Three Essays on Inequality and Political Economy

Thesis by

Samantha Myers

In Partial Fulfillment of the Requirements
for the Degree of
Doctor of Philosophy

Caltech

California Institute of Technology
Pasadena, California

2016
(Defended November 16, 2015)

To Mum, Dad, Anita and Paul

Acknowledgments

First I would like to express my sincere thanks to my advisor, Jean-Laurent Rosenthal, for his support, guidance, and patience throughout this long and sometimes painful process. It has been an honor to have the chance to work with someone whose knowledge and curiosity seem to encompass the entirety of human history, and who can see through to the heart of even the most complex problems faster than seems humanly possible.

I am indebted to my committee members, Leeat Yariv, Rod Kiewiet, and Mike Alvarez, each of whom has provided me with invaluable encouragement and advice at various stages of my degree progress. If I have overcome the hurdles that have been placed before me, it is thanks to these people.

I would like to extend my thanks to the whole Caltech faculty, who together have expanded my understanding of economics beyond what I thought was possible. I would like to thank Chris Shannon also, who provided valuable guidance.

I am grateful to everyone who provided me with comments and feedback on my papers. In particular I would like to thank Allyson Pellisier, Khai Chiong, Matt Kovach, Euncheol Shin, Matt Chao, and Kirill Polgorevski, without whom I would not have made it this far.

Thanks must of course go to all of those people who work tirelessly behind the scenes, including Sabrina Boschetti, Gail Nash, and Barbara Estrada and of course Laurel Auchampaugh (whom I am certain is part superwoman) in whose hands the entire graduate program seems to rest.

Last, but never ever least, I must thank my fiancé and colleague Jonathan Chapman. His love, support, and generosity cannot be quantified.

Abstract

This thesis consists of three essays in the areas of political economy and applied microeconomics, covering housing, inequality, and public sector pensions.

Once economies of scale are allowed into the production function, income distribution is no longer necessarily independent of efficiency in a general equilibrium framework. Chapter 2 analyzes a three-good, three-class general equilibrium model where external economies of scale exist on goods consumed primarily by the middle-class, and households have a preference for variety. In the case where the capital-labor ratio is the same across all goods, a transfer from a wealthy home to several poor homes increases consumption by the remaining citizens. Furthermore, if economies of scale are not too small, there exists a transfer from rich to middle-class households that is Pareto-optimal. Finally, the chapter demonstrates that a proportional increase in the size of the economy can offer the same benefits as redistribution.

Cities around the world are experiencing a period of rising house prices and slow development, which is often attributed to increased development regulation. In Chapter 3 I develop a model of the housing market that incorporates congestion, and show that homeowners prefer to restrict housing supply more than renters. I test this model using 40 years of census data from two Australian cities. I demonstrate that public support for regulation can be traced back to homeowners, who seek to restrict supply below market levels in order to elevate the price of their assets and reduce local crowding externalities. I find that a 10 percentage-point decrease in home ownership rates over our period would result in an increase of around 1.6 million dwellings in Sydney, enough to house 14% of the 2011 population on a two-person-per-dwelling basis. A move to a centralized governance structure for the whole city would elimi-

nate this relationship. In addition, I find that regardless of governance structure the proportion of elderly residents is negatively correlated with growth while income is, if anything, positively correlated with growth.

The present value of unfunded local government defined benefit pension liabilities has escalated since the 2008 financial crisis. Chapter 4 considers the arguments for and against switching local government pensions from defined benefit schemes to defined contribution schemes. I note that the relative generosity of government pension schemes compared with the private sector is not necessarily tied to the pension structure. Also, defined benefit pensions may sometimes be cheaper for employers than defined contribution plans, as they allow employers to pool risks that they must otherwise compensate their employees for bearing. I use propensity score matching to test the difference between total remuneration of county employees in Nebraska (defined contribution) and Kansas (defined benefit). I find that there is little difference between the cost of the total package under either plan. Furthermore, any differences are not consistent in either magnitude or sign between different sub categories of employees, or over time. I therefore conclude that there is no strong financial reason to recommend either plan structure over the other.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

In this thesis, I address three research questions in the areas of inequality and political economy. I examine three major areas of contention in current debates on the causes and outcomes of inequality: the benefits of redistribution, the costs and availability of housing, and the structure of post-retirement pensions.

In the wake of increasing income inequality, interest in its potential social and economic costs has mounted. In Chapter 2 of this thesis, I outline a situation in which an economy may benefit from redistribution. If economies become too unequal and agents have a preference for variety, they become unable to leverage economies of scale in production. Even in economies with a realistic distribution of income between rich, middle-class, and poor households, I find that the economy as a whole can benefit from an intervention that re-allocates income from a subset of very wealthy households to a subset of poor households. More powerfully, we find that a shift of the tax burden from the middle class to the rich can generate a Pareto-improvement for all members of the economy, including those households that shoulder an increase in taxation. By allocating income to those who will voluntarily choose to spend it on the right subset of goods, the economy can leverage untapped economies of scale to reduce prices.

Housing inequality has been cited as both a possible outcome of, and an input to, income inequality. Regulatory restrictions on housing supply can reduce affordability and compound intergenerational inequality. In addition, housing regulation in one jurisdiction has spillover effects on pricing and supply in other jurisdictions. In

Chapter 3, I consider the political economy behind housing regulation. Using 40 years of census data from two Australian cities I demonstrate that low supply growth is negatively correlated with the number of homeowners within a self-governed district. Between districts in a centrally governed city this correlation disappears, suggesting that the main driver of this correlation is political. I find a 10 percentage-point decrease in home ownership rates over the period would result in an increase of around 1.6 million dwellings in Sydney and the gains would be potentially even higher if housing were regulated on a Sydney-wide basis, rather than by smaller local governments. I hypothesize that these results are based on a combination of asset price protection and the aversion of residents to local congestion.

In contrast with what one may expect I do not find a negative correlation between income inequality and housing growth. In fact, if anything higher incomes in a local neighborhood are correlated with increased housing development in that area. The proportion of retirement-age (over 65) residents is, however, strongly negatively correlated with neighborhood housing growth, and this is not dependent on the political structure in question.

Notwithstanding any role in housing markets, retirees face hurdles regarding the availability and security of their pensions. While defined contribution plans have now taken the place of defined benefit plans in most private sector industries, this transformation has come with a shift of risk from employers to employees. Chapter 4 examines whether the same shift in pension structure would be beneficial in the public sector. The public sector has yet to make a wholesale move toward defined contribution pensions, and there is ongoing debate as to which structure is preferable in a public sector context. On the one hand, it is argued that defined benefit pensions represent a very large transfer from taxpayers to workers, and are prohibitively expensive. On the other hand, a move to defined contribution plans represents a substantial increase in employee risk. I find that the cost savings from switching to a defined contribution framework appear to be very small, and suggest that to the degree public sector pensions can be cast as overly generous, this is a matter of broader governance.

Chapter 2

The Size of the Pie and Who is Eating It

2.1 Introduction

In 2006, the top 1 per cent of Americans earned over 20 per cent of all income earned; more than double the share they took home in the late 1970s (Piketty and Saez, 2006). This share has grown further since the recent financial crisis: in 2010, the top 1 percent captured 93 percent of total income growth (Saez, 2009). The question of whether inequality matters, how much it matters, and what we should do about it remains highly divisive within the general population, as evidenced by various political movements (such as Occupy for example).

In contrast, economists' traditional general equilibrium models clearly tell us that distribution does not matter. Not surprisingly then, economists' attempts to demonstrate the importance of distribution have relied on the relaxation of traditional assumptions. While incorporating imperfect capital markets has led to important breakthroughs in the theory of inequality and education, investigations into the relationship between inequality and efficiency that do not rely on capital market imperfections have been more limited.

This paper examines an alternative mechanism by which distribution may affect efficiency. We already know that removing the standard assumption of convex production sets breaks the independence between distribution and efficiency (Brown and

Heal, 1979). Or in other words, if we allow firms to have access to economies of scale then changing the degree of inequality can lead to a change in social welfare. However, we have little idea of what this might mean in practice. What kind of distributions are "better"? Under what circumstances?

While this is a complex question, it is relatively easy to show, in some specific, simplified settings, that output should increase as the distribution becomes more equal. If consumers have a taste for variety, but are not able to completely satiate it then society will face a tradeoff: increasing the income of the wealthy at the expense of the middle-class will increase the variety of goods on the market, but will decrease the utilization of economies of scale.

Specifically, consider the case where consumers have non-homothetic preferences that encourage them to consume an increasingly varied bundle of goods as their income increases. Assume further that household income is not conditional on which goods are being purchased in the economy. We show that if we allow external economies of scale on the production side, a redistribution from a few very rich to a few very poor can stimulate an increase in total output. Similarly, assuming economies of scale are neither too small nor too large, we can obtain a Pareto-optimal redistribution through leveraging a small tax on the rich to provide a transfer to the middle-class. Finally, similar increases in output can be obtained by increasing the number of consumers, rather than redistributing wealth from one consumer to another.

Our results reflect the fact that rich people may consume a lot more of something than middle-class people — but *very* rich people do not necessarily consume that much more than rich people. Despite having the income of over 753 median households, Bill Gates probably doesn't own as many shirts, iPads, TVs, or cars as all of those households put together. Instead, he owns a yacht and a private jet — goods that median households do not (usually) consume. Therefore, instead of producing many more iPads, we produce a private yacht...and fail to leverage the economies of scale on either.

This paper first outlines its place in the literature, before presenting a general equilibrium model of households who have a taste for variety that they cannot afford to completely satiate. While these households do not require a varied consumption bundle in order to obtain some utility, they gain from consuming an increasing variety of goods as their income rises.

The model includes large numbers of small firms, who are all pricing at marginal cost. However, there are gains to the firm from producing more, through innovation in production techniques, learning-by-doing, and so forth. These gains cannot be captured by the firm alone, but spill over to all firms in the industry. Consequently, as industry-wide production goes up, prices fall.

Since the model has multiple equilibria, we consider local perturbations around a Pareto-superior equilibrium. For simplicity, we define two societies, one whose Lorenz curve lies completely above the other. Within each society, there are three classes: rich, middle-class, and poor. We assume that economies of scale have been exhausted on the goods consumed by the poor, and that the goods consumed only by the rich are too small in number for economies of scale to be significant. Finally, we assume that firms have a constant capital-labor ratio for all goods, since this eliminates inter-class conflict based on income sources. Given these assumptions, we show that (a) a transfer from a small number of rich households to a small number of poor households increases output and (b) if economies of scale are neither too big nor too small, there is a tax that can be levied on the rich and distributed to the middle-class that is not only output-increasing but is also Pareto-optimal.

2.2 Literature

Under standard models of convex production, the second welfare theorem guarantees that distribution cannot impact efficiency. However, this independence breaks down once production sets are non-convex(Brown and Heal, 1979). Once firms have economies of scale, there may be transfers within the economy that will improve not only social welfare, but welfare of *all* agents.

However, because any general equilibrium with internal economies of scale is sensitive to the price rules used (Brown, 1991) our model utilizes external economies of scale (often referred to as Marshallian economies of scale). One way of considering external economies of scale is as a representation of what Allen (1983) terms "collective invention", where each firm generates technological improvements through the process of expansion, which are then passed on to other firms when they too expand. So as the industry grows, technology improves and prices fall.

Chipman (1970) first showed the existence of an equilibrium in a model with external economies of scale, which was later generalized by Suzuki (1995; 1996; 2009). Since then, others have used external economies of scale in a range of models, including trade, growth, innovation, and even environmental regulation (see Mohr 2002 for the last). Key papers include Romer (1986), who incorporated external economies of scale into a model of long-run growth, and Krugman (1979; 1991), who utilized them extensively in models of trade. In general, however, this work does not seek to use economies of scale to evaluate different wealth distributions within a single society.

Technically, the model presented here bears the closest relationship to those in the development literature that seek to model how a society moves through phases of consumption. For example, our variety-seeking preferences are very similar to those utilized by Matsuyama (2002). Our preference structure produces the same 'flying geese' pattern, in which households gradually consume additional goods as prices rise (although unlike the 'flying geese' model we do not cap total consumption). Notably, Matstuyama indicates that this pattern has been demonstrated to be very consistent across countries at different stages of development. However, while Matsuyama used this technique to model product diffusion, we use the same preferences to model the importance of inequality.

Baland and Ray (1991) likewise use non-homothetic preferences to show the importance of inequality, but consider a completely different mechanism from our own. In their case, they look at whether inequality might keep the economy at an undesirable equilibrium, where unemployment is low. Our model does not rely on labor-market frictions, but is more focused on the consumption response.

Murphy et al. (1989) combine non-homothetic preferences with external economies of scale, as is done in this paper. However, their goal is to model the process of industrialization — essentially, selecting from the multiplicity of equilibria generally available in these models (see Matsuyama (1991) for details). They require firms to have an initial monopoly, and extract initial monopoly rents, in order to encourage them to enter the market. This has the effect of rendering a perfectly equal economy incapable of industrialization. Foellmi and Zweimüller (2006) take a similar approach.

Our model contains no such requirement, in part because we do not demand that products enter at the level of the rich before trickling down to the poor. It may appear as if trickle-down must happen, because the rich consume additional products that the poor do not, but we believe it is just as plausible for a firm with a truly novel product to enter at a level where it is consumed entirely by both the rich and the middle-class, or even the poor if the product has a sufficient benefit:cost ratio. Thus we do not want to place limits on *the process* by which society chooses their consumption bundle.

2.3 Model setup

Consider a general equilibrium model where households have a taste for variety in consumption and firms have external economies of scale over some range of production.

The term 'goods' refers to classes of goods rather than specific types of goods. This is important, because we are abstracting from specific tastes. One person, for instance, may prefer to have chocolate ice-cream whereas another would prefer vanilla, but the technology to produce them is almost exactly the same. Furthermore, households will have a very high elasticity of substitution between such goods, and will be excessively sensitive to relative variations in their prices. However most of this individual substitution will cancel out at the aggregate level. Consequently, we clearly distinguish the class of ice-cream from, say, iPads and other electronic tablets. As a rough guide, consider goods to be in the same class if they (a) draw on very similar technologies or (b) perform the same practical or social functions (resulting in a high elasticity of substitution between them).

2.3.1 Households

There are $i = 1, ..., n$ households, who consume $k = 1, ..., l$ goods (where a good is denoted x_k, and forms part of a bundle x). Households have varying endowments of perfectly divisible labor (L) and capital (K), which constitute their wealth. Labor attracts a wage (w), and capital attracts rents (r), which are set in competitive factor markets and make up total household income. Households have shares in firm profits (θ_j, where j denotes the firm). Consequently, households have a budget set B defined by the parameters p, L_i, K_i, θ_i (where p denotes the price vector).

Consumption is in accordance with the limitations imposed by the budget set and the nature of the household's utility function, which is continuous, monotone, and twice-differentiable. For simplicity, consumption of all goods is weakly positive (so L and K do not enter into the utility function).

2.3.1.1 Preference for variety

Households in this economy have a taste for variety in consumption. A taste for variety in this context means that utility and consumption meet four conditions.

1. No need for variety: $u(x) > 0$ *for all* $x \neq 0$.

It would be too strong to assume that variety is a necessary condition for utility. Instead, while households *prefer* variety, they will nevertheless get positive utility from any good, or subset of goods that they consume[1]. Equivalently:

$$u(x) > 0 \Leftrightarrow \exists k \in 1, ..., l \ s.t. \ x_k > 0$$

2. Weak Variety: The utility function is strictly concave in each good.

This is a standard assumption, since it implies a quasi-concave utility function and convex preferences. Since the marginal utility of each good is declining, at some point

[1]This rules out the log-linear approximation to Cobb-Douglas preferences, where utility goes to $-\infty$ as consumption of any one good goes to zero.

households would prefer to consume a different good rather than increase the size of the bundle they are currently consuming. Consequently, this can be viewed as a desire to consume more variety *eventually* (as income gets large).

3. Variety: The utility function is separable.

By Green (1961) separability (combined with weak variety) implies that all goods are normal. Consequently, as income rises, there are no goods that the household would prefer to reduce or remove from their utility function. This condition can therefore be viewed as desire to *maintain* variety as income grows.

4. Strong variety: Piecewise linear income expansion paths.

Combined with the above, this is mathematically equivalent to:

$$\frac{1}{p_x}\frac{dU}{dx}|_x = \frac{1}{p_y}\frac{dU}{dy}|_y \Rightarrow \frac{1}{p_x}\frac{dU}{dx}|_{\alpha x} = \frac{1}{p_y}\frac{dU}{dy}|_{\alpha y}$$

For all $\alpha > 0$, for fixed p_x, p_y.

This means that for all subsets of goods where consumption of each good is strictly greater than zero, income expansion paths are linear. Consequently, while a household may consume a wider range of goods over time, the ratio of the goods we are already consuming remains the same. This is a stronger version of the variety condition — rather than just looking to maintain some variety in our consumption, as income grows we will always maintain at least the degree of variety we started with.

2.3.1.2 Household problem

Given assumptions $1 - 4$, the household problem becomes:

$$\max u(x)$$

$$\begin{aligned} s.t. \qquad & x \geq 0 \\ x \in B(p, L_i, K_i, \theta_i) & = \left\{ x \in \mathbb{R}^l : p \cdot x \leq wL_i + rK_i + \sum_j \theta_{ij} \cdot (p \cdot y_j) \right\} \end{aligned}$$

where $u : \mathbb{R}^l_+ \to \mathbb{R}$: satisfies conditions 1-5

This satisfies the regularity conditions to be solvable using the Karush–Kuhn–Tucker approach.

2.3.1.3 Example utility function

Proposition 2.1: *An example of a utility function that fits assumptions 1-5 is the log-linear approximation to Cobb Douglas with an additional unit added to each x:*

$$U(x) = \sum_{k=1}^{l} \alpha_k \log(x_k + 1)$$

where $(x_k + 1)$ has been used instead of x_i to prevent utility from going to $-\infty$ when x_k is not consumed.

PROOF: (1) Since $\log(x_k + 1) \geq 0$ for all $x_k \geq 0$, $U(x) \geq 0$ for all $x \geq 0$. (2) follows trivially from the concavity the log function. (3) follows from the additivity separability of the utility function. (5) follows trivially from the fact that the log function is continuous, C^2, and monotone, so the summed log function is also.

(4) Assume without loss of generality that good k is being consumed. Then, for any other good $\hat{k}$, either:

$$\frac{1}{p_k} \frac{\alpha_k}{x_k + 1} = \frac{1}{p_{\hat{k}}} \frac{\alpha_{\hat{k}}}{x_{\hat{k}} + 1}$$

or

$$\frac{1}{p_k}\frac{\alpha_k}{x_k+1} > \frac{1}{p_{\hat{k}}}\frac{\alpha_{\hat{k}}}{1}$$

In the former case, we know from the use of the log function that the proportions of goods consumed will be the same at all levels of income. In the latter case, good $x_{\hat{k}}$ is not currently being consumed, and so is not part of the bundle of goods being considered. ■

2.3.2 Firms

There are $j = 1, ..., m$ firms in the economy. Firms produce consumption goods in quantities $y_k = \sum_{i=1}^{n} x_{ik}$. They produce these using labor (L) and capital (K), which they acquire on a frictionless input market (so there is full employment).

Each firm operates in a competitive market, but has access to economies of scale in their production functions (this is discussed in detail below).

In addition, we make the standard assumptions that (a) firms have free disposability, and (b) periods are long enough that firms can choose to produce nothing at zero cost in any period.

2.3.2.1 Economies of scale

Firms in the model have access to external economies of scale.[2] External economies of scale result in the marginal cost of production for all firms declining as the total production of the industry grows, and so are external to any individual firm. Perhaps the best example of this is innovation. As an industry expands, production processes are often subject to incremental improvements, which over time spill over from one firm to another. In addition firms benefit from learning-by-doing at other firms (in part through ongoing movement of employees), and from the construction of supportive infrastructure by third parties (such as roads, ports, etc) that become more

[2]Note that this is different from *internal* economies of scale, which are more prevalent in the industrial organization literature.

profitable as the the industry as a whole expands.

As is standard in the external economies of scale literature, we also assume that each firm is small enough to take total industry output as a given. Firms are essentially operating in a competitive environment, and set price equal to marginal cost. Firms are assumed to have cost functions that are linear at any fixed level of industry output. In addition, we assume that there is no interaction between the production processes for different goods, so that there is no difference between one firm that produces two different goods and two firms that each produce the equivalent quantity of one good.

Finally, it is necessary to assume that the marginal cost is bounded below. This is quite reasonable, since in general manufactured goods require some labor or capital to be embodied in the good itself, so the marginal costs cannot ever go to zero.

The firm's production schedule will therefore be given by the following technology function:

$$F_j : K \times L \times \mathbb{R} \to \mathbb{R},\ y_k = F_j(K_j, L_j, \sum_{j=1}^{m} y_{jk})$$

where $F_j(K_j, L_j, \sum y_{jk})$ is homogeneous of degree one in K and L, and monotonically increasing in $\sum y_{jk}$. $F_j(K_j, L_j, \sum y_{jk})$ is bounded for any level of (K_j, L_j). That is, the marginal cost does not go to zero for any level of $\sum y_{j,k}$. Furthermore, for technical reasons, I will assume that the technology function is continuous in $K_j, L_j, \sum y_k$.

Note that it is *not* essential ex-ante that F_j be the same for each j. However, if there are variations, it is only the most efficient firms that will operate at each level of demand. It is therefore sufficient to consider an industry-wide $F\left(K_j, L_j, \sum y_{jk}\right)$ that incorporates the most efficient production of any firm at each level.

2.3.2.2 Firm problem

Consequently, the firm will face the following problem:

$$\begin{aligned} \max \pi &= p \cdot y - rK - wL \\ s.t. &\quad y \geq 0 \\ &\quad y_k = F(K, L, \sum_{j=1}^{m} y_k) \end{aligned}$$

$$\text{where } F : K \times L \times \mathbb{R} \to \mathbb{R};\ w, r, K, L \in \mathbb{R}_+;\ p, y \in \mathbb{R}^l_+,$$
$$F(\lambda K, \lambda L, \sum y_k) = \lambda F(K, L, \sum y_k) \text{ and}$$
$$\sum y_k \leq \sum y'_k \Rightarrow F(K, L, \sum y_k) \leq F(K, L, \sum y'_k)$$

Since firms are therefore maximizing a continuous profit function over a bounded set in $\mathbb{R}^n$, this problem has a unique solution.

2.4 Equilibrium

Given the above setup, we first need to establish the existence of an equilibrium. We then refine the set of equilibria, and determine some key characteristics of demand and supply in equilibrium.

2.4.1 Existence

Proposition 2.2: *A competitive general equilibrium exists.*

Proof: By Suzuki (2009), an equilibrium exists so long as:

1. Consumers have continuous, convex, locally-non-satiated preferences.
2. Firm's have free disposability, possibility of no production, and a continuous technology function.
3. Consumers satisfy the minimum income condition.
4. The set of feasible allocations is bounded.

The setup directly implies all of the above conditions are met. See appendix 2.7 for full details. ∎

2.4.2 Equilibrium characteristics of demand

We assume that households are small enough that they take prices to be fixed when making their consumption decisions. This is similar to our assumption that firms are too small to consider their prices when making their production decisions. Consequently, the household problem remains the same as outlined above, even within a general equilibrium setting.

One key characteristic of demand is that the number of different *types* of goods will increase as household incomes rise. In fact, the bundle of different types of goods consumed will be subject to a range of cutoffs. Initially, households will consume only a certain group of goods (we can think of these as food and other staples) with relatively high marginal utility relative to price. As income rises, the marginal utility of staples declines, and other goods with lower marginal-utility:price ratios will become more attractive to households. They will therefore begin to consume additional goods.

An example of the number of types of goods consumed using the utility function $U(x) = \sum_{k=1}^{l} \alpha_k \log(x_k + 1)$ and a fixed set of parameters is given in the example below. As this example demonstrates, our assumptions on the utility function not only ensure that households have a desire for variety, but also that they are able to better realize that desire as their income increases.

Example 2.1: *2.1is a graph of a household's optimal consumption given this utility function and various levels of income.*

Figure 2.1: Number of different types of goods consumed as income increases*

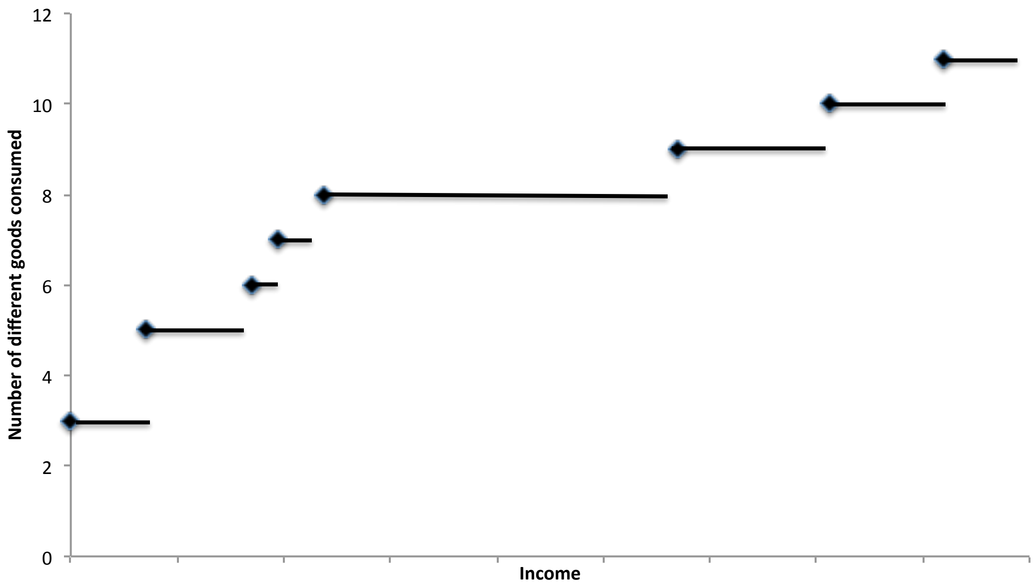

* αs range from .02 to .18, and prices are randomly generated. Initially, the consumer's income is divided between 2 goods. Then, once it reaches a threshold of around 20, the consumer starts dividing their wealth between 5 goods. Once wealth reaches 25 the consumer consumes a total of 6 goods, above 26 7 goods, and 29 8 goods and so on.

This result holds in general, as outlined in the two claims below.

Proposition 2.3: *Households will consume weakly more goods as their income grows.*

Proof: See appendix 2.7.

