps. This implies that the marginal product of labor is not equalized among firms and that the ability of the marginal manager drops relative to the Laissez-Faire economy. The pay cap does reduce inequality between managers and workers, and among managers, but increases inequality among workers. The pay cap is very effective at containing the income share going to the top one percent and, even more so, to the top 0.1 percent of earners. However, this redistribution tends to favour the "middle class" between the median and the very top, rather than the bottom half of the income distribution. Furthermore, the distortions to firm size and allocation of talent brought about by the pay-cap implies quantitatively large aggregate income losses.

The present framework stacks the deck against pay caps as the LF Economy is fully efficient. Many inefficiencies have been identified which imply that LF economies operate inside the production frontier. The quantitative results are liable to change, most likely in favor of pay-cap policies, in variations of the model in which the LF economy is not fully efficient. In future work, I plan to pursue four such variations.

The first extension will explore monopsony power in labor markets. Monopsony by employers has become a major theme in the debate on the causes of increasing inequality, with increasingly strong evidence of its importance in compressing worker incomes, especially at the bottom. Introducing monopsony power in the model would make the LF economy both inefficient and more unequal, potentially leading to an improvement of the relative attractiveness of pay caps.

In a second variation, I plan to extend the model to include non-pecuniary benefits from control. Empire-building models of managerial motivation are well established, and supported by both formal and anecdotal evidence. In the span-of-control model, such non-pecuniary benefits of control will distort firm size upwards. Since pay-caps tend to distort firm size downwards, it seems conceivable that the introduction of statutory pay ratios might be efficiency enhancing.

Third, the model assumes that both worker and managerial effort (per hour of work) are given, and invariant to the institutional setup. In economics, a long tradition on efficiency wages suggests that worker effort is increasing in pay. There is also evidence, for example in organizational behaviour, that workers respond positively to lower earnings gaps between themselves and management. Allowing for these effects should make the PC economy less inefficient compared to the LF economy.[3]

Fourth, it is well established that the principal-agent problem between managers and shareholders allows the former to extract very significant rents from the latter. A statutory Pay Ratio on managerial income could return some of these rents to the shareholders. To the extent that the shareholder pool includes some lower-income individuals, this would provide a further channel to ameliorate inequality. And to the extent that removal of rents is less distortionary of managers' incentives it may also ameliorate the efficiency losses from the Pay-Cap policy.

References

Farrell, Joseph and Suzanne Scotchmer (1988): "Partnerships," *The Quarterly Journal of Economics*, Vol. 103, No. 2 (May, 1988), pp. 279-297.

Lucas, Robert E. (1978): "On the Size Distribution of Business Firms," *The Bell Journal of Economics*, 9(2), 508–523.

Luttmer, Erzo (2010): "Models of Growth and Firm Heterogeneity," *Annual Review of Economics*, 2010, Vol. 2, 547-576

Vanek, Jaroslav (1970), *The General Theory of Labor-Managed Market Economies.* Cornell University Press.

Ward, Benjamin (1958). "The Firm in Illyria: Market Syndicalism." *The American Economic Review*, vol. 48, no. 4, pp. 566–89.

[3]In the opposite direction, many economists believe that managerial performance would decline with pay caps, though there is virtually no evidence for this.