Proposition 2.4: *Assuming conditions on the utility function hold, either (a) when the household consumes nothing, the ratio of marginal utilities for any two goods is equal the ratio of their prices, or (b) the number of goods consumed strictly increases over some range of income.*

Proof: See appendix 2.7. ■

In addition to the number of *types* of goods consumed, it is also necessary to consider how demand for *each good* varies as income increases. Demand changes in wealth for a particular good can be represented by a series of upwardly sloping linear segments, each of which has a lower slope than the previous one. Each change in the slope coincides with one or more new goods being added to the bundle.

Proposition 2.5: *Household demand for each good is locally linear in income for all except a finite number of points, and globally concave.*

PROOF: See appendix 2.7. ∎

Because demand for each good is increasing linearly within the fixed bundle of goods consumed, it is clear that any households that are consuming the same bundle of goods — regardless of household income — can be represented by a single representative consumer. This consumer, however, is *not* the representative consumer for a household that consumes a different bundle of goods. Consequently, there is not one representative consumer for the entire economy, but rather a finite set of representative consumers, as outlined below.

Proposition 2.6: *Given fixed prices, any continuous distribution of income can be reduced to a finite number of representative consumers.*

PROOF: Since the income expansion paths are linear within subsets of goods, and all households have the same preferences, all households who would consume a set number of goods have a single representative consumer.

Since there are finite goods, there are also finite subsets of goods consumed given a set of prices, ($\leq l$ to be precise) and therefore a finite number of representative consumers. ∎

2.4.3 Characteristics of production

In order to consider how equilibrium changes as we adjust the distribution, we need to have some idea as to how prices and profits change with demand. Firstly, it is worth noting that profits at all stages are zero, since production is constant in labor and capital. This simplifies the analysis considerably, since we can avoid considering any inter-class conflict that arises from the desire of factory owners to stimulate consumption of their own goods in order to elevate profits, at the expense of the consumption of goods that consumers may prefer.

Secondly, and unsurprisingly given the existence of economies of scale, if demand for any one good exogenously rises, holding wages and rents constant there will be a fall in the price of the good, as the industry generates more economies of scale.

Proposition 2.7: *Holding wages and rents constant, if a firm has external economies of scale over the necessary range, then price will be decreasing in demand.*

PROOF: See appendix 2.7. ■

2.4.4 Equilibrium characteristics and refinement

While an equilibrium exists, it will generally not be unique because production technologies are not convex. While we can do local comparative statics even if there are multiple equilibria, we wish to consider a degree of refinement so that we can determine what, if anything, we can say about the global result.

Firstly, it is worth noting that equilibria in this model are subject to coordination difficulties. Even if all households have the same level of income and the same tastes[3] each individual household may be small relative to the economy. So, if all households have coordinated on consuming good A, they will leverage economies of scale on that good, and in consequence its price may be sufficiently lower than the price of good B that no household wants to switch. However, it could also be that everyone is better off if they all switch to consuming good B.

Example 2.2: *Take 3 consumers, each with wealth $\omega = 10$. Then, let the marginal utility schedule for each good is as outlined in Table 1, and the price schedule for the production of each good (in per unit terms) be as laid out in Table 2.*

[3]If all consumers have the same wealth and tastes, the whole economy can be represented by a single agent. That agent faces a budget set that is a bounded (due to finite number of goods and bounded technology constraint) set of $\mathbb{R}^n_+$ and has a continuous utility function and so should have an optimum (i.e., unique equilibrium).

Table 2.1: Marginal Utility Schedule

units	x_1	x_2	x_3
1	3	2	3
2	2	1	2
3	1	.5	1

Table 2.2: Price Schedule

units	x_1	x_2	x_3
1	5	5	6
2	5	5	6
3	4	4	4
4	4	4	4
5	3	3	3
6	3	3	3
$\vdots$	$\vdots$	$\vdots$	$\vdots$

Then, it is an equilibrium for everyone to consume two units of good 1 and one unit of good two for a total utility of 7, and a total cost of 10. A quick inspection reveals that no one would prefer to deviate from this. However, there is an alternative equilibrium, under which everyone consumes 2 units of good 1 and 1 unit of good 3. This gives a total utility of 8 at the same cost (10) and the same level of overall economic productivity.

Therefore, as our first refinement: *we consider only those equilibria where households with the same income acting collectively as a class would not prefer to deviate.* [4]

It is important to note that this refinement does not require that individual households take into account their effects on price. Each household is too small to take into

[4] It is worth noting that low levels of inequality may be beneficial, as they can help to resolve the coordination problem associated with equilibrium selection. If all members of a group have the same wealth, there may exist an equilibrium where all consume the same good, but would be better off switching to another good with higher utility per unit and larger economies of scale.

account the effects of their own consumption on price. Rather, what we are saying is that those households with the same income cannot be 'duped' into consuming a different bundle when another stable equilibrium exists where they consume a bundle that they prefer.

However, even if we allow coordination between households with the same income, we may still encounter coordination issues *across* income levels. That is, in scenarios where all households prefer the same equilibrium, a lack of coordination between them may result in a Pareto-inferior equilibrium.

Example 2.3: *Take 2 households, one with wealth* $\omega = 8$*, and one with* $\omega = 16$*. Then, let the marginal utility schedule for each good be as outlined in Table 3, and the price schedule for the production of each good (in per unit terms) be as laid out in Table 4. Then there is one equilibrium where everyone consumes only* x_2*, and one where everyone consumes only* x_1*. Clearly the equilibrium where everyone consumes* x_1 *is preferable to both households.*

Table 2.3: Marginal Utility Schedule

units	x_1	x_2
1	14	13
2	13	12
3	12	11
4	11	10

Table 2.4: Production Schedule

units	x_1	x_2
1	10	10
2	9	9
3	8	8
4	7	7
5	7	7
$\vdots$	$\vdots$	$\vdots$

Therefore, as our second refinement: *we will consider only those equilibria that are not Pareto-dominated.*

However, we may still be concerned about a scenario where a household prefers one equilibrium, whereas poorer and/or richer households strictly prefer a different equilibrium. This can arise for many reasons. For instance, we may have inter-class conflict if goods have different capital/labor ratios. For instance, if the rich have a higher capital to labor ratio than the poor, they would prefer that the poor consume those goods which are more capital-intensive. Since we are primarily interested in the interaction between consumption, economies of scale, and inequality, this adds an unnecessary complication to the model. We would therefore assume that capital/labor ratios are independent of the consumption decision. The simplest (and certainly sufficient) method of removing this conflict is to *assume that the capital/labor ratio is constant for all goods at all periods of time.*

A constant capital/labor ratio means that a class's income will be the same in all equilibria up to a constant α which is the same for all classes. Furthermore, all prices will scale by the same factor as income. We can therefore create a scaled price vector for each equilibrium, and show that each equilibrium must have a unique (scaled) price vector.

Proposition 2.8: *Assume that the capital/labor ratio is constant for all goods. Then for any household in any two equilibria* 0 *and* 1*,* $\omega^0 = \alpha\omega^1$*, where* α *is common to all households. Further* $\frac{p^1}{\alpha}.x^1 = \omega^0$ *and* $p^0 \neq \frac{1}{\alpha}p^1$*.*

Proof: See appendix 2.7. ∎

Eliminating inter-class conflict over returns still does not ensure that we have a unique equilibrium. For example, we might have a situation where all goods have the same production schedule, but where households have different levels of income. Since goods are ranked by marginal utility we may still encounter differences in the ideal equilibria for households with different levels of income. Example 4 outlines one such case.

Example 2.4: *Take 3 households with wealth* $\omega = 7$ *and one individual with* $\omega = 31$*. Then, let the marginal utility schedule for each good be as outlined in Table 5, and* *the price schedule for the production of each good (in per unit terms) be the same as in example 3.*

Table 2.5: Marginal Utility Schedule

units	x_1	x_2
1	3	2
2	2	1.8
3	1	1.6
4	.5	1.4

Now, left to their own devices, the three households with less wealth would prefer to coordinate on all consuming good x_1 for a total utility of 3 each. In that case, if the wealthy household enters later, they will consume 3 units of good x_1, and one unit of good x_2 for a total utility of 8.

However, the wealthy household would prefer that all of the other households coordinated on consuming good x_2 (for a utility of 2 each), and themselves consume 3 units of good x_2 and one unit of good x_1, for a total utility of 8.4. Note that in both equilibria the productive capacity is the same, so we can take income as given[5]*.*

The reason for the above result is that the marginal utility curves — while both concave — cross, between 1 unit and 2 units. Consequently, at low levels of consumption one prefers to consume good A, while at high levels of consumption one prefers to consume good B. In a standard model this is not a concern, as one simply consumes good A until the marginal utility has declined sufficiently, and then consumes good B. However, once economies of scale are involved, the price of the good is positively tied to its consumption, so this strategy is no longer possible.

Proposition 2.9: *There is no crossing of marginal utilities. That is, for all $\alpha > 0$ $\frac{\delta U}{\delta x}|_{x=\hat{x}} > \frac{\delta U}{\delta y}|_{y=\hat{x}} \Rightarrow \frac{\delta U}{\delta x}|_{x=\alpha\hat{x}} > \frac{\delta U}{\delta y}|_{y=\alpha\hat{x}}$.*

PROOF: By strong variety assumption. See appendix 2.7. ■

[5]Assume for now that the marginal rate of technical substitution is the same at each level of production for both goods.

However, even this is not enough to ensure that all equilibria are Pareto-ranked. Even when marginal utilities do not cross, it is possible to encounter the same problem. Example 5 demonstrates such a scenario. If we allow marginal costs to move too closely together, or to cross, we will have even more potential for multiple equilibria that are not Pareto-dominated.

Example 2.5: *Take 3 households, one with wealth $\omega = 7$, one with $\omega = 15$ and one with $\omega = 30$. Then, let the marginal utility schedule for each good be as outlined in Table 6, and the price schedule for the production of each good (in per unit terms) be as outlined in table 7.*

Table 2.6: Marginal Utility Schedule

units	x_1	x_2
1	14	13
2	12	11
3	10	9
4	8	7

Table 2.7: Production Schedule

units	x_1	x_2
<3	20	20
3	8	8
4	8	8
5	8	8
6	8	8
>6	7	7

Clearly, the poorest individual prefers an equilibrium where everyone consumes only good x_1, since their utility is higher in that case. On the other hand, the middle-class household prefers to consume 1 unit of good x_1 and one unit of good x_2 to consuming 2 of x_1, but will not deviate on their own. Finally, the richest household wants to consume 2 of good 1 *and* 2 *of good 2, rather than $\approx 4\frac{1}{3}$ of good 1, but again they cannot deviate themselves because of the high initial cost. The two wealthier* *households can coordinate on an equilibrium that improves their outcomes, but this will be a the expense of the poorest household.*

In order to eradicate this, we make an additional assumption — that *all technologies are the same.* This is restrictive, but slightly less restrictive than it sounds, since we can manipulate the size of a unit of x (we do not need to compare, say, a gallon of milk to a t-shirt, it may be a quart of milk or a pack of t-shirts). Of course we still need to allow that the economies of scale on these goods move in essentially the same direction at essentially the same rate, but this is not significantly more restrictive than the assumptions we were already making by putting the goods into classes in the first place, and is illuminating. We should also point out that this assumption is a sufficient one — one could easily establish a range of cases where we get the same outcome without making this assumption at all.

We have already established that we can order the classes of goods based on their marginal utilities. With the additional assumption about production technologies we know that any equilibrium that does not have higher-ranked goods being consumed in greater volumes must be Pareto-dominated.

Proposition 2.10: *If there are two goods x and y s.t. $\frac{\delta U}{\delta x}|_0 > \frac{\delta U}{\delta y}|_0$, the capital/labor ratio is constant for all goods and the production technology is the same, then any equilibrium where a household does not consume $x > y$ must be Pareto dominated by another equilibrium.*

Proof: See appendix 2.7. ■

This leaves us with a subset of Pareto-superior equilibria from which to choose. Since we are looking to show that it is possible to undertake a Pareto-superior redistribution, we would like to choose the equilibrium that makes this as difficult as possible. Therefore, we want to choose the equilibrium that is best for the group we are redistributing *from* (in our case the richest household(s)). That way, the improvement they receive must be from the redistribution process, and not just from switching to an equilibrium they like better. We therefore make the final refinement: *the equilibrium must be the most preferred for the wealthiest households.*

In most cases, the equilibrium is unique. This follows from the fact that if the rich are indifferent between two equilibria, all other groups must be indifferent or prefer the equilibrium with more consumption of the higher-ranked good. This may break down if there are not as many classes as goods, but this is not interesting as we must either have a very equal (or at minimum very structured) society or a very small number of households. Alternatively, if three or more goods have untapped economies of scale then we could have an 'ends against the middle' scenario. We rule the three-or-more good case out also since in the next section we are concerned with a three-good scenario where goods consumed by the poor are consumed in such large quantities that economies of scale are exhausted.[6]

Proposition 2.11: *Assume the capital/labor ratio is constant for all goods, the production technology is the same and neither equilibrium is Pareto-dominated. Assume there is sufficient inequality between households that there are at least as many classes as goods in all equilibria. Finally, assume that there are only two goods whose economies of scale are not exhausted at all equilibria. Then if there is more than one equilibrium, there is a unique equilibrium that is preferred by the wealthiest household(s).*

PROOF: See appendix 2.7. ■

This is the equilibrium we will consider when we consider the comparative statics of redistribution in the following section.

2.5 Comparative statics and wealth distribution

Having established the existence and some key characteristics of an equilibrium, we can now return to our original question: if we change the distribution of income

[6]An alternative would be to continuously break ties in favor of the next-wealthiest household, since the poor in our subsequent comparative statics have no real sway over the equilibrium selection.

within a society, how does our equilibrium change? Will a move to a more equal society increase or decrease social welfare? Will it increase or decrease output? And can it, under some circumstances, be better for everyone?

From the sections above, it is clear that obtaining general comparative statics for this model is extremely complicated. A general solution would need to take into account what happens at every equilibrium, and would need to apply to all possible wealth distributions in an economy.

A general solution, however, is empirically unnecessary. Most societies have a skewed bell-shaped distribution of households, with large numbers of "poor"[7] with few resources, scaling up to a very few households that are very rich, who often own a very large proportion of the economy's capital.

We will therefore divide our economy into three classes of household: the poor, the middle-class, and the rich (subscripted by r, p, and m, respectively). Households in each class are identical, but levels of labor and capital endowment vary across classes. We note that even at the end of the boom in 2007, the bottom 60% of the population held only 4.2% of the total assets (Wolff, 2010). Therefore, for the purposes of the model we assume a large fraction of the population (n_p) are poor, and the poor households have only a labor endowment, and but no capital endowment ($K_p = 0$). At the same time, the top 5% of society held over 60% of total assets in 2007 (Wolff, 2010), so we assume that the rich are a very small fraction of the population (n_r) who have a very large capital endowment (K_r). The middle-class are those that remain — smaller in number than the poor, larger in number than the rich ($n_p > n_m > n_r$) — and have some degree of capital endowment (K_m). We abstract from human capital, and assume that the labor endowment (L) is the same for all households.

This leads us to the next question: if we are going to compare this society with a more equal one, what do we mean by more equal? Here we consider two possibilities: an 'intervention' model (section 2.5.1) and a taxation model (section 2.5.2). In the first case, we will select some small subset of rich households at random, and

[7] Clearly our definition of poor varies with the economy — we do not claim that the poor in sub-saharan Africa are directly comparable with the poor in the United States

reallocate the difference between their income and a middle-class income to the number of poor households, such that the reallocation is sufficient to convert those poor households into middle-class households. In the taxation case, we leverage a tax τ on rich households, from which we generate a transfer to middle-class households. In both cases, average income remains the same, as does the total number of households (although we will relax the latter in section 2.5.3).

In the first case, the change results in a Lorenz curve that lies weakly above the Lorenz curve for our original society. Thus, the societies can be strictly ordered in terms of inequality (Sen, 1973) (it also implies that the Gini coefficient is lower for our more equal society).[8] The second case does not result in a change in wealth, but results in income streams to each class that would be consistent with a more equal society.

For simplicity, in our societies we will have 3 goods: y, x, z. Consistent with section 2.4.1, prior to any intervention or taxation taking place, the poor consume good y only, the middle-class consume goods y, x, and the wealthy consume goods y, x, z. Also consistent with section 2.4.1, we note that our households are too numerous and too small to consider their own effects on the price of the goods they are consuming.

2.5.1 Output and 'intervention'

For the purposes of this section, we will consider the case where we have a change in the number of people in each class. We call this the 'intervention' model. The change in equality results from some subsection of the poor being re-allocated wealth from some subsection of the rich, converting both groups involved into middle-class households. We will denote the economy prior to the intervention as economy 1, and the economy post-intervention as economy 0. We will, however, note that our proposed 'intervention' is very small in scale, since big shifts may propel us toward alternative equilibria.

[8] This is equivalent to saying that if we order both societies from richest to poorest, and choose any person from the unequal society, then those who are richer than that person will hold more cumulative wealth than the same ranked person in the equal society (unless they are the poorest person, in which case the total capital will be the same under both societies)

Mathematically, this means that $L^1 = L^0$, $K_r^1 = K_r^0$, $K_m^1 = K_m^0$. By definition, we will also have that the proportion of households who are wealthy (n_r) and poor (n_p) are greater in the unequal society than the equal society, and the proportion of middle-class households (n_m) is greater in the equal society than in the unequal society.

The first step for determining the effect of this change is to consider the first-order effects. Specifically, if our intervention occurred without having any impact on the prices of goods, households would (unsurprisingly) increase their collective consumption of both poor and middle-class goods. This is simply a function of the fact that we are re-allocating income away from households who would spend some fraction of it on rich goods toward households who spend it only on poor and middle-class goods.

Proposition 2.12: *As a first-order effect, there will be more consumption of poor and middle-class goods in the equal society than under the unequal society, at the equal society's prices.*

PROOF: See appendix 2.7. ∎

Given this information, we can show that for everyone who is not in the households which have been displaced, it is possible to achieve an increase in total consumption. This is because the above results extend easily into the case where we have economies of scale, so there is a sharp increase in the consumption of middle-class goods. Since middle-class goods are where all of the gains from economies of scale are taking place, both the middle-class and the wealthy gain from facing cheaper prices. The key point here is that economies of scale cannot be too large (or, alternatively, the middle-class good cannot be too desirable relative to the poor good), as this would attract the poor away from the poor good in large enough numbers so as to raise the price of that good.

Proposition 2.13: *Assume that the optimal labor/capital ratio is the same for all goods, at all levels of production. Assume also there are sufficient numbers of poor* *households to exhaust economies of scale on poor goods. Assume that production quantities are too small for economies of scale to exist on rich goods. Then, if the intervention is small enough and economies of scale are not too large, every member of the equal society (excluding those who have lost wealth) will consume weakly more than the equivalent member of the unequal society (at the equal society's prices), and will therefore be better off than their counterparts in the unequal society.*

PROOF: See appendix 2.7. ■

Finally, we show that total output is higher in the equal society than in the unequal society. This results from the fact that, at the old prices, the transfer is output-neutral. However, the change in the prices is overwhelmingly negative, so society can afford to purchase more at the new prices than at the old prices.

Proposition 2.14: *Total economic output is higher at equilibrium 0 than at equilibrium 1, at the equal economy's prices.*

PROOF: See appendix 2.7. ■

Redistribution therefore increases productivity (and therefore our economy's version of GDP).

2.5.2 Output and taxation

In the taxation case, we will consider a transfer from the rich, rather than a change in the number of rich people. Mathematically, this means that $\omega_r^0 < \omega_r^1$, and $\omega_m^0 > \omega_r^1$. Since giving additional income to the poor will not increase their consumption of goods other than poor goods unless the transfer is large enough, for the purposes of this section we consider only a transfer to the middle-class. Class conflict in this case is between the middle-class and the rich, rather than between the rich and the poor. While it may seem odd that conflict is not with the poor, in fact a great deal of the conflict regarding who should bear the costs of public goods tends to be between the rich and the middle-class. We can view the redistribution here as simply a shift of that burden.

We also have that the proportion of households who are wealthy (n_r), middle-class (n_m), and poor (n_p) are equal across both societies. We will fix the transfer rate so that each wealthy consumer pays τ. Therefore, each middle-class household receives a transfer of $\frac{n_r}{n_m}\tau$.

Once again, our first step is to consider the first-order effects of the reallocation. As in the previous section, the fact that we are shifting resources from groups who purchase rich, poor, and middle-class goods to consumers of only middle-class goods raises their relative consumption.

Proposition 2.15: *As a first-order effect, there will be more consumption of middle-class and poor goods under the equal society than the unequal society at the equal society's prices.*

Proof: See appendix 2.7. ■

Given this information, what we want to know is whether we can achieve a welfare improvement for everyone for some τ. If so, then we can show that there is some Pareto-optimal reallocation.

It turns out that this is in fact possible for some τ. This reflects a certain degree of intra-class conflict. Since there are economies of scale on middle-class goods, every rich person would prefer all of the other rich people to give their money to the middle-class. However, they do not individually wish to make the same transfer.

Proposition 2.16: *Assume that the optimal labor/capital ratio is the same for all goods, at all levels of production. Assume that there are enough poor that if they choose to consume only poor goods, they can exhaust all economies of scale on poor goods. Assume also that production quantities are too small for economies of scale to exist on rich goods. Then, if economies of scale are not too small on middle-class goods, then there is some $\tau > 0$ transfer from the rich to the middle-class such that we can achieve a Pareto-improvement.*

PROOF: See appendix 2.7. ■

Corollary 2.1: *Given the framework in proposition 2.16 there is a limit to redistribution beyond which no additional transfer is Pareto-improving.*

PROOF: See appendix 2.7. ■

Corollary 2.2: *Given the framework in proposition 2.16, even if economies of scale are not large enough locally to produce a small Pareto-improving transfer, there may be a large Pareto-improving transfer, so long as the threshold in corollary 2.1 is not crossed.*

2.5.3 Welfare, distribution, and size

One of the key drivers of the claims above is that the change in distribution increases demand for certain products, which in turn drives down their price. It will therefore be obvious to the careful reader that economy size is just as important as distribution in determining, not only *total* output but output *per capita.*

Proposition 2.17: *Holding the distribution constant, a small proportional increase in the number of households will increase the amount of output per capita.*

PROOF: As the number of rich and middle-class households increases, then demand for middle-class goods must also increase (at current prices). Since there are still economies of scale to be leveraged on these goods, their price must fall. Consequently, the price of the same bundle is cheaper for all existing consumers, while their incomes are constant. They can therefore consume additional goods, and output must rise. ■

Given this impact on output per capita, it becomes immediately clear that an unequal economy that is sufficiently large relative to its more equal counterpart may actually have a larger output per capita. This is because the increase in consumption of middle-class goods from adding more households can offset the losses from making households more or less wealthy.

Note that this does *not* mean that it is more efficient for the larger economy to be unequal. It is still possible (in line with the proofs above) to make the more unequal economy produce even more output per capita by making that economy more equal also.

Proposition 2.18: *Any losses from moving toward a more unequal society can be overcome if they are matched by a sufficient increase in the* size *of the economy.*

PROOF: See appendix 2.7. ∎

2.6 Conclusion and future work

This paper examines some important circumstances under which the income distribution may affect output and efficiency. Specifically, we allow poor, rich, and middle-class households to have a taste for variety over three goods, and allow many small firms to leverage external economies of scale. Economies of scale can be leveraged only once consumption is reasonably large, and become exhausted once consumption reaches a very high level. Then, given our assumption of a constant capital-labor ratio, we achieve output gains both by redistributing the entire wealth of the rich to the poor, and from a smaller tax applied to the rich and transferred to the middle-class. Furthermore, if economies of scale are neither too large nor too small, this tax can be Pareto-optimal if set at the right level.

These results are obviously subject to the range of assumptions made. For instance, it would be useful to weaken the assumption of a constant capital/labor ratio, as this is likely to increase inter-class conflict. Allowing multiple goods whose technologies vary to be consumed by each class may offer avenues for modeling large economic shifts such as industrialization.

Finally, it would be ideal to further consider ways in which to empirically test these results. Since the model is output-focused, it may be possible to compare inter-temporal or inter-state output. Furthermore, the results on the role of economic size are very important, since we demonstrate that it is not possible to directly compare different-sized economies even with the same Gini coefficient. This may open up additional avenues for empirical consideration.

2.7 Appendix

Proposition 2.2 A competitive general equilibrium exists.

PROOF: Additional notation: $\omega_i = wL_i + rK_i$

By Suzuki (2009), an equilibrium exists so long as:

1. Consumers have continuous, convex, locally-non-satiated preferences.

2. Firms have free disposability, possibility of no production, and a continuous technology function.

3. Consumers satisfy the minimum income condition ($p\omega_i > \inf pX_i = \{px|x \in X_i\}, i = 1, ..., n$)

4. The set of feasible allocations is bounded.

(1) follows directly from the set of assumptions on the utility function. (2) follows directly from the assumptions on the technology function. This leaves (3) and (4).

Since consumers always have the option to consume nothing, for any price vector p, $0 \in X_i$, where X_i is the set of possible consumption vectors. Therefore, we must have that for every price vector $p \geq 0$ with $p \neq 0$, $\inf \{p \cdot x|x \in X_i\} = 0$. Therefore, to satisfy the minimum income condition, we need to show that for all ω_i, $p \cdot \omega_i > 0$.

Since firms have linear production functions for any given $\sum y_{kj}$, we have that firms do not ever make a profit or a loss. So $\sum_{j=1}^{m} \theta_{ij} p_k y_{jk} = 0$ for all i and for all k. Therefore, since all individuals have a labor allocation, and so long as wages are greater than zero (true, always, since marginal productivity is always >0), we must have that the $p\omega_i > 0$ for all i. Therefore, we have that for every price vector $p \geq 0$ with $p \neq 0$, $p\omega_i > \inf pX_i = \{px|x \in X_i\}, i = 1, ..., n$, so the minimum income condition is satisfied.

(4) follows from the fact that $F_j\left(K, L, \sum_{j=1}^{m} y_j\right)$ is bounded (by assumption) for any (K, L). Since $\sum K, \sum L$ are also bounded (by the endowment) we must have that:

$$\mathcal{F}=\left\{((x_i),(y_j,k_j,l_j))\in\prod_{i=1}^{n}X_i\times\mathbb{R}_+^4|\sum_{i=1}^{n}x_i\leq\sum_{j=1}^{m}(y_j-K_j-L_j)+\sum_{i=1}^{n}(K_i+L_i),y_j\leq F_j\left(K_j,L_j,\sum_{j=1}^{m}y_j\right)\right\}$$

is also bounded. ■

Proposition 2.3 Given fixed prices, households will consume a weakly greater variety of goods as their income grows.

PROOF: Let $\omega = wL + rK$, where L and K are the household's labor and capital endowments. Randomly select ω and ω' where $\omega' > \omega$. Fix an arbitrary price vector p.

By continuity of the utility function (and since perfectly divisible goods allow for a compact budget set), we know that there is a unique $x(p,\omega)$ and a unique $x(p,\omega')$ which optimize the household's utility at ω and ω', respectively.

We know that $x(p,\omega)$ is the solution to the problem:

$$\begin{aligned}\max_x &= u(x)\\ s.t.\quad p.x &\leq \omega\\ x_k &\geq 0 \quad \forall k\end{aligned}$$

Therefore, by differentiability of $u(x)$, we know that for all $k = 1, ..., l$:

$$\begin{aligned}\frac{du(x)}{dx_k}-\lambda p_k &= 0\\ x_k &> 0\end{aligned}$$

or

$$
\begin{aligned}
\frac{du(x)}{dx_k} - \lambda p_k &< 0 \\
x_k &= 0
\end{aligned}
$$

Therefore, since λ is constant, we have that:

$$
\frac{1}{p_k}\frac{du(x)}{dx_k}|_{x_k(p,\omega)} = \frac{1}{p_{\hat{k}}}\frac{du(x)}{dx_{\hat{k}}}|_{x_{\hat{k}}(p,\omega)} \tag{2.1}
$$

For all $k, \hat{k}$ with $x_k > 0$, $x_{\hat{k}} > 0$

Similarly, we know that $x(p, \omega')$ is the solution to the problem:

$$
\begin{aligned}
\max_x \quad &= \quad u(x) \\
s.t. \quad p.x &\leq \omega' \\
x_k &\geq 0 \quad \forall k
\end{aligned}
$$

Which gives exactly the same first order conditions as before. However, if good $\hat{k}$ is still consumed when the household has $\omega' > \omega$ and good k is not, then:

$$
\frac{1}{p_k}\frac{du(x)}{dx_k}|_{x_k=0} < \frac{1}{p_{\hat{k}}}\frac{du(x)}{dx_{\hat{k}}}|_{x_{\hat{k}}(p,\omega')} \tag{2.2}
$$

We also know that $\frac{du(x)}{dx_k} > 0$ for all $x_k \geq 0$. Since p is invariant to changes in ω, 2.1 and 2.2 imply that:

$$
\frac{1}{p_k}\frac{du(x)}{dx_k}|_{x_k=0} - \frac{1}{p_k}\frac{du(x)}{dx_k}|_{x_k(p,\omega)} < \frac{1}{p_{\hat{k}}}\frac{du(x)}{dx_{\hat{k}}}|_{x_{\hat{k}}(p,\omega')} - \frac{1}{p_{\hat{k}}}\frac{du(x)}{dx_{\hat{k}}}|_{x_{\hat{k}}(p,\omega)}
$$

Now, by strict concavity of u, we know that if $x_{\hat{k}}(p, \omega') > x_{\hat{k}}(p, \omega)$, we must have:

$$\begin{aligned}
\frac{1}{p_{\hat{k}}}\frac{du\left(x\right)}{dx_{\hat{k}}}|_{x_{\hat{k}}(p,\omega')} - \frac{1}{p_{\hat{k}}}\frac{du\left(x\right)}{dx_{\hat{k}}}|_{x_{\hat{k}}(p,\omega)} &< 0 \\
\Rightarrow \frac{1}{p_k}\frac{du\left(x\right)}{dx_k}|_{x_k=0} - \frac{1}{p_k}\frac{du\left(x\right)}{dx_k}|_{x_k(p,\omega)} &< 0 \\
\Rightarrow x_k\left(p,\omega\right) &< 0
\end{aligned}$$

So we have a contradiction.

In that case, the household cannot replace one good in their bundle by consuming more of another good in their bundle. In order to consume less variety of goods as price goes up, therefore, they must consume at least one good $x_{\hat{k}}$ at ω' that they don't consume at ω, and be consuming more than one good x_k at ω that they don't consume at ω'.

But then, since the household doesn't consume x_k at ω', we must have that:

$$\frac{1}{p_k}\frac{du(x)}{dx_k}|_{x_k=0} \leq \frac{1}{p_{\hat{k}}}\frac{du\left(x\right)}{dx_{\hat{k}}}|_{x_{\hat{k}}(p,\omega')}$$

However, since they consume x_k and not $x_{\hat{k}}$ at ω:

$$\begin{aligned}
\frac{1}{p_{\bar{k}}}\frac{du\left(x\right)}{dx_{\bar{k}}}|_{x_{\bar{k}}=0} &\leq \frac{1}{p_k}\frac{du(x)}{dx_k}|_{x_k(p,\omega)} \\
&< \frac{1}{p_k}\frac{du(x)}{dx_k}|_{x_k=0} \leq \frac{1}{p_{\hat{k}}}\frac{du\left(x\right)}{dx_{\hat{k}}}|_{x_{\hat{k}}(p,\omega')}
\end{aligned}$$

So we have a contradiction of concavity. Therefore, we must be consuming x_k at ω' if we consume it at ω, and we have a contradiction. ■

Proposition 2.4 Assuming conditions on the utility function hold and that prices are fixed, either (a) when the household consumes nothing, the ratio of marginal utilities for any two goods is equal the ratio of their prices, or (b) the variety of goods consumed strictly increases over some range of income.

PROOF: Given claim 4, it is sufficient to show that if (a) does not hold, there is some good that would be consumed at some ω, but would be consumed at some $\omega' > \omega$.

If (a) holds, then for all goods k, $\hat{k}$:

$$\frac{du(x)}{dx_k}|_{x_k=0} \cdot \left(\frac{du(x)}{dx_{\hat{k}}}|_{x_{\hat{k}}=0}\right)^{-1} = \frac{p_k}{p_{\hat{k}}}$$

Otherwise, there exists at least one pair of goods $x_k, x_{\hat{k}}$ s.t.

$$\frac{1}{p_k}\frac{du\,(x)}{dx_k}|_{x_k=0} \neq \frac{1}{p_{\hat{k}}}\frac{du\,(x)}{dx_{\hat{k}}}|_{x_{\hat{k}}=0}$$

In that case, there must be an ϵ such that

$$\frac{1}{p_{\hat{k}}}\frac{du\,(x)}{dx_{\hat{k}}}|_{x_{\hat{k}}=0} = \frac{1}{p_k}\frac{du\,(x)}{dx_k}|_{x_k=0} + \epsilon$$

Therefore, by continuity, there exists a $\delta > 0$ such that:

$$\frac{1}{p_{\hat{k}}}\frac{du\,(x)}{dx_{\hat{k}}}|_{x_{\hat{k}}=\delta} > \frac{1}{p_k}\frac{du\,(x)}{dx_k}|_{x_k=0}$$

■

Then, given condition 1 on the utility function, so long as $\omega < p_{\hat{k}}\delta$, the household would prefer to spend all of its income on $x_{\hat{k}}$ rather than spend any on x_k.

Proposition 2.5 Household consumption of each good is locally linear in income for all except a finite number of points, and globally concave.

PROOF: By assumption 4 on the utility function, household demand is linear each good, so long as the subset of goods consumed is fixed. By *Claim 3*, every good that is consumed at at given wealth level $\bar{\omega}$ and a fixed set of prices p, will be consumed at p for all $\omega > \bar{\omega}$.

Consequently, all goods consumed at a given level of wealth, ω will be consumed on $[\omega, \omega + \delta]$ for all $\delta > 0$. Furthermore, since the number of goods is finite, there must be an $\epsilon_1 > 0$ such that only the goods consumed at ω are consumed at $\omega + \epsilon_1$.

Furthermore, unless a new good is added at ω, there will be some for some $\epsilon_2 > 0$, s.t. the same number of goods is consumed at $\omega - \epsilon_2$ as at ω. Therefore, so long as a new good isn't added at ω, we can take $\epsilon = \min\{\epsilon_1, \epsilon_2\}$, and demand will be linear on the interval $[\omega - \epsilon, \omega + \epsilon]$.

Now, we have restricted non-linearity to the set of ω at which a new good is added, i.e., where a good x is consumed that was not consumed at $\omega - \epsilon$ for all $\epsilon > 0$. However, since the number of goods is finite, there must be a finite number of such points.

Therefore, demand is locally linear except at a finite number of points.

To show that demand is globally weakly concave, it is sufficient to show that for any good x_k, $\frac{\delta x_k^*}{\delta \omega}$ does not increase when another good is added.

Choose any good x, and any other good $\hat{x}$ that we are consuming. Assume that when we add a new good $\tilde{x}$, our consumption of x increases to αx. Then, by separability and our assumption on linear wealth expansion paths, we must have that:

$$\frac{1}{p_x}\frac{dU}{dx}|_{\alpha x} = \frac{1}{p_{\hat{x}}}\frac{dU}{dx}|_{\alpha \hat{x}}$$

Therefore, we must also increase our consumption of $\hat{x}$ by α. Since this holds for all of the goods we are consuming, we must have that consumption of all goods increases by α. Since we are adding another good, however, we must have that $\frac{\delta x_k^*}{\delta \omega}$ of some goods is declining, or we violate the budget constraint. Therefore $\frac{\delta x_k^*}{\delta \omega}$ declines as another good is added, and consumption is globally weakly concave. ■

Proposition 2.7 Holding wages and rents constant, if a firm has external economies of scale over the necessary range for a given good, then the price of that good will be decreasing in demand for that good.

PROOF: In equilibrium, $x(p^*) = \sum_{i=1}^{n} x_i(p^*) = \sum_{j=1}^{m} y_j(p^*)$.

Now, if demand increases, $x(p)$ increases to $x'(p)$ with $x'(p) > x(p)$ for all p. But then $\sum_{i=1}^{n} x_i'(p^*) > \sum_{j=1}^{m} y_j(p^*)$.

We know that $y_j(p)$ satisfies following condition:

$$p^* = MC\left(\sum_{j=1}^{m} y_j\left(p^*\right), K(y_j\left(p^*\right)), L(y_j\left(p^*\right))\right)$$

where MC is the marginal cost function.

Now, in addition to this, we know that the marginal good requires fixed amount of capital and labor for any given level of industry production. So, we have that the marginal cost function is only a function of the total industry output, or: $MC\left(\sum_{j=1}^{m} y_j\right)$

Therefore, since $\sum y\left(p^*\right) = \sum x\left(p^*\right)$:

$$p^* = MC\left(\sum_{i=1}^{n} x_i(p^*)\right)$$

But then, if demand increases to $x'\left(p\right)$, we have that $\sum_{i=1}^{n} x_i'(p^*) > \sum_{i=1}^{n} x_i(p^*)$. So firms will sell more at the current price, and marginal cost will fall. Since firms operate in a perfectly competitive environment, prices must also fall. But lower prices increases demand, and further lowers marginal cost. It remains to show that there must be some p^{**} such that:

$$p^{**} = MC\left(\sum_{i=1}^{n} x_i'\left(p^{**}\right)\right)$$

Or, equivalently, that the process has a fixed point.

Since $F(K, L, \sum y_j)$ is bounded for any (K, L), there is a minimum marginal cost that is greater than 0 for any value of $\sum y_j$. Therefore, we must have that $\exists \underline{p} > 0$ s.t.

$$\underline{p} < MC\left(\sum_{i=1}^{n} x_i'\left(\underline{p}\right)\right)$$

And since MC is continuous (by F continuous) and we have:

$$p^* > MC\left(\sum_{i=1}^{n} x_i'\left(p^*\right)\right) > MC\left(\sum_{i=1}^{n} x_i'\left(\underline{p}\right)\right) > \underline{p}$$

there must be a $p^{**} \in \left(\underline{p}, p^*\right)$ s.t. $MC\left(\sum_{i=1}^{n} x_i'\left(p^{**}\right)\right) = p^{**}$. ■

Proposition 2.8 Assume that the capital/labor ratio is constant for all goods at all periods of time. Then for any household in any two equilibria 0 and 1, $\omega^0 = \alpha\omega^1$, where α is common to all households. Further $\frac{p^1}{\alpha}.x^1 = \omega^0$ and $p^0 \neq \frac{1}{\alpha}p^1$.

PROOF: Assume not. Without loss of generality, take any two equilibria, 0 and 1. We know by the constant capital/labor ratio that:

$$\frac{w^1}{r^1} = \frac{w^0}{r^0}$$

Therefore if we have that $w^1 = \alpha w^0$, we must have that $r^1 = \alpha r^0$ and so:

$$\omega^0 = w^0 L + r^0 K = \frac{1}{\alpha}w^1 L + \frac{1}{\alpha}r^1 K = \frac{1}{\alpha}\omega^1$$

And this holds for all levels of income. Therefore, we can simply compare equilibria by holding income constant and scaling all the prices, i.e.:

$$\begin{aligned} p_x^1.x^1 &= \omega^1 \\ &= \frac{1}{\alpha}\omega^0 \\ \Rightarrow \frac{p_x^1}{\alpha}.x^1 &= \omega^0 \end{aligned}$$

From now on, for simplicity we will adopt the convention that when we refer to p_x^1 we are talking about the *scaled* version of the price, i.e. $p_x^1.x^1 = \omega^0$.

There is only one menu of *scaled* prices such that an equilibrium can be achieved. This is because for any pair of goods (x, y) either for at least one group:

$$\frac{1}{p_x}\frac{\delta U}{\delta x}|_x = \frac{1}{p_y}\frac{\delta U}{\delta y}|_y$$

In which case, since x is fixed:

$$\frac{p_y}{p_x} = \frac{\frac{\delta U}{\delta x}|_x}{\frac{\delta U}{\delta x}|_y} = \beta_y$$

where the right hand side is fixed. Further we must have that:

$$p_x x + \beta_y p_x y + \beta_z p_z z + ... = \omega^0$$

which gives a unique solution for p_x. Alternatively, for all consumers there is a pair x, y s.t.

$$\frac{1}{p_x}\frac{dU}{dx}|_x > \frac{1}{p_y}\frac{dU}{dy}|_0$$

■

But then no one is consuming any y, so the price is simply $p(0)$ (the marginal cost at 0 given w^0 which is defined by ω^0).

Therefore, we have a unique *scaled* price vector for each equilibrium allocation.

Proposition 2.9 There is no crossing of marginal utilities. That is, for all $\alpha > 0$ $\frac{\delta U}{\delta x}|_{x=\bar{x}} > \frac{\delta U}{\delta y}|_{y=\bar{x}} \Rightarrow \frac{\delta U}{\delta x}|_{x=\alpha\bar{x}} > \frac{\delta U}{\delta y}|_{y=\alpha\bar{x}}$.

PROOF: Without loss of generality, let $\bar{x} > \bar{y}$. Since wealth expansion paths are linear for a fixed variety of goods, we have that:

$$\frac{1}{p_x}\frac{\delta U}{\delta x}|_{x=\bar{x}} = \frac{1}{p_y}\frac{\delta U}{\delta y}|_{y=\bar{y}} \Rightarrow \frac{1}{p_x}\frac{\delta U}{\delta x}|_{x=\alpha\bar{y}} = \frac{1}{p_y}\frac{\delta U}{\delta y}|_{y=\alpha\bar{y}}$$

Note that this holds even when another good is added, because the utility function is separable. This implies:

$$\frac{1}{p_x}\frac{\delta U}{\delta x}|_{x=\bar{x}} > \frac{1}{p_y}\frac{\delta U}{\delta y}|_{y=\bar{y}} \Rightarrow \frac{1}{p_x}\frac{\delta U}{\delta x}|_{x=\alpha\bar{x}} \neq \frac{1}{p_y}\frac{\delta U}{\delta y}|_{y=\alpha\bar{y}}$$

We need to rule out the case where, for some α:

$$\frac{1}{p_x}\frac{\delta U}{\delta x}|_{x=\bar{x}} > \frac{1}{p_y}\frac{\delta U}{\delta y}|_{y=\bar{y}} \text{ and } \frac{1}{p_x}\frac{\delta U}{\delta x}|_{x=\alpha\bar{x}} < \frac{1}{p_y}\frac{\delta U}{\delta y}|_{y=\alpha\bar{y}}$$

Assume that we can have this scenario. Now, since the marginal utility function is continuous, there must be a $\alpha' < \alpha$ s.t.

$$\frac{1}{p_x}\frac{\delta U}{\delta x}|_{x=\alpha'\bar{x}} = \frac{1}{p_y}\frac{\delta U}{\delta y}|_{y=\alpha'\bar{y}}$$

But then, by the first condition, we must have that $\beta = \frac{\alpha}{\alpha'}$ satisfies:

$$\frac{1}{p_x}\frac{\delta U}{\delta x}|_{x=\beta\alpha'\bar{x}} = \frac{1}{p_y}\frac{\delta U}{\delta y}|_{y=\beta\alpha'\bar{y}} \Leftrightarrow \frac{1}{p_x}\frac{\delta U}{\delta x}|_{x=\alpha\bar{x}} = \frac{1}{p_y}\frac{\delta U}{\delta y}|_{y=\alpha\bar{y}}$$

So we have a contradiction. Therefore, setting $\bar{x} = \bar{y}$, it must be that:

$$\frac{1}{p_x}\frac{\delta U}{\delta x}|_{x=\bar{x}} > \frac{1}{p_y}\frac{\delta U}{\delta y}|_{y=\bar{x}} \Rightarrow \frac{1}{p_x}\frac{\delta U}{\delta x}|_{x=\alpha\bar{x}} > \frac{1}{p_y}\frac{\delta U}{\delta y}|_{y=\alpha\bar{x}}$$

■

Proposition 2.10 If there are two goods x and y s.t. $\frac{\delta U}{\delta x}|_0 > (=)\frac{\delta U}{\delta y}|_0$, any equilibrium where an individual does not consume $x > (=)y$ must be Pareto dominated by another equilibrium.

PROOF: Assume not. Without loss of generality, take any two equilibria, 0 and 1. We also use the subset r to denote the consumption of the wealthiest group. First we can eliminate any equilibria that are Pareto-dominated.

We know that goods are ranked, and have the same technologies. So we know that:

$$\frac{1}{p(0)}\frac{\delta U}{\delta x}|_0 > \frac{1}{p(0)}\frac{\delta U}{\delta y}|_0 \Rightarrow \frac{1}{p(0)}\frac{\delta U}{\delta x}|_x > \frac{1}{p(0)}\frac{\delta U}{\delta y}|_x$$

Without loss of generality let $\frac{1}{p(0)}\frac{dU}{dx}|_0 > \frac{1}{p(0)}\frac{dU}{dy}|_0$. Then, for any given level of consumption:

$$\frac{1}{p\left(\sum x\right)}\frac{\delta U}{\delta x}|_{x_r} > \frac{1}{p\left(\sum x\right)}\frac{\delta U}{\delta y}|_{x_r}$$

Therefore, we would always have that $y_r < x_r$ unless, $\sum y > \sum x$.

Now if $\sum y > \sum x$, however, all groups prefer to move back to consuming more x. Why? Because if we collectively reverse our consumption of x and y, the prices of the two goods are the same as before, but our higher marginal utility means that we prefer the consumption of x, i.e., if $x^1 = y^0$ and $y^1 = x^0$:

$$p\left(\sum x^1\right)x_i^1 + p\left(\sum y^1\right)y_i^1 = p\left(\sum y^1\right)x_i^0 + p\left(\sum x^1\right)y_i^0 \quad \blacksquare$$

but $U\left(x_i^0, y_i^0\right) > U\left(x_i^1, y_i^1\right)$

So if the society as a whole are consuming more of y than x, all groups would prefer an equilibrium where we switched those two goods around.

Therefore, if we are looking for Pareto-optimal equilibrium, we can exclude this set.

Proposition 2.11 Assume the capital/labor ratio is constant for all goods, the production technology is the same and neither equilibrium is Pareto-dominated. Assume there is sufficient inequality between households that there are at least as many classes as goods in all equilibria. Then if there is more than one equilibrium, there is a unique equilibrium that is preferred by the wealthiest household(s).

PROOF: Without loss of generality choose two equilibria and some group i. Take any good x s.t. i consumes more x in equilibrium 0. Then take any good y s.t. i consumes more y in equilibrium 1. Because preferences are the same for all households, if one household wants to increase consumption of one good and decrease consumption of another when prices change, no household will want to do the reverse. Therefore $\sum x^0 > \sum x^1$ and $\sum y^0 < \sum y^1$.

Now, since i is indifferent between 0 and 1:

$$\frac{p\left(\sum y^1\right)}{p\left(\sum x^1\right)} = \frac{\frac{\delta U}{\delta y}\big|_{y_i^1}}{\frac{\delta U}{\delta x}\big|_{x_i^1}} < \frac{\frac{\delta U}{\delta y}\big|_{y_i^0}}{\frac{\delta U}{\delta x}\big|_{x_i^0}} = \frac{p\left(\sum y^0\right)}{p\left(\sum x^0\right)}$$
$$\Rightarrow \frac{p\left(\sum x^1\right)}{p\left(\sum x^0\right)} > \frac{p\left(\sum y^1\right)}{p\left(\sum y^0\right)}$$

Therefore either $p\left(\sum x^1\right) > p\left(\sum x^0\right)$ or $p\left(\sum y^1\right) < p\left(\sum y^0\right)$.

Now if there are no economies of scale on x, then, we know that for each household in i

$$p_x^1.x_i^0 + p_y^1 y_i^0 < p_x^0 x_i^0 + p_y^0 y_i^0$$

(note here that we assume each individual household in y is too small to make a change on the price. This does *not* affect our statement about coordination, since coordination refers to choice of equilibria, not equilibrium stability).

So therefore i can afford to consume as much x and more y at 1 as they did at 0, with some to spare. So equilibrium 1 is no longer stable.

We can construct an identical argument to show that there must be economies of scale on y, or the equilibrium must be unique.

Therefore, for there to be two equilibria, we must have that there are economies of scale on all goods that change.

First we will consider the case where only x and y change. Then assume that x is the higher-ranked good (so 0 is the equilibrium in which we consume more of the higher-ranked good).Then, for a contradiction, assume that the wealthiest household r is indifferent between 0 and 1, and both are Pareto-superior.

First, anyone who is *not* consuming y but is consuming x will strictly prefer equilibrium 0, as in equilibrium 0 they face the lowest price. By our assumption that there is sufficient inequality to generate a class of households for each good at all equilibria, there must be a household consuming x but not y in equilibrium 1 (and therefore also in equilibrium 0). Therefore, there must be an $l < r$ who consumes both x and y at least in equilibrium 1 and prefers equilibrium 1 to equilibrium 0, otherwise 1 is Pareto-dominated by 0.

Further l must consume less goods than r, because otherwise l and r are covered by the same representative consumer, and is therefore also indifferent between the two equilibria. However, in that case, by separability and our linear expansion paths, we know that:

$$\begin{aligned} \frac{1}{p_x^0}\frac{dU}{dx}\Big|_{x_r^0} &= \frac{1}{p_y^0}\frac{dU}{dy}\Big|_{y_r^0} \\ \Rightarrow \frac{1}{p_x^0}\frac{dU}{dx}\Big|_{\frac{\omega_r}{\omega_l}x_r^0} &= \frac{1}{p_y^0}\frac{dU}{dy}\Big|_{\frac{\omega_r}{\omega_l}y_r^0} \\ \frac{\frac{dU}{dx}|_{x_r^0}}{\frac{dU}{dx}|_{\frac{\omega_r}{\omega_l}x_r^0}} &= \frac{\frac{dU}{dy}|_{y_r^0}}{\frac{dU}{dy}|_{\frac{\omega_r}{\omega_l}y_r^0}} \end{aligned}$$

But then also

$$\frac{\frac{dU}{dx}|_{x_r^1}}{\frac{dU}{dx}|_{\frac{\omega_r}{\omega_l}x_r^1}} = \frac{\frac{dU}{dy}|_{y_r^1}}{\frac{dU}{dy}|_{\frac{\omega_r}{\omega_l}y_r^1}}$$

■

So the utility ratios must be the same. So if we are indifferent between $(x_r^1, y_r^1) \sim (x_r^0, y_r^0)$ then we must have that $(x_l^1, y_l^1) \sim (x_l^0, y_l^0)$.

Once again, by our piecewise linear expansion paths, we know that if r is indifferent between two bundles, l must be indifferent between two bundles with the same ratios of goods (otherwise our expansion paths would break down). Therefore, if l is consuming x and y in both equilibrium 0 and equilibrium 1, they must be indifferent between both equilibria, and again we have a contradiction. So it must be that l is consuming both x and y in equilibrium 1, but only x in equilibrium 0.

But in that case, l can afford to consume $\left(\frac{\omega_l}{\omega_r}x_r^0, \frac{\omega_l}{\omega_r}y_r^0\right)$, and are be indifferent between that bundle and their bundle at equilibrium 1 $\left(\frac{\omega_l}{\omega_r}x_r^1, \frac{\omega_l}{\omega_r}y_r^1\right)$. But instead, at 0 they choose $(x_l^0, 0)$, which must therefore be preferred to $\left(\frac{\omega_l}{\omega_r}x_r^1, \frac{\omega_l}{\omega_r}y_r^1\right)$. Therefore l also prefers equilibrium 0.

Proposition 2.12 As a first-order effect, there will be more consumption of poor and middle-class goods in the equal society than under the unequal society, at the equal society's prices.

PROOF: Since everyone in both societies is consuming the poor goods, the demand curve for any one good x must be concave across its entire relevant range.

Poor goods: By the weak concavity of the demand function, we must have that $n_p^0 y\left(p^0, k_p\right) + n_m^0 y\left(p^0, k_m\right) + n_r^0 y\left(p^0, k_r^0\right) \geq n_p^1 y\left(p^0, k_p\right) + n_m^1 y\left(p^0, k_m\right) + n_r^1 y\left(p^0, k_r^0\right)$. So we have that the total level of consumption is higher under the equal society if we fix the prices at the equilibrium prices of the equal society.

Middle-class goods: Because society 1 has the same amount of capital, and the same number of people as the unequal society:

$$\left(n_r^1 - n_r^0\right) k_r = \left(n_m^0 - n_m^1\right)\left(k_m - k_p\right) - \left(n_r^0 - n_r^1\right) k_p$$

which is > 0 since by the increase in inequality we have that $n_r^1 k_r > n_r^0 k_r$.

$$\begin{aligned}
\Rightarrow \left(n_r^1 - n_r^0\right)\left(k_r - k_p\right) &= \left(n_m^0 - n_m^1\right)\left(k_m - k_p\right) \\
\Rightarrow \left(n_r^1 - n_r^0\right)\frac{\left(k_r - k_p\right)}{\left(k_m - k_p\right)} &= \left(n_m^0 - n_m^1\right)
\end{aligned}$$

Fixing the price vector at p^0, the difference in consumption of middle-class goods between the two societies is:

$$\begin{aligned}
D &= n_r^1 x_m\left(p^0, k_r\right) + n_m^1 x_m\left(p^0, k_m\right) - n_r^0 x_m\left(p^0, k_r\right) - n_m^0 x_m\left(p^0, k_m\right) \\
&= \left(n_r^1 - n_r^0\right) x_m\left(p^0, k_r\right) - \left(n_r^1 - n_r^0\right)\frac{\left(k_r - k_p\right)}{\left(k_m - k_p\right)} x_m\left(p^0, k_m\right)
\end{aligned}$$

Note that $k_p = 0$. Therefore, for a contradiction, assume that $D > 0$. Then:

$$\begin{aligned}
\left(n_r^1 - n_r^0\right) x_m\left(p^0, k_r\right) - \left(n_r^1 - n_r^0\right) \frac{(k_r - k_p)}{(k_m - k_p)} x_m\left(p^0, k_m\right) &> 0 \\
\Rightarrow \frac{x_m\left(p^0, k_r\right)}{x_m\left(p^0, k_m\right)} &> \frac{(k_r - k_p)}{(k_m - k_p)} \qquad (2.3) \\
\Rightarrow \frac{x_m\left(p^0, k_r\right)}{x_m\left(p^0, k_m\right)} &> \frac{k_r}{k_m} \text{ as } k_p \to 0
\end{aligned}$$

But we know that:

$$\begin{aligned}
\Rightarrow \frac{p^0 x\left(p^0, k_r\right)}{p^0 x\left(p^0, k_m\right)} &> \frac{k_r}{k_m} \\
&> \frac{r k_r + wL}{r k_m + wL} \\
&\geq \frac{p^0 y\left(p^0 k_r\right)}{p^0 y\left(p^0 k_m\right)}
\end{aligned}$$

Therefore, x must be increasing faster than y as income increases, which contradicts our assumption of linear wealth expansion paths.

Note that we can prove this more generally (for $k_p > 0$), so long as the rich are sufficiently wealthy.

Fix k_p, k_m. Then, if we increase k_r, we have (by 2.3):

$$\frac{\delta \frac{(k_r - k_p)}{(k_m - k_p)}}{\delta k_r} = \frac{1}{(k_m - k_p)} > 0$$

which is necessarily constant.

We also have:

$$\frac{\delta \frac{x\left(p^0, k_r\right)}{x\left(p^0, k_m\right)}}{\delta k_r} = \frac{1}{x\left(p^0, k_m\right)} \frac{\delta x\left(p^0, k_r\right)}{\delta k_r} > 0$$

$$\frac{\delta^2 \frac{x\left(p^0, k_r\right)}{x\left(p^0, k_m\right)}}{\delta k_r^2} = \frac{1}{x_m\left(p^0, k_m\right)} \frac{\delta^2 x\left(p^0, k_r\right)}{\delta k_r^2} < 0 \text{ through large enough range of } k_r$$

Since the right hand side of (1) is increasing at a constant rate, whereas the left hand side is increasing at a declining rate, at some point the two must cross. Therefore, so long as k_r is sufficiently large, we must have that D is negative, and consequently the total consumption of middle-class goods is smaller under the unequal society than under the equal society. The point will depend on values of k_p (the smaller k_p, the smaller the crossing point), the taste for variety (this will determine the concavity of x_m), and the quantity of income from labor (if $x(p^0, 0) > 0$, then the crossing point will require a higher k_r). ■

Proposition 2.13 Assume that the optimal labor/capital ratio is the same for all goods, at all levels of production. Assume also there are sufficient numbers of poor households to exhaust economies of scale on poor goods. Assume that production quantities are too small for economies of scale to exist on rich goods. Then, if the intervention is small enough and economies of scale are not too large, every member of the equal society (excluding those who have lost wealth) will consume weakly more than the equivalent member of the unequal society (at the equal society's prices), and will therefore be better off than their counterparts in the unequal society.

PROOF: Since there is an constant optimal capital/labor ratio, we know that by proposition 2.8 that there is only one set of *scaled* prices (prices will scale by the same factor as income, so we can ignore the effects of changes in income). We will work with scaled prices for the remainder of this proof.

It is sufficient to show that $p^0.x^1 < p^0.x^0$ for everyone. This is clearly sufficient for the change in output, and is also sufficient for a change in welfare since otherwise a household could defect and purchase the bundle that they purchased at 1. So, for a contradiction, we need to show that $p^0.x^1 > p^0.x^0$ for someone.

This is equivalent to saying that:

$$\begin{aligned}
p^0.x^1 - p^0.x^0 &\geq 0 \\
\Rightarrow p^0.x^1 &\geq w^0 L + r^0 K \\
&= \frac{w^0}{w^1}\left[w^1 L + r^1 K\right] \\
&= \frac{w^0}{w^1} p^1.x^1
\end{aligned}$$

So we must have that $p^0 \geq \frac{w^0}{w^1} p^1$ for some good.

Now, we know that for all rich goods, $p_r^0 \leq \frac{w^0}{w^1} p_r^1$. This is simply because the price of the rich goods cannot rise any higher. Therefore, the good that increases in price must be either a poor good or a middle-class good.

By the proposition 2.12, we know that if we hold prices constant, consumption of middle-class goods will rise, under this scenario. The middle-class good will therefore only increase in price if either the rich or the middle-class choose to consume *less* of the middle-class good than they did in the unequal equilibrium (the poor cannot consume less, since they weren't consuming any before).

Remember that in equilibrium 1, rich households faced the first-order condition:

$$\begin{aligned}
\frac{1}{p\left(n_m^1 x_m^1 + n_r^1 x_r^1\right)} \frac{dU}{dx}|_{x_r^1} &= \frac{1}{p\left(n_r^1 z_r^1\right)} \frac{dU}{dz}|_{z_r^1} \\
&= \frac{1}{p\left(n_r^0 z_r^1\right)} \frac{dU}{dz}|_{z_r^1}
\end{aligned}$$

However, after the re-allocation, the rich face the following:

$$\frac{1}{p\left(n_m^0 x_m^1 + n_r^0 x_r^1\right)} \frac{dU}{dx}|_{x_r^1} > \frac{1}{p\left(n_m^1 x_m^1 + n_r^1 x_r^1\right)} \frac{dU}{dx}|_{x_r^1}$$

So the rich will substitute away from rich goods and toward middle-class goods. Both the middle-class and the rich will also substitute away from poor goods, since:

$$\begin{aligned}
\frac{1}{p\left(n_m^0 x_m^1 + n_r^0 x_r^1\right)} \frac{dU}{dx}|_{x^1} &> \frac{1}{p\left(n_m^1 x_m^1 + n_r^1 x_r^1\right)} \frac{dU}{dx}|_{x^1} \\
&= \frac{1}{p\left(n_m^1 y_m^1 + n_r^1 y_r^1 + n_p^1 y_p^1\right)} \frac{dU}{dx}|_{y^1} \\
&= \frac{1}{p\left(n_m^0 y_m^1 + n_r^0 y_r^1 + n_p^0 y_p^1\right)} \frac{dU}{dx}|_{y^1}
\end{aligned}$$

Where the last equality follows from the fact that the first-order effect is to increase demand for poor goods (for which prices cannot fall).

Thus, we have a second-order increase in demand for middle-class goods, and their price will fall. Subsequent-order effects follow the same pattern. However, since economies of scale are bounded, and marginal cost is declining, there must be a point at which

$$\frac{1}{p\left(n_m^0 x_m^0 + n_r^0 x_r^0\right)} \frac{dU}{dx}|_{x^1} = \frac{1}{p\left(n_m^0 y_m^1 + n_r^0 y_r^1 + n_p^0 y_p^1\right)} \frac{dU}{dx}|_{y^0}$$

So we have convergence on a lower price.

It remains to consider the impact of prices on poor goods. Since the poor have sufficient numbers (and have not lost any income), by our equilibrium refinement if they choose an equilibrium where they consume less poor goods than would be sufficient to exhaust all the economies of scale, then they must be better off. By our construction, if they choose to consume less poor goods, they must be consuming more middle-class goods, and since they didn't do that in equilibrium 1 the prices of middle-class goods in equilibrium 0 must be lower than in equilibrium 1. However, if the poor do choose to make the switch, then:

$$u(z_p^0, x_p^0, 0) > u(z_p^1, 0, 0)$$

So the middle-class will consume $(\frac{\omega_m^0}{\omega_p} z_p^0, \frac{\omega_m^0}{\omega_p} x_p^0, 0)$, by linear expansion paths. Now if this is not as good as the middle-class's previous option, then:

$$\begin{aligned}
u(\frac{\omega_m^0}{\omega_p}z_p^0, \frac{\omega_m^0}{\omega_p}x_p^0, 0) &< u(z_m^1, x_m^1, 0) \\
\Rightarrow u(z_p^0, x_p^0, 0) &< u(\frac{\omega_p}{\omega_m}z_m^1, \frac{\omega_p}{\omega_m}x_m^1, 0) \\
&< u(\frac{\omega_p}{\omega_m}z_m^1, \frac{\omega_p}{\omega_m}x_m^1, 0) \\
&\leq u(z_p^1, 0, 0)
\end{aligned}$$

So the middle-class must be better off too. The same construction can be used to show that the rich must also be better off. Therefore, even if the prices on poor goods rise everyone will be better off. ■

Proposition 2.14 That total economic output is higher at equilibrium 0 than at equilibrium 1, at the equal economy's prices.

PROOF: Assume not. Then, it must be that:

$$\begin{aligned}
n_r^1 p^1 \cdot x_r^1 + n_m^1 p^1 \cdot x_m^1 + n_r^1 p^1 \cdot x_p^1 &> n_r^0 p^0 \cdot x_r^0 + n_m^0 p^0 \cdot x_m^0 + n_p^0 p^0 \cdot x_p^0 \\
\Rightarrow \left(n_r^1 - n_r^0\right) p^1 \cdot x_r^1 + \left(n_p^1 - n_p^0\right) p^1 \cdot x_p^1 &> n_r^0 \left(p^0 \cdot x_r^0 - p^1 \cdot x_r^1\right) + \left(n_m^0 - n_m^1\right) p^1 \cdot x_m^1 \\
&\quad + n_m^0 \left(p^0 \cdot x_m^0 - p^1 \cdot x_m^1\right) + n_r^0 \left(p^0 \cdot x_p^0 - p^1 \cdot x_p^1\right)
\end{aligned}$$

But, since we know that we maintained the same average wealth when we redistributed, and the consumption bundles for each group are lower under 1 than 0, we know that:

$$\begin{aligned}
\left(n_r^1 - n_r^0\right) p^0 \cdot x_r^0 + \left(n_p^1 - n_p^0\right) p^0 \cdot x_p^0 &= \left(n_m^0 - n_m^1\right) p^0 \cdot x_m^0 \\
\Rightarrow \left(n_r^1 - n_r^0\right) p^1 \cdot x_r^1 + \left(n_p^1 - n_p^0\right) p^1 \cdot x_p^1 &< \left(n_m^0 - n_m^1\right) p^0 \cdot x_m^0
\end{aligned}$$

Combining this result with the above gives:

$$\begin{aligned}
\Rightarrow \left(n_m^0 - n_m^1\right) p^0 \cdot x_m^0 &> n_r^0 \left(p^0 \cdot x_r^0 - p^1 \cdot x_r^1\right) + \left(n_m^0 - n_m^1\right) p^1 \cdot x_m^1 \\
&\quad + n_m^0 \left(p^0 \cdot x_m^0 - p^1 \cdot x_m^1\right) + n_r^0 \left(p^0 \cdot x_p^0 - p^1 \cdot x_p^1\right) \\
\Rightarrow n_m^1 \left(p^1 \cdot x_m^1 - p^0 \cdot x_m^0\right) &> n_r^0 \left(p^0 \cdot x_r^0 - p^1 \cdot x_r^1\right) + n_r^0 \left(p^0 \cdot x_p^0 - p^1 \cdot x_p^1\right) \\
\Rightarrow 0 > n_m^1 \left(p^1 \cdot x_m^1 - p^0 \cdot x_m^0\right) &> n_r^0 \left(p^0 \cdot x_r^0 - p^1 \cdot x_r^1\right) + n_r^0 \left(p^0 \cdot x_p^0 - p^1 \cdot x_p^1\right) > 0 \blacksquare
\end{aligned}$$

And we have a contradiction.

Proposition 2.15 As a first-order effect, there will be more consumption of middle-class and poor goods under the equal society than the unequal society at the equal society's prices.

PROOF: We know that (because society 1 is more unequal):

$$\begin{aligned}
n_r \omega_r^1 + n_m \omega_m^1 &= n_r \omega_r^0 + n_m \omega_m^0 \\
n_r \omega_r^1 &> n_r \omega_r^0
\end{aligned}$$

Since, at fixed prices, the change in demand for poor goods with respect to income for the middle-class is higher than the change in demand with respect to wealth for the rich, we must have that:

$$n_r y \left(p^0, \omega_r - \tau\right) + n_m y \left(p^0, \omega_m + \frac{n_r}{n_m} \tau\right) > n_r y \left(p^0, \omega_R\right) + n_m y \left(p^0, \omega_m\right)$$

$$\begin{aligned}
&\Rightarrow n_r y \left(p^0, \omega_r - \tau\right) + n_m y \left(p^0, \omega_m + \frac{n_r}{n_m} \tau\right) + n_p y \left(p^0, \omega_p\right) \\
&> n_r y \left(p^0, \omega_r\right) + n_m y \left(p^0, \omega_m\right) + n_p y \left(p^0, \omega_p\right)
\end{aligned}$$

So we have that the total level of consumption of poor goods is higher under the equal society if we fix the prices.

Fixing the price vector at the equal society's prices, the difference in consumption (at equal prices) of middle-class goods between the two societies is:

$$\begin{aligned} D &= n_r x\left(p^0, \omega_r^0\right) + n_m x\left(p^0, \omega_m^0\right) - n_r x\left(p^0, \omega_r^1\right) - n_m x\left(p^0, \omega_m^1\right) \\ &= n_r\left[x\left(p^0, \omega_r^0\right) - x\left(p^0, \omega_r^1\right)\right] + n_m\left[x\left(p^0, \omega_m^0\right) - x\left(p^0, \omega_m^1\right)\right] \end{aligned}$$

By the concavity of the demand function:

$$\begin{aligned} n_m\left[x\left(p^0, \omega_m + \frac{n_r}{n_m}\epsilon\right) - x\left(p^0, \omega_m\right)\right] &> x\left(p^0, \omega_m + \epsilon\right) - x\left(p^0, \omega_m\right) \\ &> x\left(p^0, \omega_r\right) - x\left(p^0, \omega_r - \epsilon\right) \\ &> n_r\left[x\left(p^0, \omega_r\right) - x\left(p^0, \omega_r - \epsilon\right)\right] \end{aligned}$$

Which implies that:

$$\begin{aligned} n_r\left[x\left(p^0, \omega_r\right) - x\left(p^0, \omega_r - \epsilon\right)\right] &< n_m\left[x\left(p^0, \omega_m + \frac{n_r}{n_m}\epsilon\right) - x\left(p^0, \omega_m\right)\right] \\ \Rightarrow n_r\left[x\left(p^0, \omega_r^1\right) - x\left(p^0, \omega_r^0\right)\right] &< n_m\left[x\left(p^0, \omega_m^0\right) - x\left(p^0, k_m^1\right)\right] \\ \Rightarrow 0 &< n_m\left[x\left(p^0, \omega_m^0\right) - x\left(p^0, \omega_m^1\right)\right] - n_r\left[x\left(p^0, \omega_r^1\right) - x\left(p^0, \omega_r^0\right)\right] \\ \Rightarrow 0 &< D \end{aligned}$$

So we consume more middle-class goods under the more equal society. ∎

Proposition 2.16 Assume that the optimal labor/capital ratio is the same for all goods, at all levels of production. Assume that there are enough poor that if they choose to consume only poor goods, they can exhaust all economies of scale on poor goods. Assume also that production quantities are too small for economies of scale to exist on rich goods. Then, if economies of scale are not too small on middle-class goods, then there is some $\tau > 0$ transfer from the rich to the middle-class such that we can achieve a Pareto-improvement.

PROOF: Since there is an constant optimal capital/labor ratio, we know that by proposition 2.8 that there is only one set of *scaled* prices (prices will scale by the same factor as income, so we can ignore the effects of changes in income). We will work with scaled prices for the remainder of this proof.

For the poor, $\tau = 0$. Since there are enough poor to exhaust economies of scale on poor goods (z), then regardless of what they choose as their bundle after redistribution (equilibrium 1), they can (as a class) afford their bundle under the pre-redistribution equilibrium (equilibrium 0). Therefore they must be weakly better off under the new equilibrium.

For middle class, $\tau < 0$. Therefore, it is sufficient to show that prices do not rise, since in that case the middle class can also afford their former bundle. Rich good prices cannot rise. Middle-class good prices cannot rise unless poor good or rich good prices rise, since at those prices the middle-class will purchase more middle-class goods (and if they do not, there is a Pareto-improving equilibrium where they do), see proposition 2.15.

Since the poor have sufficient numbers (and have not lost any income), by our equilibrium refinement if they choose an equilibrium where they consume less poor goods than would be sufficient to exhaust all the economies of scale, then they must be better off. By our construction, if they choose to consume less poor goods, they must be consuming more middle-class goods, and since they didn't do that in equilibrium 1 the prices of middle-class goods in equilibrium 0 must be lower than in equilibrium 1. However, if the poor do choose to make the switch, then:

$$u(z_p^0, x_p^0, 0) > u(z_p^1, 0, 0)$$

So the middle-class will consume $(\frac{\omega_m^0}{\omega_p} z_p^0, \frac{\omega_m^0}{\omega_p} x_p^0, 0)$, by linear expansion paths. Now if this is not as good as the middle-class's previous option, then:

$$
\begin{aligned}
u(\frac{\omega_m^0}{\omega_p}z_p^0, \frac{\omega_m^0}{\omega_p}x_p^0, 0) &< u(z_m^1, x_m^1, 0) \\
\Rightarrow u(z_p^0, x_p^0, 0) &< u(\frac{\omega_p}{\omega_m}z_m^1, \frac{\omega_p}{\omega_m}x_m^1, 0) \\
&< u(\frac{\omega_p}{\omega_m}z_m^1, \frac{\omega_p}{\omega_m}x_m^1, 0) \\
&\leq u(z_p^1, 0, 0)
\end{aligned}
$$

So the middle-class must be better off too. The same construction can be used to show that the rich must also be better off. Therefore, even if the prices on poor goods rise everyone will be better off. Further, if that is the case middle-class good prices must still fall.

So the middle class must be strictly better off.

If rich good prices do not rise and poor good prices do not rise, then middle-class good prices must fall, so long as:

$$
\frac{1}{p\,(n_m x_m^0 + n_r x_r^0)} \frac{dU}{dx}|_{x_m^0} > \frac{1}{p\,(n_r y_r^1)} \frac{dU}{dy}|_0
$$

This is because, holding the other prices constant, increasing the income of the middle-class will result in more income being spent on middle-class goods, which will in turn lower their price (see proposition 2.15) for details.

For the rich, $\tau > 0$. Therefore, they will lose a percentage of each type of good at the initial stages. For the rich to be better off, in that case, we must have the middle-class prices fall enough to offset both the tax and any price increase on poor goods.

Then, for the rich:

$$
U = u(x, y, z) = u\left(z, x, \left[\frac{\omega - \tau - p_z z - p_x x}{p_y}\right]\right)
$$

So:

$$\frac{dU}{d\tau} = \frac{du}{dz}\frac{dz}{d\tau}+\frac{du}{dx}\frac{dx}{d\tau}+\frac{du}{dy}\frac{1}{p_y}\left(-\frac{dp_z}{d\tau}z - \frac{dp_x}{d\tau}x - p_z\frac{dz}{d\tau} - p_x\frac{dx}{d\tau} - 1\right)-\frac{1}{p_y^2}\frac{dp_y}{d\tau}\frac{du}{dy}\left(\omega - \tau - p_z z - p_x x\right)$$

$$\begin{aligned}
\Rightarrow \frac{dU}{d\tau}|_{\tau=0} &= \frac{du}{dz}\frac{dz}{d\tau} + \frac{du}{dx}\frac{dx}{d\tau} + \frac{du}{dy}\frac{1}{p_y}\left(-\frac{dp_z}{d\tau}z - \frac{dp_x}{d\tau}x - p_z\frac{dz}{d\tau} - p_x\frac{dx}{d\tau} - 1\right) \\
&= \left(\frac{du}{dz} - \frac{du}{dy}\frac{p_z}{p_y}\right)\frac{dz}{d\tau} + \left(\frac{du}{dx} - \frac{du}{dy}\frac{p_x}{p_y}\right)\frac{dx}{d\tau} + \frac{du}{dy}\frac{1}{p_y}\left(-\frac{dp_z}{d\tau}z - \frac{dp_x}{d\tau}x - 1\right) \\
&= -\frac{du}{dy}\frac{1}{p_y}\left(\frac{dp_z}{d\tau}z + \frac{dp_x}{d\tau}x + 1\right)
\end{aligned}$$

Then, we know that if p_z doesn't change:

$$\Rightarrow \frac{dU}{d\tau} = -\frac{du}{dy}\frac{1}{p_y}\left(\frac{dp_x}{d\tau}x + 1\right)$$

Which is weakly positive so long as:

$$\begin{aligned}
\frac{dp_x}{d\tau}x + 1 &\leq 0 \\
\Rightarrow \frac{dp_x}{d\tau} &\leq -\frac{1}{x}
\end{aligned}$$

If p_z does change it must increase. So in that case, economies of scale on x need to be larger:

$$\frac{dp_x}{d\tau} \leq -\frac{1}{x}\left(1 + \frac{dp_z}{d\tau}z\right)$$

Therefore, so long as the economies of scale are sufficiently large $\frac{dp_x}{d\tau} \leq -\frac{1}{x_r}\left(1 + \frac{dp_z}{d\tau}z\right)$, since the rich are consuming some rich goods in equilibrium 1 there must be some transfer that is sufficiently small that there is a Pareto optimal redistribution.

Finally, we need to consider whether it is possible to obtain the same result without redistribution. In that case, however, because the wealthy are consuming rich goods in a fixed quantity, if given more income at this stage they will spend more of it on rich goods and less on middle-class goods. Therefore, the price of middle-class goods must rise, and we cannot sustain the equilibrium. ■

Corollary 2.1Given the framework in proposition 2.16 there is a limit to redistribution beyond which no additional transfer is Pareto-improving.

PROOF: At equilibrium 1, the middle class have the following constraint:

$$\frac{1}{p_{x^1}}\frac{dU}{dx}|_{x_m^1} > \frac{1}{p_{y^1}}\frac{dU}{dy}|_0$$

After the transfer, the middle-class consume more x, so

$$\frac{dU}{dx}|_{x_m^0} < \frac{dU}{dx}|_{x_m^1}$$

However, price is also falling so

$$p_x^0 < p_x^1 \Rightarrow \frac{1}{p_{x^0}} > \frac{1}{p_{x^1}}$$

and $\frac{1}{p_{y^1}}\frac{dU}{dy}|_0$ remains constant.

now, we know that there is a lower limit to p_x^0 (>0). So at some point, p_x^0 must stop declining. However, $\frac{dU}{dx}|_{x_m^0} \to 0$ as $x_m^0 \to \infty$. Since $\frac{1}{p_{y^1}}\frac{dU}{dy}|_0 = \frac{1}{p_{y^0}}\frac{dU}{dy}|_0$ and is constant, at some point we must have that, as the transfer gets large:

$$\frac{1}{p_{x^0}}\frac{dU}{dx}|_{x_m^0} < \frac{1}{p_{y^0}}\frac{dU}{dy}|_0$$

■

At that point, the middle-class will start to consume rich goods. By linear expansion paths the middle-class will consume goods in the same ratios as the rich. At that point, any additional transfer from the rich to the 'middle class' will result in the same aggregate ratios being consumed as before that additional transfer (since we are in equilibrium). Thus no additional transfer can result in a price change for any goods, so the rich will simply lose the value of the transfer and not gain anything in return. Thus the transfer cannot be Pareto-improving.

Corollary 2.2 Given the framework in proposition 2.16, even if economies of scale are not large enough locally to produce a small Pareto-improving transfer, there may be a large Pareto-improving transfer, so long as the threshold in corollary 2.1 is not crossed.

PROOF: Since economies of scale are not decreasing at a strictly decreasing rate, it is possible that while there is no small transfer from the rich to the middle-class, there is a large transfer.

By the same arguments as in proposition 2.16, the poor (weakly) and middle-class (strictly) must always be better off after the transfer. It remains to show that that can also be true for the rich.

It is sufficient to show that if economies of scale are sufficiently large, then for the rich:

$$\begin{aligned} p^0.x^1 - p^0.x^0 &\leq 0 \\ \Rightarrow p^0.x^1 &\leq w^0 L + r^0 K - \frac{n_r}{n_i}\tau \\ &= \frac{w^0}{w^1}\left[w^1 L + r^1 K\right] - \frac{n_r}{n_i}\tau \\ &= p^1.x^1 - \frac{n_r}{n_i}\tau \\ \Rightarrow \left(p^0 - p^1\right).x^0 &\leq -\frac{n_r}{n_i}\tau \end{aligned}$$

Now we know that: $p_y^0 = p_y^1$, $p_z^0 \geq p_z^1$ and $p_x^0 \leq p_x^1$.

Therefore, for the rich we need that:

$$\begin{aligned} \left(p_x^0 - p_x^1\right)x^0 + \left(p_z^0 - p_z^1\right)z^0 &\leq -\tau \\ \Rightarrow \left(p_x^0 - p_x^1\right) &\leq -\frac{1}{x^0}\left(\tau + \left(p_z^0 - p_z^1\right)z^0\right) \end{aligned}$$

■

Which is essentially exactly the same condition as in proposition 2.16. However it is not required to be true for consecutive values of τ. This is because $p_x^0 - p_x^1$ may be small for low values of τ(since $\frac{dp_x}{dx}$ is not strictly declining, so $\frac{dp_x}{d\tau}$ need not be either) but larger for large values of τ.Of course, once the threshold in corollary 2.1 is crossed, it is no longer clear that there will be a benefit from the transfer.

Further by the same argument as in proposition 2.16, this cannot be supported by another equilibrium.

Proposition 2.18 Any losses from moving toward a more unequal society can be overcome if they are matched by a sufficient increase in the *size* of the economy.

PROOF: "intervention" Case: Take $\frac{n_m^0}{n_m^1}$. Then, if we multiply the number of people in the unequal society by this fraction, we will have:

$$\begin{aligned}
\frac{n_m^0}{n_m^1} n_r^1 &> n_r^0 \text{ rich households;} \\
\frac{n_m^0}{n_m^1} n_p^1 &> n_p^0 \text{ poor households; and} \\
n_m^0 &= n_m^0 \text{ middle-class households}
\end{aligned}$$

"Taxation" Case: Simply add sufficient additional players to both the rich and the middle-class, such $\hat{n}_m \omega_m = n_m (\omega_m + \tau)$. Add a proportional number of players to all other classes.

In either case, as a first-order effect, the price of middle-class goods will fall to their same level as under the respective equal equilibria. Since there are more members of other classes also, the total movement of purchasing power into middle-class goods will be weakly higher. Therefore, the price of middle-class goods will be weakly lower, and we are done. ∎

Chapter 3

Political Economy of Housing Supply — Evidence from Australian Cities

3.1 Introduction

Steadily climbing housing prices and increasing housing shortages have become key features of major cities across the world. These trends have been ongoing for several decades, interrupted only briefly by crises as occurred in 2008. These trends are not limited to well-known US examples such as New York and San Francisco, but extend to major cities globally, including London, Sydney, and Vancouver.

While a body of literature ascribes these dynamics to increased regulation (see Glaeser et al. (2005b) for example, or Quigley and Rosenthal (2005) for a summary), this merely pushes the question back one step. Why does this regulation exist at all? And if regulation is in fact the driver of housing price growth and slow supply responses, why has regulation become more restrictive over time?

This paper considers these questions. First, what are the demographic drivers of increased housing development regulation (and therefore slower housing growth)? And secondly, what political structures support this process? To answer these questions, we develop a model of housing markets, where residents are averse to congestion. Our model predicts that homeowners will tend to be more averse to growth than renters, and therefore more likely to vote for any one of the various broad categories of regulation we consider. We also model voting, to demonstrate that neighborhood-level preferences will matter more when the government is decentralized.

Australia is a good location to test the predictions of the model. In Australia, unlike the US, local governments (the primary town planning regulators) do not provide essential public safety or schooling services. This allows us to test the impact of demographic factors on housing regulation without worrying about the overwhelming selection effects that would otherwise result. In addition, Australia offers us variation in the degree of governance centrality, and rich 5-yearly census data.

We find that rates of home ownership are closely tied to slower growth in cities with a decentralized governance structure. Our finding is consistent with existing theories that homeowners oppose growth, as growth reduces their expected home valuations and increases home price risk. We do *not* find the same relationship between neighborhoods in a city managed by a single local government, which supports our hypothesis that homeowners act as housing 'insiders' who construct artificial barriers to entry to reduce local congestion.

In addition, we find that the presence of large numbers of pensioners is negatively correlated with housing growth in a given neighborhood regardless of the governance structure. Perhaps surprisingly, higher incomes seem to be correlated with slightly higher rates of housing development.

Section 3.2 presents the current literature on housing growth and regulation, and also provides some historical background. Section 3.3 outlines a model of development that incorporates congestion preferences, and examines the predictive of various demographical variables and governance structures on housing outcomes. Section 3.4 describes the data used for testing these predictions in major Australian cities. Section 3.5 presents our results.

3.2 Literature Review

Recent years have seen the accumulation of a preponderance of evidence that regulation is a driving force in determining housing supply and prices. Glaeser et al. (2005b) first demonstrated that for Manhattan, at least 50 percent of the value of the mean owned apartment in Manhattan in 2003 was the result of regulatory constraints. Similar results have followed for other areas, including Boston (Glaeser et al., 2006) and California (Quigley and Raphael, 2005).

There is little consensus on why this relationship exists. It appears to affect some places more than others, specifically coastal cities Gyourko (2009). This has led some to argue that the limiting factor is geography, not regulation (Saiz, 2010), but this stands in contrast with the finding that the relationship between land regulation and housing prices has strengthened over time (Glaeser et al., 2005a).

One hypothesis is that "homevoters" use their political power to enhance their asset values. Externalities from new developments may have positive or negative effects on housing prices, and existing residents will be opposed to anything that either reduces their house price in expectation (Cooley and LaCivita, 1982) or increases their exposure to housing price risk in the event that they are using their homes as a hedge against consumption or productivity shocks (Ortalo-Magné and Rady, 2002; Ortalo-Magné and Prat, 2007). Risk aversion is higher for those whose wealth is concentrated in their home, rather than diversified across a range of assets (Fischel, 2001).

Empirically, Dubin et al. (1992) analyzed voting data from growth control measures proposed in San Diego, and found that homeowners were in fact more likely to vote in support of those measures. They also found an increase in support for growth control from those living in areas with high levels of traffic congestion.

There is also considerable debate over what conditions would allow homeowners to have the most political power. On the one hand, Ellickson (1977) suggests that small areas with limited local government issues may be more likely to exclude development, and for this reason Gottlieb and Glaeser (2008) argue that a national policy should be implemented to reduce the power of local homeowners. On the other hand, White (1975) and Hamilton (1978) proposed the "monopoly zoning hypothesis": larger towns will have stricter land use regulations as they can more easily exploit market power. Extensive econometric analysis has failed to resolve this question.[1]

An alternative view is that slower housing development is the result of inequality.

[1] Fischel (1980); Thorson (1996); or see Quigley and Rosenthal (2005) for a summary.

There is evidence to suggest that income inequality may raise the cost of housing for the lowest income earners (Matlack and Vigdor, 2008), and there is a long-run correlation between UK housing prices and income inequality (Green and Shaheen, 2014). Inequality is usually hypothesized to reduce 'filtering', where new middle-class housing deteriorates and eventually becomes housing for the poor (Quigley and Raphael, 2004), thereby limiting demand rather than supply. However, Quigley and Rosenthal (2005) argue that wealthy communities may also prefer stricter regulation which may in turn elevate housing prices.

One final characteristic that may be important is age. It has been well established that those over 65 are more likely to be members of a political party, and to vote in voluntary elections (Verba et al. 1995). Higher rates of civic participation may mean that their preferences effectively carry more weight than those of other citizens in local government decision-making. In addition, there is evidence that, at least in the case of homeowners, the elderly are less likely to move (Burkhauser et al., 1995). This could increase their incentives to ensure that their existing neighborhood meets their desires in terms of congestion and development. While it is possible that the incentives could work in either direction — the elderly may be more opposed to congestion and change than younger residents, or alternatively be more price-sensitive on average — we consider this an important avenue of investigation.

Beyond simple correlations these theories have suffered from a lack of clear evidence. The legal framework behind housing supply is extremely complicated — housing is governed directly by zoning laws, and indirectly by environmental, heritage, health and safety and other legislation, depending on the jurisdiction. On-paper regulation may also not translate directly into what occurs in practice. This makes regulation incredibly difficult to measure (Quigley and Rosenthal, 2005).

In addition, there are difficulties in establishing the size of the impact of any given regulation. There may be a long gap between the implementation of regulation and the production of new housing. Furthermore, a single jurisdiction provides very little variation in the proportions of owners, renters, wealthy, poor, young, and elderly voters required to test their effects on home growth. By leveraging within-city and across-city differences over a relatively long time period, this paper seeks to fill that gap.

The majority of the literature on local governments assumes a localized economy, and thus incorporates production as well as housing consumption in determining a general equilibrium. We would argue that this is approach is misguided in the case of a modern city that is divided into various 'sub cities', which are separately governed. In many cases, the 'city' as conceptualized by the strength of economic ties across a contiguous area has no government of its own. Instead, administration of the area is divided between smaller cities, which have little or no general oversight.

This is highly visible in in the Australian context. In recent years, the Australian Bureau of Statistics has defined a Greater Capital City Statistical Area (see figure 3.1) for each city, which denotes the area that includes the whole urban sprawl, including those areas where people regularly commute to the middle of the city for recreation or employment. The Sydney Greater Capital City Statistical Area includes all or part of 40 separate local government areas, and there is no single authority responsible for the area as a whole.

This creates an assortment of citizen 'clubs' which take little or no account of the externalities they impose upon their neighbors. For instance, it is established that housing regulation in any one area has spillovers for housing prices in other close geographical areas (Levine, 1999). These impacts need not be considered by those who choose to impose that regulation. One could consider this a 'voting externality'; since geographical areas are not isolated from one another, all political decisions in a decentralized system have impacts on people outside the voting population. However, unlike economic externalities, the construction of the political system does not allow for these externalities to be offset by economic transfers.

3.2.1 Local government and housing in Australia

In the event that there is a correlation between one or more of the proposed demographic variables and housing growth, we would like to know whether this is partially

Figure 3.1: Sydney Greater Capital City Statistical Area — density map

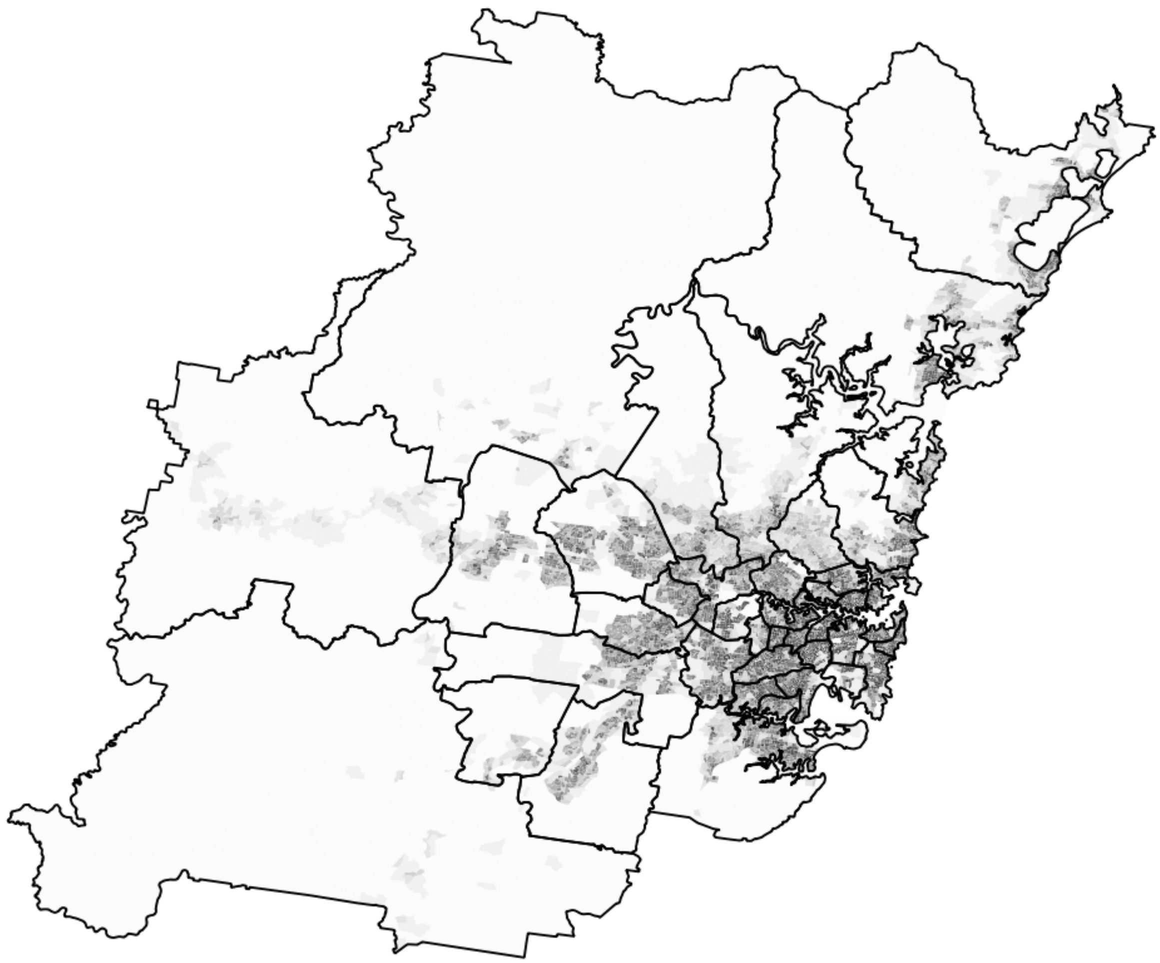

the result of the governance structure. In particular, we would like to know if decentralizing the planning process results in stronger feedback from local preferences to local outcomes. In order to test this, we need to be able to compare areas with different governance structures.

The primary responsibility for zoning and housing approvals in Australia generally rests with local governments, which are similar in form and function to US cities and counties. However, as this is the result of a devolution of power on behalf of the States, its exact form is thus slightly different in each State.[2]

[2]The Australian Constitution does not recognize the existence of local governments, and allocates all power over zoning (as a residual power) to the States, although some environmental matters of

Unlike US counties Australian local governments are in general not responsible for schooling, police, health, transportation, or fire services — these services are supplied by the State governments. Australian local governments tend to provide more localized services — local roads, local water and sewerage infrastructure[3], street lighting, and footpaths.[4]

As a result, in many cities local governments have little impact on household's locational choices. As noted in the final report of the Independent Local Government Review Panel (2013) (page 14) "On the whole, people appear satisfied with the performance of local government...However, the overall level of awareness and understanding about the role and functions of councils is quite low..." As a result, we can ignore clustering on the basis of local public good provision, which would otherwise generate endogeneity.

Nonetheless, turnout for local government elections in some parts of Australia is comparatively high. Under the *Local Government Act 1993*[5] and the *City of Brisbane Act 1924*[6] residents in Sydney and Brisbane respectively are compulsorily required to vote in local council elections or pay a fine.[7] As a result, turnout averages 83 percent in New South Wales local government elections (Australian Bureau of Statistics, 2010). Similarly, turnout for Brisbane City Council elections was 86 percent in 2008 and 82 percent in 2012 (Electoral Commission of Queensland, 2012). This substantially reduces complications from voter selection.

---

national importance are still dealt with at the federal level.

[3]Within cities with multiple local governments, local governments tend to source water from the same source, and return sewerage to the same endpoint, and thus they tend to provide only the local portion of the water and sewerage infrastructure.

[4]Brisbane City represents an unusual departure from this model, as they have historically been responsible for the entire greater capital city statistical area, and as such have partnered with the Queensland State Government to provide additional services such as transportation and major roads. However, as all the Brisbane data comes from the same local government area, the level at which these services are provided should not result in variation between areas, and should therefore not affect our results.

[5]Compulsory voting for local elections was first required under the *Local Government (Electoral Provisions) Act 1947*.

[6]This clause later fell under the *Local Government Act 2009* and the *Local Government Electoral Act 2011*, but the effect has remained consistent.

[7]Slightly different provisions exist (and have existed) for the Sydney City local council. However, since this area is removed from our data for other reasons, these variations have no effect on our results.

One key area of local government responsibility is town planning, zoning, and building approvals, although their involvement in this area is recent and evolving. The first legislation allowing (and, in fact, requiring) town plans did not appear in Australia until after WWII. New South Wales passed the *Local Government (Town and Country Planning) Amendment Act 1945,* giving their local governments the authority to create town plans and manage development[8] (Hamnett and Freestone, 1999). Thus we have a devolved governance structure that is responsible for a sizable proportion of housing development regulation, and which varies in structure from city to city.

Of course, zoning legislation is much more complicated than a simple town plan. Zoning legislation is supplemented by environmental and heritage legislation, which can be used to slow or block development at the local, State, or national level. This legislation is even more recent — New South Wales passed the first *Heritage Act* and the *Environmental Planning and Assessment Act* in 1977 and 1979 respectively, and the national government passed the *Environmental Protection and Biodiversity Conservation Act* in 2001. To complicate matters even further, exceptions are routinely given to all of the above legislation.

The increase in the volume and complexity of legislation makes it very difficult to determine either the source of the legislation or its impact on housing supply. As a result, we take an indirect approach, and examine the differences between the *results* of the different types of governance structures. By tracking the evolution of development over time, we demonstrate how the different structures and demographics lead to different development outcomes.

3.3 Supply, demographics, and governance structure

In this section we will outline a single period model of housing markets. We will then consider how owners and renters will react to various categories of regulation. Finally, we will consider how these preferences are likely to translate into outcomes

[8]Victoria and Tasmania released similar legislation at around the same time.

in a democratic political system, depending on whether city governance is centralized or decentralized.

3.3.1 Single period model

Consider a city with a single neighborhood. The utility of living in that neighborhood in a period is given by:

$$U = u\left(h, y\right) - c\left(n, G\right)$$

where h is the size of the house/land, y is consumption of other goods, G is the exogenous allocation of local public goods in the neighborhood (for example beaches etc), and n is the number of houses in the neighborhood.

Given a budget constraint of m_i, and a house price of p, we have that in any given period:

$$U = u\left(h, m - ph\right) - c\left(n, G\right)$$

We assume that utility is concave in h and y, and that the congestion function c is convex in n and concave in G. Income m is distributed in the population according to $F(m)$. We make no assumptions on $F(m)$ other than it is fixed over time, and is positive only for positive values of m.

The assumption that the utility function is additively separable in congestion is akin to stating that all public goods are equally accessible to all residents of the neighborhood, and that congestion affects all residents equally regardless of wealth. Furthermore, we assume that for all $G < \infty$, as $n \to \infty$, $c(n, G) \to 0$, or in other words as if the population is large enough an individual's value from any public goods will tend to zero. Residents and prospective residents do not recognize their own effect on the congestion function, as urban areas are characterized by a large number of dwellings per square kilometer, and each new home is likely to represent less than a 1% change in this variable.

Each period individuals choose whether to live in the neighborhood or outside the city, and whether to purchase a home at a price $P(p)$, or whether to rent a home at cost p. The nature of the neighborhood (in terms of public goods, congestion, or demographic makeup) does not impact this decision, but rather we assume it is a function of exogenous finance constraints. The area of the neighborhood is then costlessly divided into multiple parcels depending on the preferences of the individual residents, and the cost of dwelling construction is then incurred. This creates the additional constraint for each area that:

$$\int_{\tilde{m}}^{\infty} h^* (m, p) \, df(m) = A$$

Note that this means we do not distinguish between four-apartment block on a one hectare block and four single-dwelling units each on $1/4$ hectare.

Finally, if the individual lives outside the neighborhood we assume they obtain a fixed endowment of housing, and no public goods. Their utility is therefore given by:

$$U = u(H, m)$$

where H is the fixed endowment of housing.

A single-period free market equilibrium for this model exhibits monotonicity in income. Essentially, a higher-income individual will never want to live outside the neighborhood while a lower-income individual lives in it, because non-housing consumption has diminishing marginal utility in income. Thus, as income rises an individual would like to diversify into consuming housing and public goods by moving into the neighborhood.

Proposition 3.1: *If we have two individuals with income m and $\hat{m} > m$, then if individual with income m wants to live in the neighborhood then individual with income $\hat{m}$ wants to live in the neighborhood also.*

PROOF: By concavity of utility function — see appendix 3.8. ∎

We can utilize this monotonicity to demonstrate that a single-period equilibrium exists. Essentially, because the outside option is fixed, individuals will sort monotonically by income, and there will be a cutoff point above which all individuals will live in the neighborhood and below which they will live outside it.

Proposition 3.2: *There exists a pure strategy Nash equilibrium if there is only one area.*

Proof: See appendix 3.8 for full proof. Essentially we can divide the process into two stages — in the first stage individuals choose whether to live in the neighborhood, and in the second stage they maximize their utility given the price of land in the neighborhood. Given the monotonicity of the utility function we can establish a cutoff point and a price. ∎

3.3.2 Regulation and responses to growth

Having established the existence of a market equilibrium in the previous section, we can now compare it with a regulated equilibrium. If we allow residents (and only residents) to vote, they may choose to vote for regulation that limits supply. We can use our model to determine whether preferences for regulation (as opposed to the market equilibrium) vary depending whether an individual owns or rents their home, how wealthy they are, and so forth. This allows us to determine whether individuals are likely to support or oppose policies that limit growth relative to what they expect the status quo to be in the next period.

We require additional assumptions. First, and most importantly, we assume that the price of housing is perfectly predictable. Therefore if we assume a particular set of expectations regarding price and then adjust regulation, we can predict the impact of that on price and therefore establish the preferences for regulation.

We also maintain our assumption regarding the perfect divisibility of housing; in fact, for the purposes of this section we do not distinguish between housing and lot size. We also assume that housing can be costlessly re-allocated every period in any fraction. This is clearly an abstraction from reality, as in reality the existing housing stock has value from period to period. However, introducing costs of division should bias against our findings, as larger blocks are easier to divide than smaller blocks, creating a reversion to the mean across areas as growth occurs. Based on this assumption there is essentially no difference between an owner selling their property and buying another of equal size in the same area, and the owner keeping their original property.

Finally we assume that all residents of the neighborhood have the right to vote for regulation in the neighborhood, but those outside the area do not.

The menu of possible regulation is extensive, far too extensive for us to consider all possibilities. Instead we identify three restrictions that incorporate many of the most common forms regulation as subsets. First, we consider changes in the area of the neighborhood, A. This covers two common areas of regulation: staggered land release and environmental regulation. If residents oppose new land releases in their area, they are essentially requesting regulation to reduce A in the next period. Similarly, preventing home production by creating space for trees or parkland, however this is paid for, reduces A. We assume that land cannot be forcibly reclaimed from existing residential development.

Due to their aversion to congestion and the fact that home costs are sunk, existing homeowners will support regulation that restricts future land releases relative to a free market equilibrium. On the other hand, renters will prefer to allow land releases, as they will be compensated for any increased congestion through lower prices. Thus we should expect opposition to land releases in areas withe higher rates of homeownership.

Proposition 3.3: *Homeowners will support regulation that restricts any increase in A. If housing is a constant factor share of income, renters will oppose it.*

Proof: Full proof is in appendix 3.8. Essentially, increased congestion would decrease both utility and prices, and thus be negative for owners as they have sunk costs. Renters, on the other hand, will be fully compensated through the price mechanism.

Note that housing being less than a constant factor share of income is not sufficient to render this untrue, but the outcome would then be dependent on the exact form of the utility function. ■

Second, we consider a blanket restriction on n, or limiting the number of houses in an area. Essentially, the government indicates that no new homes may be built unless an equivalent number of homes are destroyed. It is similar to stating that all homes must be detached single-family dwellings, although not entirely. Once again, we assume that the area cannot 'go backwards'; that is, we cannot forcibly remove anyone from the area, only prevent new entry.

In this case, our restrictions are no longer both binding — in order to restrict the number of people in the area prices must rise, but then we can no longer fill the area as at the higher price residents will want to reduce the size of their houses. Homeowners will therefore support this policy only if there is a third party to purchase any excess land. Renters will once again prefer the free market equilibrium.

Proposition 3.4: *Homeowners will support regulation that restricts any increase in n so long as there is a third party to purchase any excess land. If housing is a constant factor share of income, renters will oppose it.*

Proof: Full proof is in appendix 3.8. Essentially, the negative impact of congestion makes homeowners who wish to stay oppose development, and the price effect makes those who wish to leave oppose development.

Note that housing being less than a constant factor share of income is not sufficient to render this untrue, but the outcome would then be dependent on the exact form of the utility function. ■

The final policy limits the smallest possible size of homes, or in our model sets a minimum $\underline{h}$. This is perhaps the most well known, and is most commonly seen in height and block size restrictions. This also covers a great deal of heritage legislation which prevents existing homes from being divided into smaller dwellings.

In this case, the impact on prices is mixed. Owners who wish to stay will obviously prefer to reduce congestion. Furthermore, since they do not have to leave if the minimum exceeds their current block size, they will (unlike renters) also not oppose the policy due to concerns about having to leave the area. However, the effects of such a policy on price are less clear-cut.

Proposition 3.5: *Independent of the effects on price, owners are more likely than renters to support minimum size regulations. The effect on price depends on the nature of the distribution, the relative concavity of the utility function, the value of the outside option, and the degree of impact of congestion. Higher prices in response to a minimum size regulation will result in owners supporting the policy, lower prices will result in renters supporting the policy.*

PROOF: Full proof is in appendix 3.8. ∎

We would argue that the value of the outside option is increasing less rapidly with income than we have modeled here (as one is likely to lose one's city job if one does not live in the city), which increases our chance of a price increase in response to a minimum size regulation. Furthermore, while here we have allowed owners to 'pool resources' to extract the optimal number of new housing units, in reality these are likely to be spread across the city, thus meaning that instead of two houses being turned into three we must have one house remaining the same size and the other split into two. This reduces the positive impact on price for any one owner of allowing the group to divide homes into smaller units, as it forces departing owners to compete with each other to attract new demand. Finally, we would argue that in fact housing demand is pretty price inelastic, especially as one moves down the income distribution, thus supporting the view that prices are likely to rise as minimum lot sizes increase.

Overall, these three propositions suggest that a given neighborhood with more owner-occupiers relative to renters will have a political preference for more restrictive regulation. We therefore predict that a larger number of owners relative to renters will reduce the quantity of new houses produced in the neighborhood in the next period.

In contrast, wealth alone also does not necessarily lead a resident to prefer more restrictive regulation. On the one hand, wealthy individuals would prefer to pay more than less wealthy individuals for less congestion, as the utility from congestion is fixed whereas utility from housing and other goods are declining in income. On the other hand, if congestion is higher in poorer neighborhoods, there may be a larger return to the expenditure of resources to reduce congestion (depending on the marginal value of decreased congestion).

For those owners who are selling, the incentives are equally mixed. While a wealthy owner has more to gain from a price increase in absolute terms, they also have a lower marginal utility from increasing their wealth, so they may care less. Furthermore, if they are divided into areas where they are consuming more public goods but not larger houses, giving up a fixed y for a larger p may not result in much change. Therefore our model (in contrast with some of the literature) does not unambiguously predict that regulation will be tighter in neighborhoods that are wealthier.

Finally, we need to consider age. Like wealth, the effects of age may run in two conflicting directions. On the one hand, the expected length of future tenure in the home is likely to decline with age. On the other hand, for some elderly residents their home represents the bulk of their assets, and they have limited future labor income. Further, they may have a strong bequest motive, and wish to maximize the value of the house they pass to their heirs. As a result, they may be more concerned about price effects and may be more supportive of regulation. Furthermore, the elderly may have stronger preferences against congestion (through a different c function), for historical or age-related reasons.

In either case, as noted in section 3.2, the there is considerable evidence that the elderly have higher rates of civic participation. Thus their preferences may be more represented. Overall, however, our model does not predict the direction of those preferences.

3.3.3 Decentralization

Even were our data to support our predictions, a single neighborhood city such as described above would not allow us to identify whether the relationships were the result of supply factors or demand factors as theorized. We would therefore like to exploit a comparison between the impact of those variables on multiple-neighborhood cities with different governance structures.

The first step is to ensure that an equilibrium exists for a city with multiple neighborhoods. Obviously any given equilibrium may not be unique, as households are interested in the interaction between public goods and congestion. If we have two areas we may have that the area with more public goods has higher congestion and a lower price, or lower congestion and a higher price.

We can simplify this problem by noting that our previous results on monotonicity extend naturally to the case where we have multiple neighborhoods, since individuals with higher incomes will prefer to live in less dense and/or higher endowment areas. Furthermore, the neighborhoods where the wealthy live will have higher prices per unit of land than areas where the poor live (proposition 3.6). Therefore individuals will generally sort themselves according to income, except for the set of knife-edge cases where two or more areas have differences in congestion that completely balance out the difference between their public good endowments (proposition 3.7).

Proposition 3.6: *If we have two individuals with incomes m and $\hat{m} > m$, and individual m strictly (weakly) prefers neighborhood B to neighborhood A while individual with income $\hat{m}$ strictly (weakly) prefers the reverse, then we must have that $p_A > p_B$ $(p_A = p_B)$.*

Proof: By concavity of utility function — see appendix 3.8. ∎

Proposition 3.7: *If we have two individuals with income m and $\tilde{m} > m$, and individual m strictly (weakly) prefers A to B while individual $\tilde{m}$ strictly (weakly) prefers B to A, then any individual with income $\hat{m} > m$ strictly (weakly) prefers neighborhood A to neighborhood B.*

PROOF: By proposition 3.6, we know that if m strictly (weakly) prefers A to B and individual $\tilde{m}$ strictly (weakly) prefers B to A, then $p_A > p_B$ ($p_A = p_B$). But by the same argument, if $\hat{m} \geq m$ prefers B to A, we must have that $p_B > p_A$, which would be a contradiction. So we must have that individual $\hat{m}$ prefers A to B. ■

Proposition 3.8 demonstrates that we can therefore use the same sorting mechanism we outlined in propositions 3.6 and 3.7 to construct at least one equilibrium.

Proposition 3.8: *For a city with multiple areas, there is at least one separating equilibrium for $n > k$, and at most one pooling equilibrium.*

PROOF: See appendix 3.8. ■

We have shown in the previous section that some types of voters prefer a regulated equilibrium to a free market equilibrium, whereas others do not. These preferences are defined by their demographic characteristics. So far, however, we have only looked at the case where the city is composed of one neighborhood. We need to know if the same preferences result if the city is divided into multiple neighborhoods.

We argue that if the demographic characteristics of voters affect their taste for regulation, then we would expect a stronger regulatory response at the level at which regulation is set. Essentially, if a neighborhood does not have the jurisdiction to create its own housing regulation in our model, the number of owners versus renters (or wealthy vs poor, or old vs young) does not matter. Since everyone in the whole city can vote, the city as a whole may exhibit these same relationships[9], but the pivotal voter will never be in any given neighborhood. Since all neighborhoods prefer to reduce congestion in *their* neighborhood, they will prefer to relocate development away from themselves, and into other neighborhoods.

Furthermore, if we consider the multiple area case, prices will tend to need to adjust by more in response to regulation, as reducing land availability, increasing minimum housing size, and/or limiting the number of residents creates congestion in

[9]Voters may have preferences against congestion at multiple levels — neighborhood and city as a whole. But preferences for reduced congestion from the city as a whole should not cause noticeable variation between areas.

other neighborhoods, reducing the relative attractiveness of the outside option as the outside option.

Therefore, we would expect our predicted relationships to hold in neighborhoods where governance is decentralized to the neighborhood level, but not if regulation is set more centrally.

3.4 Data and empirical strategy

3.4.1 Data

We selected data for local government and statistical local areas in the Australian cities of Sydney and Brisbane. The Australian Bureau of Statistics (ABS) collected the data as part of the Australian Census program in 1933, 1947, 1954, and then in 5-yearly intervals from 1961 to 2011. We obtained data from 1933 to 1961 from the original ABS census publications. The Australian Data archive provided us with most of the data for the period 1966-1991 in electronic form, translated from the original microfiche. Some demographic information from the 1986 census was translated directly by us from the microfiche at the University of Queensland archives. We obtained the remaining data from the ABS online census portal.

From the data, we selected the two larger cities for analysis: Brisbane and Sydney. We chose these cities for several reasons.

- They are two of the three largest cities by population in Australia.
- Both cities have compulsory voting at the local government level, and very similar rates of voter turnout.
- Both cities are situated on the east coast on natural harbors.
- Both cities and are limited spatially by mountain ranges to the west.
- Both were originally founded as British convict colonies.

Australia's population is concentrated, with around 60 percent of people living in just five major cities in 2014 (calculated from Australian Bureau of Statistics, 2015). While we may also wish to include Australia's second-largest city (Melbourne) in our analysis, this city went through a complete restructuring of local government boundaries in 1993, resulting in an insufficient quantity of data to extract robust results. Australia's fourth and fifth largest cities, Perth and Adelaide, do not have compulsory voting at the local government level, and has low turnout rates (33 percent in the former case) so they were also not considered (Australian Bureau of Statistics, 2010). The remaining Australian cities are much smaller, housing less than 3 percent of the Australian population each (Australian Bureau of Statistics, 2015).

The boundaries of Sydney are based on the greater capital city statistical areas for 2011, which are determined by the ABS based on economic linkages and activity. Sydney is composed of several local government areas (LGAs) which are separately governed, and there is no central city-wide authority. Data for these cities is therefore reported on a local government basis in each census. The exact number of local government areas varies over time, but was 43 in 2011. Brisbane, on the other hand, has historically been governed by a single local government[10], and because of this the ABS reports Brisbane data by suburban areas (termed statistical local areas or SLAs), of which there were over 100 in 2011.

The ABS also provided historical maps for the years 1981-2011, which we used to calculate the total area of the SLAs and LGAs in each period. We also used these maps to determine which areas had been merged and divided over time. Every attempt has been made to ensure long-term consistency in the geographical areas examined, and more detail can be found in appendix 3.9.

From the census, we selected key demographic and dwellings variables. These included the number and type of dwellings (exact definitions and number of categories varied from census to census, but included houses, flats, semi-detached dwellings, villas, etc). Other variables included age (by 5-year intervals), and tenancy arrange-

[10]By 2011, the greater capital city statistical area for Brisbane included some areas outside that were outside the Brisbane local government area. However, this is, for the most part, a handful of suburbs in neighboring local government areas.

ments (divided into owned outright, owned and mortgaged, rented privately, and rented from a government authority).

The ABS did not collect income data until 1976, which is also the year that the SLA boundaries for Brisbane were stabilized. We therefore limited our analysis to 1976 onward. Unfortunately, the income data bins (intervals) were inconsistent, in that they did not represent a fixed fraction of households in each year. We therefore created a new bin that represents as close to the top 15 percent of households in the city (we also replicated this for families and individuals) as possible. The same data was used to calculate the median income in each area in each year.

We calculated similar variables using binned rent and mortgage data, for the years in which these were available (1933 onward, excluding 1966 and 1971 for rent data and 1976 onward for mortgage data). Once again, the variable contains some noise, due to variation in the bin cutoffs.

Construction costs were calculated using ABS data. Each city has its own construction cost index published quarterly by the ABS. We took the average of this construction index for the 5-year period preceding each census, and then discounted it by the relevant (again city-based) price index. This measures the real variation in building costs over time.

One concern with our model is that at very low levels, congestion preferences may in fact be positive. This is partly the result of an interaction effect between the availability of basic infrastructure (such as street-lighting, roads, and public water supply) and the number of people. This makes fringe suburb growth more difficult to predict. In addition, new greenfields space may be partially controlled by the State government, particularly around the edges of the city, and even once sold large estates may be held for long periods by developers in order to generate monopoly rents. We therefore limit our analysis to areas which are 'established' suburban areas. As a proxy for 'established', we use Sydney LGAs that are designated as part of Sydney city in 1911, assuming that if significant parts of those areas were urban in 1911 they would be entirely urban by 1976. For Brisbane we take a later profile — since the SLAs are much smaller we assume that these are almost entirely urban by the time they are established by the ABS as separate SLAs, and we therefore consider everything that was a separate SLA in 1954.[11]

3.4.2 Empirical strategy

Our empirical strategy exploits three types of variation. Firstly, there is variation in the demographic variables across neighborhoods within the same city. Secondly, there is variation within neighborhoods over time. And finally, there is variation between neighborhoods in cities with different governance structures. This allows us to determine both which demographic variables are likely to result in supply restrictions and under which political structure those variables are important.

Growth can be measured either in terms of the percentage change in the number of houses or in the change in housing density. While this changes the interpretation of the coefficients, the two are mathematically equivalent (for proof see appendix 3.8). We prefer the density measure, as it automatically adjusts for the size of the local neighborhood.

In growth terms, therefore, we are interested in testing:

$$Density - l.Density = \alpha + X\beta + \beta_6 ownership + \beta_7 high\_income + \beta_8 median\_income + \beta_9 pensionage + \epsilon$$

where any combination of $\beta_6, \beta_7, \beta_8, \beta_9 < 0$ are significant.

Some of these variables are likely to be correlated with each other. For example, the percentage of people with high income is positively correlated with the median income. As regards the remaining the remaining variables, homeownership rates are slightly positively correlated with median income (0.28) and slightly negatively correlated with the proportion of elderly residents (-0.28). Median income is slightly negatively correlated (-0.28) with the proportion of elderly residents.[12] Therefore we consider both the fully specified model and the partially specified model for each of

[11] In general we find that those suburbs had a density greater than 4 houses per hectare in 1976, whereas most of the remainder had a density of less than 2 houses per hectare.

[12] These statistics are calculated from a mix of the Brisbane and Sydney data, limited to suburban areas.

the key variables of interested.

So far we have not distinguished the impact of these variables on demand versus their impact on supply (via regulation). To do so, we exploit the variation between the two governance structures. In the case of just two cities, this can be easily tested by including an interaction term as follows:

$$Density - l.Density = \alpha + X\beta + \gamma_1 city + \gamma_2 demographic + \gamma_3 city\_demographic + \epsilon$$

If we compare two different cities, one where the lowest level of regulation is set at the neighborhood level, and another where it is set at the city-wide level, we would expect the impact of the demographic variables to be much larger in the former case.

In order to test the above relationships, we run ordinary least squares on our panel dataset.[13] We allow for fixed effects for local government area and statistical local area, and also for year.[14] We cluster our standard errors by local government area or statistical local area.

3.5 Results

The stock of housing does not remain constant over time — there is some divergence. This can be seen from figure 3.2, which shows the growth trajectory for four Sydney suburbs since 1966. There is some difference in the variation across the two cities, as can be seen from figure 3.3. Sydney has faster density growth in some areas, which creates a bias against our hypothesis.

In the absence of political economic factors, we would expect housing growth to respond in a theoretically consistent manner to market-based supply and demand fac-

[13] We also considered using a matching methodology. This involves matching areas in Brisbane and Sydney on area, density and year, and then running ordinary least squares on the Sydney, ownership and interaction variables, including a binary variable to take into account the fixed effects of each pair. We used wards instead of statistical local areas, to ensure there was common support on the size of the area. However, matching produced no improvement in the balance statistics. Also, matching even with a relatively large caliper dramatically reduced the size of the data set (down to 304 observations). Consequently the results have not been included.

[14] We considered running each year as a separate cross section, but this limited us to too few data points, especially for Sydney.

Figure 3.2: Housing density of four Sydney LGAs over time

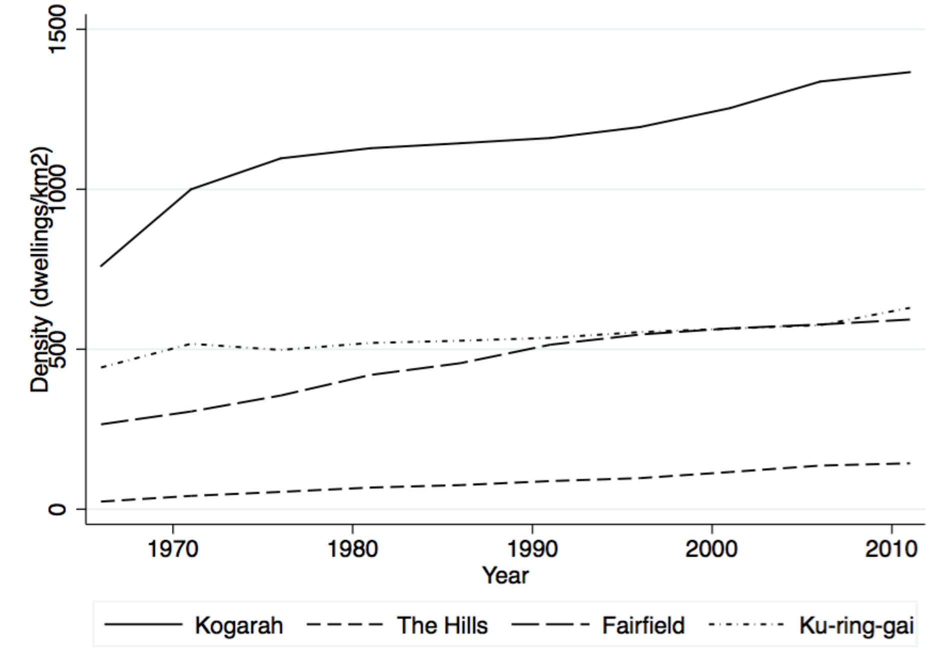

Figure 3.3: Change in density in Brisbane statistical local areas and Sydney local government areas

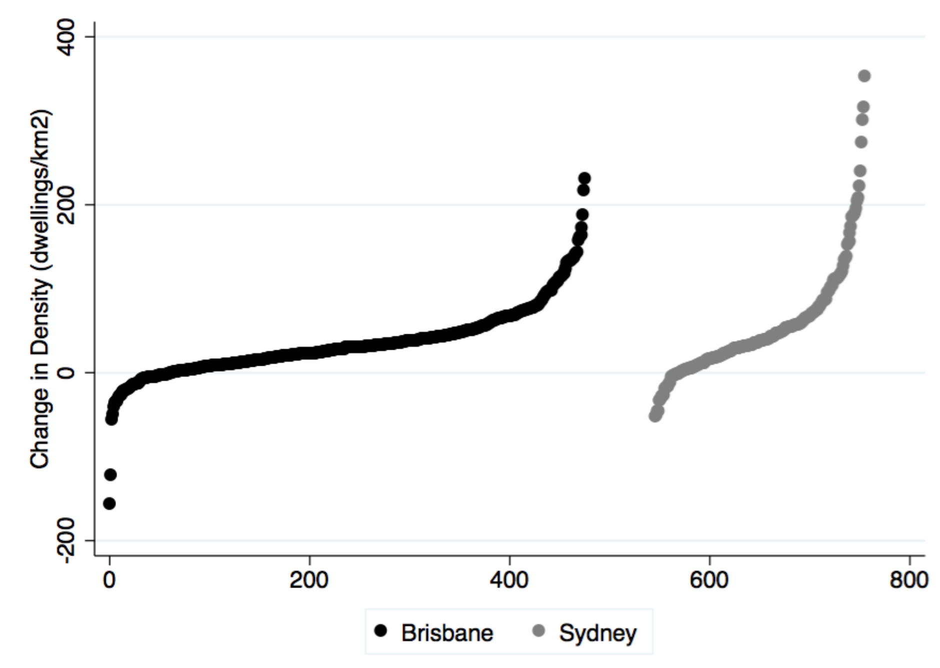

tors. For instance, new housing should increase with population growth and decline with construction costs. In fact, we find that density growth is higher when construction costs are higher, and population changes have zero significant impact on housing growth (see table 3.1). Obviously if housing is not positively responsive to population growth, one would expect higher prices and housing 'shortages' to result[15].

[15]When most people discuss housing shortages they are not actually referring to an excess of demand at the current market price, rather they are referring to the unavailability of housing at a sufficiently low price point to be sustainably affordable to lower income households.

Table 3.1: Responsiveness of housing supply to population and construction costs

	(1)
	Density - L1.Density
L1.Density (dwellings/km^2)	-0.144***
	(0.0356)
Area (km^2)	-35.00***
	(12.08)
Population - L1.Population	-15.14
	(13.46)
Construction Cost Index	3.984***
	(1.385)
Observations	638

Note: Sydney data is based on local government areas (LGAs), Brisbane data is based on statistical local areas (SLAs). Standard errors are clustered by LGA or SLA, and are presented in parentheses in the table. Regressions include fixed effects for each LGA and SLA and year. L1 refers to the first lagged value of the variable. x denotes an interaction effect. ** $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$*

Table 3.2 presents the correlations between growth and various demographic characteristics across the two types of jurisdiction. In accordance with our theory, we find that home ownership rates reduce housing growth in Sydney suburbs (decentralized system), but not in Brisbane suburbs (centralized system). A 1 percent increase in ownership rates across Sydney results in around 20,000 fewer dwellings than would otherwise be built every five years. Over a 40 year period, that is 160,000 dwellings, enough to house 1.4 percent of Sydney's current population on a two-person-per-dwelling basis (and almost sufficient to compensate for the current predicted 'housing shortage'). If instead we consider a 5% drop (the difference between aggregate Australian and US rates) we would expect to see supply growth of 800,000 dwellings, or around 7% of the current population.

Of course, the results suggest that home ownership is not the only factor in play. Our results also suggest that the same result would not necessarily occur in a centralized political structure. Setting approvals at the level of the city as a whole, rather than at the neighborhood level, appears to completely eliminate the negative relationship between home ownership and growth. Even when we consider Brisbane 'wards'[16] (similar to parliamentary constituencies) instead of statistical local areas, we cannot recover the relationship that we find for Sydney.

The relationship between ownership and lower growth applies to detached single-family housing as well as high-rise. However, at the outermost fringes of the cities we find that the relationship starts to break down. This is not surprising, as a lot of the land on the city fringes is still owned by the State (rather than local) governments, who choose whether to release it.

On the other hand if we examine the data at the block level we find that ownership is negatively correlated with growth regardless of the jurisdictional structure, but we suggest that this is misleading. When dealing with such small spaces (sometimes as small as a city block) that micro factors become more important. For instance if one is in the process of building new houses one will clearly have fewer owner-occupiers in residence.

In addition to ownership, we are also interested in the effects of income and income inequality on housing supply. We consider whether the spatial differences in housing growth are driven by spatial differences in income in two ways. Firstly, we look at whether development levels are affected by average incomes. Secondly, we consider whether spatial differences are affected by the proportion of residents who

[16] Brisbane elects one councillor per ward, these together form the council. The mayor is elected separately. Wards are considerably larger than statistical local areas.

are wealthy (in the top 15% of incomes). We find that neither average incomes nor inequality appear to play a role in reducing development. Contrary to the theory, income and/or an increase in the number of relatively wealthy residents appears to be, if anything, positively related to development in a neighborhood. In addition to the results presented here, we also extended our analysis back to 1954 by using rent as a proxy for income (the correlation between income and rent from the years in which both sets of data are available is .8). Since the correlation is not clearly sensitive to jurisdiction, we cannot ascertain whether this is related to supply or demand factors.

One counterargument that could be raised is that income is a poor proxy for wealth, and that it is wealth that really matters. Even if we had household assets at this level of detail, however, we would be reluctant to use them. The main residence usually represents a large proportion of household assets. Since home values are based on local house prices which are in turn a function of demand and supply, the analysis would be circular. Instead, we split households into those who have repaid their mortgage (which we assume are on average wealthier) and those who have not. We find that more outright owners does not result in slower growth.

Finally, we find that an increase in the number of elderly residents (over 65) has a substantial negative impact on future housing growth. This is true regardless of the jurisdictional structure. Once again we cannot distinguish whether this is driven by supply or demand factors, but we do note that higher proportions of elderly residents are also correlated with higher property taxes.

Table 3.2: Impact of homeownership rates, median income, inequality and proportion of elderly residents on changes in housing density

	(1) Dens-L1.Dens	(2) Dens-L1.Dens	(3) Dens-L1.Dens	(4) Dens-L1.Dens	(5) Dens-L1.Dens
L1.Density (dwellings/km^2)	-0.140***	-0.134***	-0.141***	-0.155***	-0.162***
	(0.038)	(0.031)	(0.033)	(0.034)	(0.035)
Area (km^2)	-34.31***	-30.90**	-31.50**	-29.69**	-30.34**
	(12.08)	(12.26)	(12.94)	(13.26)	(12.68)
L1.Homeowners (%)	0.501				-0.314
	(0.720)				(0.623)
L1.Homeowners x Sydney	-1.796**				-1.986***
	(0.770)				(0.707)
L1.Top Earners (%)		0.308			
		(0.601)			
L1.Top Earners x Sydney		0.308			
		(0.779)			
L1.Median Income ($'000)			0.495**		0.225
			(0.232)		(0.261)
L1.Median Income x Sydney			-0.019		0.117
			(0.294)		(0.277)
L1.Over 65 (%)				-1.679***	-1.491***
				(0.470)	(0.501)
L1.Over 65 x Sydney				-3.676	-4.298*
				(2.387)	(2.192)
Observations	638	638	638	638	638

Note: Sydney data is based on local government areas (LGAs), Brisbane data is based on statistical local areas (SLAs). Standard errors are clustered by LGA or SLA, and are presented in parentheses in the table. Regressions include fixed effects for each LGA and SLA and for each year. Sydney fixed effects are absorbed by LGA fixed effects. L1 refers to the first lagged value of the variable. x denotes an interaction effect.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

3.6 Conclusion

A high rate of home ownership in any given period is not unambiguously good. In fact, one side effect of increasing home ownership now is a considerable reduction in future housing supply, and associated higher housing prices. We find that a 10 percentage-point decrease in home ownership rates over our period would result in an increase of around 1.6 million dwellings in Sydney, enough to house 14% of the 2011 population on a two-person-per-dwelling basis. Without additional forms of intervention, this interaction may make high rates of home ownership unsustainable in the longer term, as higher prices and limited access to finance reduce housing affordability.

Second, while decentralization allows local areas to more accurately gauge and enforce the preferences of their citizens, local districts are not completely independent and these preferences may impose costly externalities on residents of other districts. Local governments form clubs which may choose to prevent development in their area, even if it is supported by the majority of the citizens in the city as a whole. These concerns should be taken into account when considering the role of local governments moving forward.

Our results offer an avenue towards uniting two conflicting strands of literature. It has been argued by one school of thought that the causes of housing supply limitations are geographic. On the other hand, there is evidence to suggest that they are regulatory. Rather than considering the city as a whole, we look at the impact of demographic variables at the level of the regulatory body. If local regulatory bodies look to restrict development in their area to prevent congestion, we may still see significant growth in housing across the city as a whole, if the city is able to expand outwards, as regional areas with low density may encounter fewer limitations from a political economy perspective. This provides a possible avenue for future research.

While we primarily consider voting as the mechanism for amalgamating local political preferences, we reject the idea that wealthy residents directly increase regulation through lobbying. We find the presence of more high income earners tends to be, if anything, positively associated with growth. However, we do not look at the relative impact of home owner associations or the possibilities for corruption at the local government level, we leave these extensions for further work.

We also find that a large proportion of elderly residents is strongly correlated with lower growth. This result is concerning, as much of the western world is experiencing rapid aging. As we do not have any significant variation in this result across the two political structures we are considering, we cannot determine whether this is driven by supply factors (such as regulation). Lower growth in areas with more elderly people may simply be the result of lower demand from outsiders or lower mobility in response to price changes from residents. However, we consider this to be very relevant, and it should be considered more throughly in future research.

3.7 Appendix: Robustness

This appendix includes a range of additional robustness checks. These include adjustments for missing data, adjustments to reflect variance, adjustments to time periods, and others.

Rent

As noted previously, income data is limited to the 1976-2011 period. However, rent is highly correlated with income over this period (the correlation between being the proportion of rent payers in the top 15 percent and the proportion of income earners in the top 15 percent is around .8). Therefore, we can use rent as a proxy for high income.

As table 3.3 outlines, even if we extend the analysis back to 1954, places with a high proportion of top rent payers have faster growth. This supports our existing results that income is not linked to supply restrictions. Furthermore, the effect is stronger for Brisbane, suggesting that it is in fact the result of the response of supply to higher demand rather than to political pressures.

Table 3.3: Impact of proportion of flats with high rent on changes in housing density

	(1) Dens-L1.Dens (Sydney)	(2) Dens-L1.Dens (Brisbane)	(3) Dens-L1.Dens (Joint)
L1.Density (dwellings/km^2)	-0.170**	-0.143***	-0.150***
	(0.066)	(0.041)	(0.034)
Area (km^2)	-26.33***	-31.54	-25.84***
	(7.396)	(22.98)	(9.401)
L1.High Rent (%)	68.72	107.2**	111.7**
	(69.87)	(49.82)	(45.30)
L1.High Rent x Sydney			-24.50
			(78.29)
Observations	138	408	546

Note: Sydney data is based on local government areas (LGAs), Brisbane data is based on statistical local areas (SLAs). Standard errors are clustered by LGA or SLA, and are presented in parentheses in the table. Regressions include fixed effects for each LGA and SLA and for each year. Sydney fixed effects are absorbed by LGA fixed effects. L1 refers to the first lagged value of the variable. x denotes an interaction effect.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Wealth, income and ownership

It could be argued that income is not the key variable in determining the impact of inequality — rather, wealth is the main variable of interest. Income and wealth are often poorly correlated, and it is especially difficult to use income as a proxy for wealth when income data is not stratified by age. However, it is also very difficult to separate wealth from home ownership, since one's home often represents a large proportion of one's wealth. We therefore make the assumption that owner-occupiers with a mortgage are less wealthy than those without a mortgage on average, and compare the two.[17]

However, as table 3.4 demonstrates, we actually get the opposite results. In fact, the proportion of mortgages seems more negatively correlated than outright ownership. A Wald test indicates no difference between the joint coefficients for mortgagees compared with outright owner-occupiers. However, once the variables are interacted with the Sydney dummy they are significantly different.

[17]For completeness, outright ownership rates are negatively correlated with the proportion of elderly residents.

Table 3.4: Impact of proportion of homeowners with a mortgage and the proportion of outright owners (no mortgage) on changes in housing density

	(1)
	Density - L1.Density
L1.Density(dwellings/km^2)	-0.134***
	(0.038)
Area (km^2)	-34.43***
	(12.03)
L1.Mortgage (%)	0.720
	(0.719)
L1.Owned outright (%)	0.206
	(0.744)
L1.Mortgage x Sydney	-2.101***
	(0.764)
L1.Owned Outright x Sydney	1.832**
	(0.900)
Observations	638

Note: Sydney data is based on local government areas (LGAs), Brisbane data is based on statistical local areas (SLAs). Standard errors are clustered by LGA or SLA, and are presented in parentheses in the table. Regressions include fixed effects for each LGA and SLA and for each year. Sydney fixed effects are absorbed by LGA fixed effects. L1 refers to the first lagged value of the variable. x denotes an interaction effect.

** $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$*

High density vs detached housing

It could be argued that in fact people are only averse to high density housing (such as multi-story apartments), rather than housing as a whole. We therefore ran two additional regressions, separating out high density housing (apartments) from detached housing. Per the results presented in table 3.5, we find that ownership results in less growth in either form of housing, and again that this result is only significant for Sydney.

Table 3.5: Impact of homeownership rates on changes in housing density on high-rise dwellings and detached dwellings

	(1)	(2)
	Highrise - L1.Highrise	House - L1.House
L1.Density (dwellings/km^2)	-0.012***	-0.026
	(0.003)	(0.019)
Area (km^2)	-1.167*	-1.797
	(0.624)	(2.299)
L1.Homeowners (%)	0.041	-0.178
	(0.063)	(0.181)
L1.Homeowners x Sydney	-0.369***	-0.304
	(0.073)	(0.202)
Observations	638	638

Note: Sydney data is based on local government areas (LGAs), Brisbane data is based on statistical local areas (SLAs). Standard errors are clustered by LGA or SLA, and are presented in parentheses in the table. Regressions include fixed effects for each LGA and SLA and for each year. Sydney fixed effects are absorbed by LGA fixed effects. L1 refers to the first lagged value of the variable. x denotes an interaction effect. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Fringe vs established suburban growth

As noted above, one concern with our model is that it may not apply as well to neighborhoods at the city fringe. We therefore limited our analysis to suburban areas. The regressions in table 3.6 show the same analysis for fringe suburbs, and we find that the relationship does in fact disappear as anticipated.

Table 3.6: Impact of homeownership rates on changes in housing density in fringe suburbs

	(1) Dens-L1.Dens (Sydney)	(2) Dens-L1.Dens (Brisbane)	(3) Dens-L1.Dens (Joint)
L1.Density (dwellings/km^2)	-0.373*	-0.154***	-0.162***
	(0.210)	(0.0475)	(0.0474)
Area (km^2)	-1.028*	-1.511	-0.929***
	(0.528)	(4.751)	(0.348)
L1.Homeowners (%)	-0.870	0.793	0.391
	(1.195)	(0.692)	(0.658)
L1.Homeowners x Sydney			0.705
			(1.045)
Observations	128	273	401

Note: Sydney data is based on local government areas (LGAs), Brisbane data is based on statistical local areas (SLAs). Standard errors are clustered by LGA or SLA, and are presented in parentheses in the table. Regressions include fixed effects for each LGA and SLA and for each year. Sydney fixed effects are absorbed by LGA fixed effects. L1 refers to the first lagged value of the variable. x denotes an interaction effect.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Rates and approval delays

One way in which local governments can reduce development is simply by delaying approvals. We therefore considered the number of days to approval (available for Sydney in 2011) and compared that with our independent variables in 2006 and 2011. We don't find evidence for or against approval delays (see table 3.7), but we do find evidence that demographic factors impact rates. Median income of residents appears to have no impact on rates at all, but places with high proportions of elderly have higher rates on average.

Table 3.7: Impact of median income, age and homeownership rates on average residential rates and approval delays

	(1)	(2)
	Rates ($)	Approval Time (days)
Median Income ($'000)	0.743	-0.004
	(1.68)	(0.186)
Over 65 (%)	41.06**	1.342
	(16.91)	(1.876)
Homeowners (%)	-9.047	-0.336
	(6.643)	(0.737)
Observations	22	22

Note: Sydney data is based on local government areas (LGAs). Standard errors are clustered by LGA, and are presented in parentheses in the table. x denotes an interaction effect.

** p < 0.10, ** p < 0.05, *** p < 0.01*

Binned Ownership

One possible area of concern is that ownership rates may be generally negatively correlated with density growth, but that the relationship may not be linear. An approach to dealing with this is to divide ownership into intervals called bins. All coefficients for ownership interacted with Sydney are negative (see table 3.8). The only significant variable ownership rates between 70 and 80 percent; however, there are also relatively few observations in each ownership category.

Micro areas

Given our argument that the difference between cities is due to different jurisdictions, we wish to check that when we limit our analysis to the block level, we don't find any difference between Brisbane and Sydney, and the demographic variables we've identified should cease to matter. We find that the effect does not go away when we shrink our analysis to the block level (table 3.9), although we believe this is due to different factors operating at such a small level.

Table 3.8: Impact of homeownership rates on changes in housing density using homeownership bins

	(1) Density - L1.Density
L1.Density (dwellings/km^2)	-0.155***
	(0.039)
Area (km^2)	-34.12***
	(12.41)
L1.Homeowners 50%-60%	-11.06
	(9.476)
L1.Homeowners 60%-70%	-2.930
	(11.84)
L1.Homeowners 70%-80%	12.84
	(14.92)
L1.Homeowners 80%-100%	24.49
	(19.83)
L1.Homeowners 50%-60% x Sydney	-3.366
	(10.33)
L1.Homeowners 60%-70% x Sydney	-19.82
	(14.15)
L1.Homeowners 70%-80% x Sydney	-34.07*
	(18.06)
L1.Homeowners 80%-100% x Sydney	-10.39
	(19.34)
Observations	638

Note: Sydney data is based on local government areas (LGAs), Brisbane data is based on statistical local areas (SLAs). Standard errors are clustered by LGA or SLA, and are presented in parentheses in the table. Regressions include fixed effects for each LGA and SLA and for each year. Sydney fixed effects are absorbed by LGA fixed effects. L1 refers to the first lagged value of the variable. x denotes an interaction effect.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3.9: Impact of homeownership rates on changes in housing density at block level

	(1)	(2)	(3)	(4)
	Dens-L1.Dens	Dens- L1.Dens	Dens- L1.Dens	Dens- L1.Dens
L1.Density (dwellings/1000m^2)	0.051	-0.055	0.004	0.050
	(0.043)	(0.041)	(0.045)	(0.042)
Area ('000m^2)	-16495.8*	-820524.1***	-16532.4***	-16840.3*
	(8865.5)	(201033.7)	(4528.9)	(8940.4)
L1.Homeowners (%)	-1028.3***	-363.0***	-1229.6***	-203.8*
	(256.8)	(98.33)	(425.9)	(105.3)
Sydney				61891.6***
				(17422.9)
L1.Homeowners (%) Sydney				-836.1***
				(231.1)
Observations	1111	338	1262	1449

Note: Data is based on collection districts, which are small areas that the Australia Bureau of Statistics reported data for in 2006. These districts were renamed S1 districts in 2011 and boundaries were redrawn. Observations include only those areas for which the boundaries did not change. Sydney represents the fixed effect for Sydney. There are no other fixed effects as data is only available for 2006 and 2011. L1 refers to the first lagged value of the variable. x denotes an interaction effect. Area is in thousands of metres squared not km squared per the other regressions, and this also applies to density.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Larger Brisbane areas

One final thing we may be concerned about is that the areas in Sydney are very much larger than the areas we are looking at in Brisbane. Furthermore, Brisbane is divided into 'wards', which may approximate the Sydney areas. We therefore look at the equivalent regression using approximations of Brisbane's Wards, rather than its statistical districts (table 3.10). While the significance shrinks, likely due to our smaller quantity of data, the level of the effect is approximately the same.

Table 3.10: Impact of homeownership rates, median income, inequality and proportion of elderly residents on changes in housing density using wards instead of statistical local areas for Brisbane

	(1) Dens-L1.Dens	(2) Dens-L1.Dens	(3) Dens-L1.Dens	(4) Dens-L1.Dens	(5) Dens-L1.Dens
L1.Density (dwellings/km^2)	-0.414**	-0.397***	-0.409***	-0.427***	-0.450***
	(0.175)	(0.145)	(0.145)	(0.140)	(0.155)
Area (km^2)	-10.15**	-10.12**	-10.08**	-10.21**	-10.11**
	(4.128)	(4.049)	(4.107)	(3.978)	(4.042)
L1.Homeowners (%)	-1.205				-1.997
	(2.824)				(2.352)
L1.Homeowners x Sydney	-1.698*				-2.018*
	(1.002)				(1.015)
L1.Top Earners (%)		1.680			
		(1.612)			
L1.Top Earners x Sydney		0.993			
		(1.101)			
L1.Median Income ($'000)			-0.226		0.112
			(0.524)		(0.459)
L1.Median Income x Sydney			0.866*		0.491
			(0.513)		(0.400)
L1.Over 65 (%)				-0.355	0.459
				(1.913)	(1.925)
L1.Over 65 x Sydney				-9.222**	-9.844**
				(4.220)	(4.356)
Observations	304	304	304	304	304

Note: Sydney data is based on local government areas (LGAs), Brisbane data is based on council wards, which have been calculated by merging statistical local areas. Standard errors are clustered by LGA or ward, and are presented in parentheses in the table. Regressions include fixed effects for each LGA and ward. Sydney fixed effects are absorbed by LGA fixed effects. L1 refers to the first lagged value of the variable. x denotes an interaction effect.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3.11: Relationship between the number of high-rise dwellings and proximity to local government area boundaries using block level data in Sydney

	(1) Small Area Highrise
Distance to Boundary	-669.2***
	(192.7)
Total Highrise in LGA	0.003***
	(0.000)
Total Area of LGA (km^2)	-1.59e-08**
	(7.30e-09)
Observations	7102

Note: Data is based on collection districts, which are small areas that the Australia Bureau of Statistics reported data for in 2006. Distance was calculated using QGIS software and maps provided by the Australian Bureau of Statistics.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Congestion

If congestion is one of the drivers of preferences for reduced regulation we should expect that development would be clustered on the edge of local government areas, as this externalizes some of the congestion effect onto other neighborhoods. We used mapping software to calculate the distance from each block in Sydney in 2006 to the edge of the local government area, and then tested whether, conditional on the number of high-rise in an area, the blocks with more high-rise were likely to be closer to the edge of the local government area. We found that high-rise were more likely to be clustered closer to the edge for suburban local government areas (see table 3.11 and figure 3.4).

Figure 3.4: High-rise in Sydney

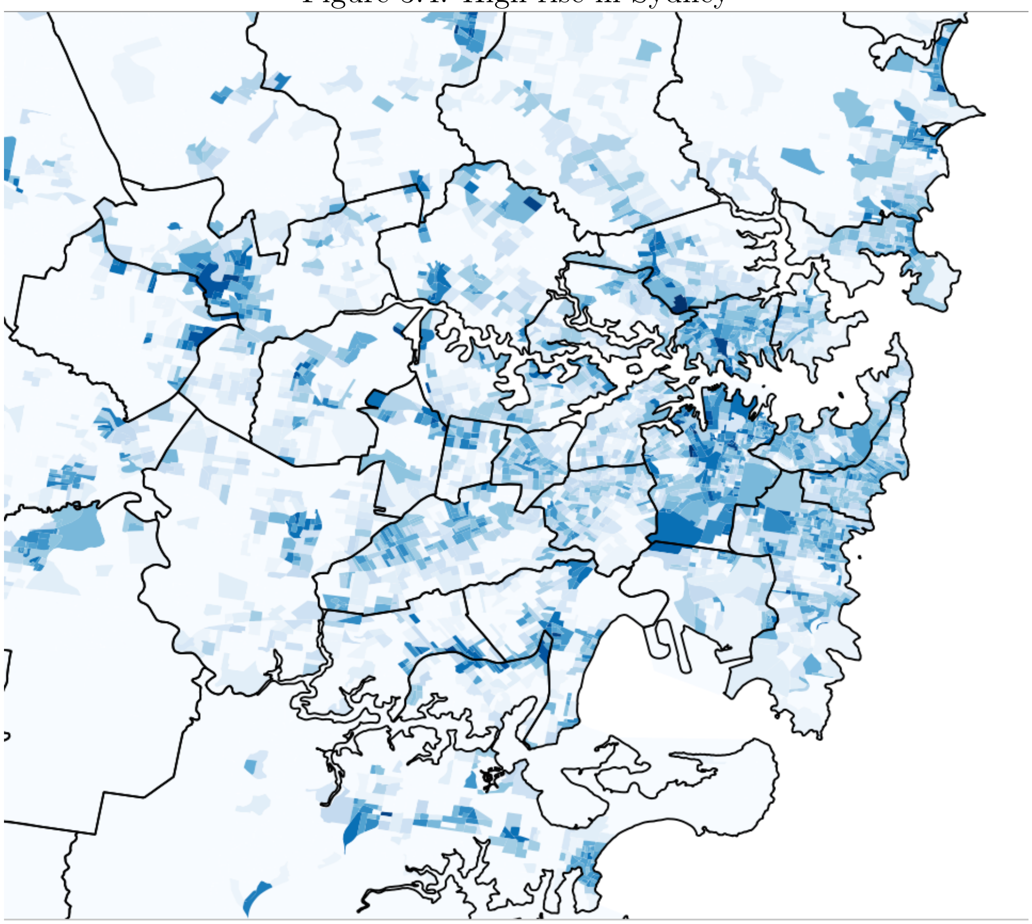

3.8 Appendix: Math

Proofs

Proposition 3.1 If we have two individuals with income m and $\hat{m} > m$, then if individual with income m wants to live in the neighborhood then individual with income $\hat{m}$ wants to live in the neighborhood also.

PROOF: Assume not. Then, since the individual with income m has utility such that:

$$U = u\left(h_m^*, m - ph_m^*\right) + c\left(n, G\right) > u\left(H, m\right)$$

and the individual with utility $\hat{m}$ has utility such that:

$$U = u\left(h_{\hat{m}}^*, \hat{m} - ph_{\hat{m}}^*\right) + c\left(n, G\right) < u\left(H, \hat{m}\right)$$

then we must have that:

$$u\left(h_{\hat{m}}^*, \hat{m} - ph_{\hat{m}}^*\right) - u\left(h_m^*, m - ph_m^*\right) < u\left(H, \hat{m}\right) - u\left(H, m\right)$$

However, we know that:

$$\left(\hat{m} - ph_m^*\right) - \left(m - ph_m^*\right) = \hat{m} - m$$

so by the fact that the utility function is concave in y:

$$u\left(h_m^*, \hat{m} - ph_m^*\right) - u\left(h_m^*, m - ph_m^*\right) > u\left(H, \hat{m}\right) - u\left(H, m\right)$$

Furthermore, by optimality, we have that:

$$u\left(h_{\hat{m}}^*, \hat{m} - ph_{\hat{m}}^*\right) > u\left(h_m^*, \hat{m} - ph_m^*\right)$$

and therefore we must have that:

$$
\begin{aligned}
&u\left(h^*_{\hat{m}}, \hat{m} - ph^*_{\hat{m}}\right) - u\left(h^*_{m}, m - ph^*_{m}\right) \\
>&u\left(h^*_{m}, \hat{m} - ph^*_{m}\right) - u\left(h^*_{m}, m - ph^*_{m}\right) > u\left(H, \hat{m}\right) - u\left(H, m\right)
\end{aligned}
$$

So we have a contradiction. ■

Proposition 3.2 There exists a pure strategy Nash equilibrium if there is only one area.

Proof: To solve for equilibrium, we can divide the problem into two stages. In the first stage, individuals choose between living in the city and the outside option. In the second stage they choose how much housing to consume. It is sufficient to show this for $\underline{u}(m)$ monotonically weakly increasing in m, since a fixed $\underline{u}$ is a special case of this.

Equilibrium in stage 2 is therefore determined as follows. Each individual maximizes their utility function.

$$\max_{y,h} u(h, y) + c(s, L) + \lambda(m - y - ph)$$

$$
\begin{aligned}
\Rightarrow \quad & u'_h - \lambda p = 0 \\
& u'_y - \lambda = 0 \\
& m = y + ph + P
\end{aligned}
$$

$$\Rightarrow \frac{u'_h}{u'_y} = p$$

Given the assumptions on concavity, we have that there exists a unique function for each individual i

$$h^*(m_i, p)$$

which is increasing continuously in m_i and declining continuously in p. Further we know that $h^*(m,0) = \infty$.

On the supply side, we have that

$$\int_i h_i^*(m_i, p) = A$$

Since each $h_i^*(m,p)$ is declining monotonously in p_i and $h^*(m,0) = \infty$, $\int_i h_i^*(m_i,p)$ must have those same properties. Consequently there is a unique equilibrium p such that $\int_i h_i^*(m_i,p) = A$.

It remains to show that a pure strategy Nash equilibrium can also be achieved once we add stage one.

From proposition 3.1, we know that if an individual with income m prefers living in the city, the individual with $\hat{m} > m$ will also prefer living in the city.

Consequently, we know that:

$$u\left(y^*\left(\tilde{m},p\right), h^*\left(\tilde{m},p\right)\right) + c\left(\int_{\tilde{m}}^{\infty} f(m), G\right)$$

is strictly increasing in $\tilde{m}$, and is continuous in $\tilde{m}$.

Therefore, by continuity of m and $\underline{u}(m)$, and since we know that $\underline{u}(0) > u(0)$ if everyone lives in the area, and $\underline{u}(max(m)) < u(max(m))$ if only the person with maximum income lives in the area, there must be at least one $\tilde{m}$ s.t.

$$u\left(y^*\left(\tilde{m},p\right), h^*\left(\tilde{m},p\right)\right) + c\left(\int_{\tilde{m}}^{\infty} f(m), G_j\right) = \underline{u}(H,\tilde{m})$$

and

$$\int_{\tilde{m}}^{\infty} h^* df(m) = A$$

Therefore there exists at least one pure strategy equilibrium where everyone above some cutoff $\tilde{m}$ lives in the city, and everyone else lives outside. ■

Proposition 3.3 Homeowners will support regulation that restricts any increase in A.

If housing is a constant factor share of income, renters will oppose it.

PROOF: If homeowners choose to stay in their homes in the next period, their utility will be:

$$u(h,m)+c(n,G)$$

in which case, any increase in n will reduce their utility. Thus in that case, they will prefer to reduce n as much as possible to maximize utility. Thus they will support any regulation that increases $\tilde{M}$.

Alternatively, if they move outside the city, they will obtain the following utility:

$$u\left(H,m+\rho(p)\,h\right)$$

where ρ is the discount rate over the infinite future.

In that case, utility is maximized by choosing the development level that maximizes ph.

We know that p is given by a combination of two equations:

$$u\left(h^*\left(\tilde{M},p\right),\tilde{M}-ph^*\left(\tilde{M},p\right)\right)+c\left(n,G\right)=u\left(H,\tilde{M}\right) \tag{3.1}$$

and

$$\int_{\widetilde{M}}^{\infty} h^*\left(m,p\right)dm=A \tag{3.2}$$

If we increase A then we know from 3.2 that either p must decrease or if not $\tilde{M}$ must decrease. However, we know from the first equation that if $\tilde{M}$ declines then by continuity we have that p declines also.

Fully differentiating 3.1 with respect to A gives us:

$$\frac{du}{dh^*}\frac{dh^*}{dp}.\frac{dp}{dA}+\frac{du}{dh^*}.\frac{dh^*}{d\tilde{M}}.\frac{d\tilde{M}}{dA}-\frac{du}{dy^*}\frac{dp}{dA}h^*+\frac{du}{dy^*}.\frac{d\tilde{M}}{dA}-\frac{du}{dy^*}.p.\frac{dh^*}{dp}.\frac{dp}{dA}-\frac{du}{dy^*}.p.\frac{dh^*}{d\tilde{M}}.\frac{d\tilde{M}}{dA}+\frac{dc}{dn}.\frac{dn}{dA}=\frac{du}{dy}.\frac{d\tilde{M}}{dA}$$

$$\frac{dh^*}{dp}.\frac{dp}{dA}.\left[\frac{du}{dh^*}-\frac{du}{dy^*}.p\right]+\frac{dh^*}{d\tilde{M}}.\frac{d\tilde{M}}{dA}\left[\frac{du}{dh^*}-\frac{du}{dy^*}.p\right]-\frac{du}{dy^*}\frac{dp}{dA}h^*+\frac{dc}{dn}.\frac{dn}{dA}-\frac{d\tilde{M}}{dA}\left(\frac{du}{dy\left(\tilde{M}\right)}-\frac{du}{dy^*\left(\tilde{M},p\right)}\right)>0$$

$$\begin{aligned}
\Rightarrow \frac{dh^*}{dp}.\frac{dp}{dA}.\left[\frac{du}{dh^*}-\frac{du}{dy^*}.p\right]+\frac{dh^*}{d\tilde{M}}.\frac{d\tilde{M}}{dA}\left[\frac{du}{dh^*}-\frac{du}{dy^*}.p\right]-\frac{du}{dy^*}\frac{dp}{dA}h^* &> -\frac{dc}{dn}.\frac{dn}{dA}\\
\Rightarrow \frac{dh^*}{dp}.\frac{dp}{dA}.\left[\frac{du}{dh^*}-\frac{du}{dy^*}.p\right]+\frac{dh^*}{d\tilde{M}}.\frac{d\tilde{M}}{dA}\left[\frac{du}{dh^*}-\frac{du}{dy^*}.p\right]-\frac{du}{dy^*}\frac{dp}{dA}h^* &> -\frac{dc}{dn}.\frac{dn}{dA}\\
\Rightarrow \frac{du}{dy^*\left(\tilde{M},p\right)}\left(\frac{dp}{dA}h^*\left(\tilde{M},p\right)\right) &< \frac{dc}{dn}.\frac{dn}{dA}<0
\end{aligned}$$

Alternatively, $\tilde{M}$ cannot decline if p stays constant. So we have that both must decrease.

Thus owners will block anticipated land releases.

Renters, unlike owners, will have utility in the next period of:

$$u\left(h^*\left(p,m\right),m-ph^*\right)+c(n,G)$$

if they stay in the area, as they will need to pay for the housing. Alternatively, they can have the utility:

$$u\left(H,m\right)$$

Since the outside option isn't changing and renters are already in the area, they clearly prefer to remain in the area rather than exit. If the total population is growing, then in the next period some existing renters will not make the 'cutoff' if A does not increase. They will prefer the outcome that maximizes their internal utility.

We also know that utility is decreasing in p but also decreasing in n, both of which increase as we increase A. So it remains to see which effect dominates. Maximizing renter utility with respect to A gives us:

$$\begin{aligned}
\frac{dU^*}{dA} &= \frac{du}{dh^*}\frac{dh^*}{dp}.\frac{dp}{dA} - \frac{du}{dy^*}\frac{dp}{dA}h^* - \frac{du}{dy^*}.p.\frac{dh^*}{dp}.\frac{dp}{dA} + \frac{dc}{dn}.\frac{dn}{dA} \\
&= \frac{dh^*}{dp}.\frac{dp}{dA}\left(\frac{du}{dh^*} - \frac{du}{dy^*}.p\right) - \frac{du}{dy^*}\frac{dp}{dA}h^* + \frac{dc}{dn}.\frac{dn}{dA} \\
&= -\frac{du}{dy^*}\frac{dp}{dA}h^*(m,p) + \frac{dc}{dn}.\frac{dn}{dA} \\
&> \frac{du}{dy^*\left(\tilde{M},p\right)}\frac{dp}{dA}h^*\left(\tilde{M},p\right) - \frac{du}{dy^*\left(m,p\right)}\frac{dp}{dA}h^*(m,p) \\
&> \frac{dp}{dA}\left[\frac{du}{dy^*\left(\tilde{M},p\right)}h^*\left(\tilde{M},p\right) - \frac{du}{dy^*\left(m,p\right)}h^*(m,p)\right]
\end{aligned}$$

Therefore renters will oppose regulation that limits area growth so long as $\frac{d\frac{du^*}{dy^*}h^*(m,p)}{dm} > 0$.

$$\begin{aligned}
\frac{d\frac{du^*}{dy^*}h^*(m,p)}{dm} &= \frac{du^*}{dm}.\frac{dh^*}{dm} + \frac{du^{2*}}{dm^2}h^* \\
&= \frac{d^2u}{dy^2}\left(1 - p\frac{dh}{dm}\right)h + \frac{du^*}{dy}.\frac{dh}{dm}
\end{aligned}$$

which is definitely positive if

$$\frac{1}{p} = \frac{dh}{dm}$$

Or essentially if housing is a constant factor share of income. Note that housing being less than a constant factor share of income is not sufficient to render this untrue, but the outcome would be dependent on the utility function. ■

Proposition 3.4 Homeowners will support regulation that restricts any increase in n so long as there is a third party to purchase any excess land. If housing is a constant factor share of income, renters will oppose it.

PROOF: If homeowners choose to stay in their homes in the next period, their utility will be:

$$u(h,m)+c(n,G)$$

in which case, any increase in n will reduce their utility. Thus in that case, they will prefer to reduce n as much as possible to maximize utility.

Alternatively, if they move outside the city, they will obtain the following utility:

$$u\left(H, m+\rho(p)h\right)$$

In that case, utility is maximized by choosing the development level that maximizes p.

We know that p is given by a combination of two equations:

$$u\left(h^*\left(\tilde{M},p\right), \tilde{M}-ph^*\left(\tilde{M},p\right)\right)+c\left(n,G\right)=u\left(H,\tilde{M}\right) \tag{3.3}$$

and

$$\int_{\widetilde{M}}^{\infty} h^*\left(m,p\right)dm = A \tag{3.4}$$

Completely differentiating 3.3 gives us:

$$\frac{du}{dh^*}\left(\frac{dh^*}{d\tilde{M}}\frac{d\tilde{M}}{dn}+\frac{dh^*}{dp}.\frac{dp}{dn}\right)+\frac{du}{dy^*}\frac{d\tilde{M}}{dn}-\frac{du}{dy^*}.p\frac{dh^*}{d\tilde{M}}\frac{d\tilde{M}}{dn}-\frac{du}{dy^*}h^*\frac{dp}{dn}-\frac{du}{dy^*}p\frac{dh^*}{dp}\frac{dp}{dn}+\frac{dc}{dn}=\frac{du}{dy}.\frac{d\tilde{M}}{dn}$$

$$\frac{du}{dh^*}\left(\frac{dh^*}{d\tilde{M}}\frac{d\tilde{M}}{dn}+\frac{dh^*}{dp}\frac{dp}{dn}\right)-\frac{du}{dy^*}p\frac{dh^*}{d\tilde{M}}\frac{d\tilde{M}}{dn}-\frac{du}{dy^*}h^*\frac{dp}{dn}-\frac{du}{dy^*}p\frac{dh^*}{dp}\frac{dp}{dn}+\frac{dc}{dn}-\frac{d\tilde{M}}{dn}\left(\frac{du}{dy\left(\tilde{M}\right)}-\frac{du}{dy^*\left(\tilde{M},p\right)}\right)>0$$

$$\begin{aligned}
\Rightarrow \frac{dh^*}{d\tilde{M}}\frac{d\tilde{M}}{dn}\left[\frac{du}{dh^*} - \frac{du}{dy^*}.p\right] + \frac{dh^*}{dp}\frac{dp}{dn}\left[\frac{du}{dh^*} - \frac{du}{dy^*}.p\right] - \frac{du}{dy^*}h^*\frac{dp}{dn} &> -\frac{dc}{dn} \\
\Rightarrow \frac{dp}{dn} &< \frac{dc}{dn}\left[\frac{du}{dy^*}h^*\right]^{-1} \\
&< 0
\end{aligned}$$

which implies that as we limit n prices must rise.

Note, however, that $\tilde{M}$ must also increase (since that is how we limit n), so the second equation must fail to bind. Therefore, some homeowners may be unwilling to sell their houses since no-one will buy them. To keep this constraint in place, therefore, there must be an external party who is forced to purchase the slack.

Renters, unlike owners, will have utility in the next period of:

$$u\left(h^*\left(p, m\right), m - ph^*\right) + c(n, G)$$

if they stay in the area, as they will need to pay for the housing. Alternatively, they can have the utility:

$$u\left(H, m\right)$$

Since the outside option isn't changing and renters are already in the area, they clearly prefer to remain in the area rather than exit. If the total population is growing, then in the next period some existing renters will not make the income 'cutoff' if n does not increase. They will prefer will prefer the outcome that maximizes their internal utility.

Differentiating renter utility with respect to n gives us:

$$
\begin{aligned}
\frac{dU}{dn} &= \frac{du}{dh^*}\frac{dh^*}{dp}\frac{dp}{dn} - \frac{du}{dy^*}\frac{dh^*}{dp}\frac{dp}{dn}p - \frac{du}{dy^*}\frac{dp}{dn}h^* + \frac{dc}{dn} \\
&= \frac{dh^*}{dp}\frac{dp}{dn}\left[\frac{du}{dh^*} - \frac{du}{dy^*}p\right] - \frac{du}{dy^*}\frac{dp}{dn}h^* + \frac{dc}{dn} \\
&= -\frac{du}{dy^*}\frac{dp}{dn}h^* + \frac{dc}{dn} \\
&> \frac{dc}{dn}\left(\frac{du}{dy^*\left(\tilde{M},p\right)}h^*\left(\tilde{M},p\right)\right)^{-1}\frac{du}{dy^*\left(m,p\right)}h^*\left(m,p\right) - \frac{dc}{dn} \\
&> \frac{dc}{dn}\left[\frac{du}{dy^*\left(m,p\right)}h^*\left(m,p\right)\left(\frac{du}{dy^*\left(\tilde{M},p\right)}h^*\left(\tilde{M},p\right)\right)^{-1} - 1\right] \\
&> 0 \qquad \blacksquare
\end{aligned}
$$

where the last inequality follows from the same argument presented in proposition 3.3, since $\frac{d\frac{du^*}{dy^*}h^*(m,p)}{dm} > 0$ so $\frac{du}{dy^*(m,p)}h^*\left(m,p\right)\left(\frac{du}{dy^*\left(\tilde{M},p\right)}h^*\left(\tilde{M},p\right)\right)^{-1} > 1$

Therefore, renters will oppose regulations that place restrictions on n.

Proposition 3.5 Independent of the effects on price, owners are more likely than renters to support minimum size regulations. The effect on price depends on the nature of the distribution, the relative concavity of the utility function, the value of the outside option, and the degree of impact of congestion. Higher prices in response to a minimum size regulation will result in owners supporting the policy, lower prices will result in renters supporting the policy.

Proof: Since no-one is increasing $\underline{h}$ over the existing stock, we know that owners cannot lose what they have. Once again they can only choose whether to continue to live in the area or sell and move out. Consequently, if they choose to stay then their utility is:

$$u\left(h, y\right) + c\left(n, G\right)$$

where everything except n is fixed. Therefore, they will support any regulation that increases n. On the other hand, if they choose to leave, they will only support regulation that increases p.

Renters on the other hand will have utility in the next period of:

$$u\left(h^*\left(p,m\right),m-ph^*\right)+c(n,G)$$

if they stay in the area, as they will need to pay for the housing. Since they have no staying power, however, their current consumption is largely irrelevant. If, after the expected level of growth, current period renters no longer wish to live in the area we will ignore them, as they have no clearly defined preferences either way.

If they do choose to leave next period, renters will have the utility:

$$u\left(H,m\right)$$

Since the outside option isn't changing with growth, any renters who are already living in the area clearly prefer to remain in the area rather than exit. They will therefore prefer the degree of regulation that maximizes their internal utility.

Anyone who wants to consume $h^* < \underline{h}$ would prefer not to increase $\underline{h}$ unless they are above the new minimum sizing cutoff.

If we place a lower bound on $\underline{h}$ then in order for the total area constraint to bind we must have that:

$$\int_{\tilde{M}}^{\underline{M}} \underline{h}dm + \int_{\underline{M}}^{\infty} h^*\left(m,p\right)dm = A \tag{3.5}$$

and for the lowest-income individual to live in the area we must have that

$$u\left(\underline{h},\tilde{M}-p\underline{h}\right)+c\left(n,G\right)=u\left(H,\tilde{M}\right) \tag{3.6}$$

We also define $\underline{M}$ such that

$$\begin{aligned} u\left(h^*\left(\underline{M},p\right),\underline{M}-ph^*\left(\underline{M},p\right)\right) &= u\left(\underline{h},\underline{M}-p\underline{h}\right) \\ \Rightarrow h^*\left(\underline{M},p\right) &= \underline{h} \end{aligned}$$

From 3.5 we know that ceteris paribus increasing $\underline{h}$ increases total demand, and that therefore either $\tilde{M}$ must increase or p must increase.

Differentiating 3.6 with respect to $\underline{h}$ gives us:

$$\begin{aligned} \frac{du}{dy\left(\tilde{M},H\right)}\frac{d\tilde{M}}{d\underline{h}} &= \frac{du}{d\underline{h}}+\frac{du}{dy\left(\tilde{M},\underline{h},p\right)}\frac{d\tilde{M}}{d\underline{h}}-\frac{du}{dy\left(\tilde{M},\underline{h},p\right)}p-\frac{du}{dy\left(\tilde{M},\underline{h},p\right)}\frac{dp}{d\underline{h}}\underline{h}+\frac{dc}{dn}\frac{dn}{d\tilde{M}}\frac{d\tilde{M}}{d\underline{h}} \\ \Rightarrow \frac{du}{dy\left(\tilde{M},\underline{h},p\right)}\frac{dp}{d\underline{h}}\underline{h} &= \frac{d\tilde{M}}{d\underline{h}}\left[\frac{du}{dy\left(\tilde{M},\underline{h},p\right)}-\frac{du}{dy\left(\tilde{M},H\right)}\right]+\frac{dc}{dn}\frac{dn}{d\tilde{M}}\frac{d\tilde{M}}{d\underline{h}}-\left[\frac{du}{dy\left(\tilde{M},\underline{h},p\right)}p-\frac{du}{d\underline{h}}\right] \end{aligned}$$

$$\begin{aligned} \Rightarrow \frac{dp}{d\underline{h}} &= \left(\frac{du}{dy\left(\tilde{M},\underline{h},p\right)}\underline{h}\right)^{-1}\left[\frac{d\tilde{M}}{d\underline{h}}\left[\frac{du}{dy\left(\tilde{M},\underline{h},p\right)}-\frac{du}{dy\left(\tilde{M},H\right)}+\frac{dc}{dn}\frac{dn}{d\tilde{M}}\right]-\left[\frac{du}{dy\left(\tilde{M},\underline{h},p\right)}p-\frac{du}{d\underline{h}}\right]\right] \\ \Rightarrow \frac{d\tilde{M}}{d\underline{h}} &= \left[\frac{du}{dy\left(\tilde{M},\underline{h},p\right)}\frac{dp}{d\underline{h}}\underline{h}+\frac{du}{dy\left(\tilde{M},\underline{h},p\right)}p-\frac{du}{d\underline{h}}\right]\left[\frac{du}{dy\left(\tilde{M},\underline{h},p\right)}-\frac{du}{dy\left(\tilde{M},H\right)}+\frac{dc}{dn}\frac{dn}{d\tilde{M}}\right]^{-1} \end{aligned}$$

If $\frac{d\tilde{M}}{d\underline{h}}\leq 0$ then we have that $\frac{dp}{d\underline{h}}<0$, which violates what we have already found from 3.5. Therefore $\frac{d\tilde{M}}{d\underline{h}}>0$.

Therefore, if renters are evenly distributed through the population, some of them will not make the cutoff for the higher minimum.

For the remaining owners and renters, their housing decision is dependent on the change in price. Owners who are leaving prefer a higher price, as this is the only impact that the policy will have on their utility. Differentiating renter optimal utility with respect to n gives us:

$$
\begin{aligned}
\frac{dU}{d\underline{h}} &= \frac{du}{dh^*}\frac{dh^*}{dp}\frac{dp}{d\underline{h}} - \frac{du}{dy^*}\frac{dh^*}{dp}\frac{dp}{d\underline{h}}p - \frac{du}{dy^*}\frac{dp}{d\underline{h}}h^* + \frac{dc}{dn}.\frac{dn}{d\tilde{M}}.\frac{d\tilde{M}}{d\underline{h}} \\
&= \frac{dp}{d\underline{h}}\frac{dh^*}{dp}\left[\frac{du}{dh^*} - \frac{du}{dy^*}p\right] - \frac{du}{dy^*}\frac{dp}{d\underline{h}}h^* + \frac{dc}{dn}.\frac{dn}{d\tilde{M}}.\frac{d\tilde{M}}{d\underline{h}} \\
&= -\frac{du}{dy^*(m,p)}\frac{dp}{d\underline{h}}h^*(m,p) + \frac{dc}{dn}.\frac{dn}{d\tilde{M}}.\frac{d\tilde{M}}{d\underline{h}}
\end{aligned}
$$

So renters would clearly prefer less development, unless $\frac{dp}{d\underline{h}} < 0$.

It remains to consider the likelihood that $\frac{dp}{d\underline{h}} > 0$.

Since staying in the area is exogenous, we can simply deduct the total consumption of those that stay from the area, giving us an adjusted version of 3.5:

$$
\int_{\tilde{M}}^{\underline{M}} \underline{h}dm + \int_{\underline{M}}^{\infty} h^*(m,p)\,dm = A \tag{3.7}
$$

where A is the remaining area after owners who are staying have been removed, and m is the distribution with those people removed. Differentiating 3.7 with respect to $\underline{h}$ (using Leibniz rule) gives us:

$$
\begin{aligned}
\frac{d}{d\underline{h}}\left(\underline{h}\left[\int_{\tilde{M}}^{\infty} dm - \int_{\underline{M}}^{\infty} dm\right] + \int_{\underline{M}}^{\infty} h^*(m,p)\,dm\right) &= \left[\int_{\tilde{M}}^{\infty} dm - \int_{\underline{M}}^{\infty} dm\right] - \underline{h}.\frac{d\tilde{M}}{d\underline{h}} + \frac{d\underline{M}}{d\underline{h}}\underline{h} + \int_{\underline{M}}^{\infty} \frac{dh^*}{dp}\frac{dp}{d\underline{h}}dm - \frac{d\underline{M}}{d\underline{h}}h^*(\underline{M},p) \\
&= 0 \\
\Rightarrow -\int_{\underline{M}}^{\infty} \frac{dh^*}{dp}\frac{dp}{d\underline{h}}dm &= \int_{\tilde{M}}^{\infty} dm - \int_{\underline{M}}^{\infty} dm - \underline{h}.\frac{d\tilde{M}}{d\underline{h}} \\
&= \int_{\tilde{M}}^{\underline{M}} dm - \underline{h}.\frac{d\tilde{M}}{d\underline{h}}
\end{aligned}
$$

$$
\Rightarrow \frac{dp}{d\underline{h}} = -\left[\int_{\tilde{M}}^{\underline{M}} dm - \underline{h}.\frac{d\tilde{M}}{d\underline{h}}\right]\left(\int_{\underline{M}}^{\infty} \frac{dh^*}{dp}dm\right)^{-1}
$$

Essentially this means that the decrease (increase) in consumption of the people who are consuming their optimal housing (more than $\underline{h}$) must be equal to the increase (decrease) in consumption of those who would like to consume less than $\underline{h}$ but must instead consume $\underline{h}$ less the decrease (increase) in consumption from those at the bottom leaving the area.

Plugging in the value for $\frac{d\tilde{M}}{d\underline{h}}$ we obtain the following:

$$-\frac{dp}{d\underline{h}}\int_{\underline{M}}^{\infty}\frac{dh^*}{dp}dm$$

$$=\int_{\tilde{M}}^{M}dm-\underline{h}\left[\frac{du}{dy\left(\tilde{M},\underline{h},p\right)}\frac{dp}{d\underline{h}}\underline{h}+\frac{du}{dy\left(\tilde{M},\underline{h},p\right)}p-\frac{du}{d\underline{h}}\right]\left[\frac{du}{dy\left(\tilde{M},\underline{h},p\right)}-\frac{du}{dy\left(\tilde{M},H\right)}+\frac{dc}{dn}\frac{dn}{d\tilde{M}}\right]^{-1}$$

$$\Rightarrow\frac{dp}{d\underline{h}} = \left[\int_{\tilde{M}}^{M}dm-\underline{h}\left[\frac{du}{dy\left(\tilde{M},\underline{h},p\right)}p-\frac{du}{d\underline{h}}\right]\left[\frac{du}{dy\left(\tilde{M},\underline{h},p\right)}-\frac{du}{dy\left(\tilde{M},H\right)}+\frac{dc}{dn}\frac{dn}{d\tilde{M}}\right]^{-1}\right]*$$
$$\left[\underline{h}^2\frac{du}{dy\left(\tilde{M},\underline{h},p\right)}\left[\frac{du}{dy\left(\tilde{M},\underline{h},p\right)}-\frac{du}{dy\left(\tilde{M},H\right)}+\frac{dc}{dn}\frac{dn}{d\tilde{M}}\right]^{-1}-\int_{\underline{M}}^{\infty}\frac{dh^*}{dp}dm\right]^{-1}$$

The sign on this is determined by

$$\int_{\tilde{M}}^{M}dm-\underline{h}\left[\frac{du}{dy\left(\tilde{M},\underline{h},p\right)}p-\frac{du}{d\underline{h}}\right]\left[\frac{du}{dy\left(\tilde{M},\underline{h},p\right)}-\frac{du}{dy\left(\tilde{M},H\right)}+\frac{dc}{dn}\frac{dn}{d\tilde{M}}\right]^{-1}$$

where the first term is positive and the second term is negative.

Essentially, the change in price is balanced by the nature of the distribution, the relative concavity of the utility function, the value of the outside option and the degree of impact of congestion. ∎

Proposition 3.6 If we have two individuals with incomes m and $\hat{m} > m$, and individual m strictly (weakly) prefers neighborhood B to neighborhood A while individual with income $\hat{m}$ strictly (weakly) prefers the reverse, then we must have that $p_A > p_B$ $(p_A = p_B)$.

PROOF: We know that individual with income m has

$$u\left(h^*\left(m,p_A\right),m-p_Ah^*\left(m,p_A\right)\right)-c\left(n_A,G_A\right)<u\left(h^*\left(m,p_B\right),m-p_Bh^*\left(m,p_B\right)\right)-c\left(n_B,G_B\right)$$

whereas individual with income $\tilde{m}$ prefers:

$$u\left(h^*\left(\hat{m},p_A\right),\hat{m}-p_Ah^*\left(\hat{m},p_A\right)\right)-c\left(n_A,G_A\right)>u\left(h^*\left(\hat{m},p_B\right),\hat{m}-p_Bh^*\left(\hat{m},p_B\right)\right)-c\left(n_B,G_B\right)$$

Then assume that $p_A < P_B$. Then we have that:

$$\begin{aligned} u\left(h^*\left(m,p_B\right),m-p_Bh^*\left(m,p_B\right)\right) &< u\left(h^*\left(m,p_B\right),m-p_Ah^*\left(m,p_B\right)\right) \\ &< u\left(h^*\left(m,p_A\right),m-p_Ah^*\left(m,p_A\right)\right) \end{aligned}$$

So we must have, for m to prefer B to A, that:

$$c\left(n_B,G_B\right)<c\left(n_A,G_A\right)$$

But for $\hat{m}$ to prefer A to B while m prefers B to A, in that case:

$$\begin{aligned} &u\left(h^*\left(\hat{m},p_A\right),\hat{m}-p_Ah^*\left(\hat{m},p_A\right)\right)-u\left(h^*\left(\hat{m},p_B\right),\hat{m}-p_Bh^*\left(\hat{m},p_B\right)\right) \\ >&u\left(h^*\left(m,p_A\right),m-p_Ah^*\left(m,p_A\right)\right)-u\left(h^*\left(m,p_B\right),m-p_Bh^*\left(m,p_B\right)\right) \\ >&0 \end{aligned}$$

However, based on the utility function we also know that, given a fixed area choice:

$$\frac{u'_h}{u'_y}=p$$

And furthermore, taking the second derivative of utility with respect to income and price:

$$\begin{aligned}
\frac{du}{dm} &= u_h' \frac{dh^*}{dm} + u_y' \left(1 - p\frac{dh^*}{dm}\right) \\
\Rightarrow \frac{d^2u}{dmdp} &= u_h' \frac{d^2h^*}{dmdp} + u_y'' \left(1 - p\frac{dh^*}{dm}\right)^2 - u_y' \left(p\frac{d^2h^*}{dmdp}\right) - u_y' \frac{dh^*}{dm} \\
&= \frac{d^2h^*}{dmdp}\left[u_h' - u_y' p\right] + u_y'' \left(1 - p\frac{dh^*}{dm}\right)^2 - u_y' \frac{dh^*}{dm} \\
&= u_y'' \left(1 - p\frac{dh^*}{dm}\right)^2 - u_y' \frac{dh^*}{dm} \\
&< 0
\end{aligned}$$

Therefore, we know that:

$$\begin{aligned}
& u\left(h^*\left(\hat{m}, p_A\right), \hat{m} - p_A h^*\left(\hat{m}, p_A\right)\right) - u\left(h^*\left(\hat{m}, p_B\right), \hat{m} - p_B h^*\left(\hat{m}, p_B\right)\right) \\
< \; & u\left(h^*\left(m, p_A\right), m - p_A h^*\left(m, p_A\right)\right) - u\left(h^*\left(m, p_B\right), m - p_B h^*\left(m, p_B\right)\right)
\end{aligned}$$

And so we have a contradiction. ∎

Proposition 3.8 There is an equilibrium for multiple areas.

PROOF: By propositions 3.6 and 3.7, all equilibria are weakly separating. Furthermore, since all areas have different levels of public goods G, there must be only one equilibrium which equates all values of $c(n, G)$. If $c(n, G)$ is higher or lower for an area, the price will change. Therefore, we must have that there is at most one pooling equilibrium.

It remains to show that there is at least one separating equilibrium. This we will do by construction.

Let $h_j^*\left(m, p_j\right)$ be the optimal housing for any given individual conditional on their having chosen area j. By the maximum theorem $h_j^*\left(m, p_j\right)$ is continuous in m and p_j. Let $f\left(m\right)$ be the distribution of m.

Let there be k locations. Number these so that $G_1 > G_2 > ... > G_k$. Let $\tilde{M}_k$ be the lowest income level found in group k. Then we have $2k$ unknowns, since $\tilde{M}_j$ is unknown for all j and p_j is unknown for all j.

Then, choose an $\underline{M} < \tilde{M}_1 < \bar{M}$ at random. Then, we know that:

$$\int_{\tilde{M}_1}^{\infty} h^*\left(m, p_1\right) f\left(m\right) dm = A_1$$

Since $h^*(m, p_1)$ is continuous in p_1, and the distribution is continuous, then the above is continuous and strictly decreasing in p_1, we know that it must have a unique solution for p_1.

By the same argument for each subsequent area, we have that:

$$\int_{\tilde{M}_j}^{\tilde{M}_{j-1}} h^*\left(m, p_j\right) f\left(m\right) dm = A_j \tag{3.8}$$

which gives us a continuous function from $p_j \rightarrow \tilde{M}_j$ given a fixed $\tilde{M}_{j-1}$. Furthermore, since we have that $h^*\left(m, p_j\right)$ is strictly decreasing in p_j and $\tilde{M}_j$ is strictly decreasing in h^*, we know that $\tilde{M}_j$ is strictly decreasing in p_j and so we have a 1:1 function. Furthermore, as $p_j \rightarrow 0, \tilde{M}_j \rightarrow \tilde{M}_{j-1}$ and as $p_j \rightarrow \infty, \tilde{M}_j \rightarrow \underline{M}$.

In addition, we know that:

$$u\left(y^*\left(\tilde{M}_1, p_1\right), h^*\left(\tilde{M}_1, p_1\right)\right) - c\left(\int_{\tilde{M}_1}^{\infty} f(m)dm, G_1\right) = u\left(y^*\left(\tilde{M}_1, p_2\right), h^*\left(\tilde{M}_1, p_2\right)\right) - c\left(\int_{\tilde{M}_2}^{\tilde{M}_1} f(m)dm, G_2\right) \tag{3.9}$$

which has a fixed left hand side. This gives us a continuous function from $p_2 \rightarrow \tilde{M}_2$ that is continuous and strictly increasing. In this case, as $p_2 \rightarrow 0$ $\tilde{M}_2 \rightarrow \underline{M}$ and as $p_2 \rightarrow \infty, \tilde{M}_j \rightarrow \tilde{M}_1$.

Therefore 3.8 and 3.9 must have a unique crossing point, and must have a unique crossing point so long as $\tilde{M}_1 > \underline{M}$. Therefore given a fixed $\tilde{M}_1$ we have a unique pair p_2 and $\tilde{M}_2$.

Furthermore, by 3.8 we know that if $\tilde{M}_1$ increases, then either $\tilde{M}_2$ increases or p_2 increases (or both). However, by 3.9 and since $p_1 > p_2$ from proposition 3.7, we know that:

$$\Rightarrow \frac{d\tilde{M}_2}{d\tilde{M}_1} > 0$$

Therefore even if p_2 is increasing $\tilde{M}_2$ is increasing also.

By an iterated version of the same argument, we can continue to allocate individuals to areas until either (a) $\tilde{M}_j = \underline{M}$ or (b) we obtain a p_k and $\tilde{M}_k$.

In the case of (a), raise $\tilde{M}_1$ by $\epsilon < \bar{M} - \tilde{M}_1$ and repeat.

We know that $\tilde{M}_j$ is increasing in $\tilde{M}_{j-1}$ and therefore for all $\tilde{M}_j$, and therefore $\tilde{M}_j$ is increasing in $\tilde{M}_1$ for all j. Therefore, by iteratively increasing $\tilde{M}_1$, we will eventually reach case (b).

In the case of (b), we then have the constraint that

$$u\left(y^*\left(\tilde{M}_k, p_k\right), h^*\left(\tilde{M}, p_k\right)\right) - c\left(\int_{\tilde{M}_k}^{\tilde{M}_{k-1}} f(m), G_j\right) = \underline{u}\left(H, \tilde{M}_k\right)$$

which may or may not hold. If so then we are done. Otherwise, if

$$u\left(y^*\left(\tilde{M}_k, p_k\right), h^*\left(\tilde{M}, p_k\right)\right) - c\left(\int_{\tilde{M}_k}^{\tilde{M}_{k-1}} f(m), G_j\right) > \underline{u}\left(H, \tilde{M}_k\right)$$

go back and decrease $\tilde{M}_1$ by $\epsilon < \bar{M} - \tilde{M}_1$. If

$$u\left(y^*\left(\tilde{M}_k, p_k\right), h^*\left(\tilde{M}_k, p_k\right)\right) - c\left(\int_{\tilde{M}_k}^{\tilde{M}_{k-1}} f(m), G_j\right) < \underline{u}\left(H, \tilde{M}_k\right)$$

go back in increase $\tilde{M}_1$ by $\epsilon < \bar{M} - \tilde{M}_1$.

Since we know that increasing (decreasing) $\tilde{M}_1$ increases (decreases) $\tilde{M}_k$, it remains to show that

$$\frac{du\left(y^*\left(\tilde{M}_k, p_k\right), h^*\left(\tilde{M}_k, p_k\right)\right) - c\left(\int_{\tilde{M}_k}^{\tilde{M}_{k-1}} f(m), G_j\right)}{d\tilde{M}_1} > \frac{d\underline{u}\left(H, \tilde{M}_k\right)}{d\tilde{M}_1} \text{ for all } \tilde{M}_1$$

In which case we must have convergence to an equilibrium, since the inside utility from changing $\tilde{M}_1$ is changing faster than the outside utility. Furthermore, since $u(0) < \underline{u}(0)$, we know that we won't always have everyone living outside, and since $u\left(\tilde{M}\right) > \underline{u}\left(\tilde{M}\right)$ we also won't have everyone inside.

Solving for the right hand side gives us:

$$\frac{d\underline{u}\left(H,\tilde{M}_k\right)}{d\tilde{M}_1} = \frac{du}{dy}\frac{d\tilde{M}_k}{d\tilde{M}_1}$$

Solving for the left hand side gives us:

$$\begin{aligned}
&\frac{du\left(y^*\left(\tilde{M}_k,p_k\right),h^*\left(\tilde{M}_k,p_k\right)\right)-c\left(\int_{\tilde{M}_k}^{\tilde{M}_{k-1}}f(m),G_j\right)}{d\tilde{M}_1}\\
= \;& \frac{du}{dy^*}\left(\frac{d\tilde{M}_k}{d\tilde{M}_1}-p\frac{dh^*}{d\tilde{M}_k}\frac{d\tilde{M}_k}{d\tilde{M}_1}-\frac{dp}{d\tilde{M}_k}\frac{d\tilde{M}_k}{d\tilde{M}_1}h^*-p\frac{dh^*}{dp}\frac{dp}{d\tilde{M}_k}\frac{d\tilde{M}_k}{d\tilde{M}_1}\right)+\frac{du}{dh}\left(\frac{dh^*}{d\tilde{M}_k}\frac{d\tilde{M}_k}{d\tilde{M}_1}+\frac{dp}{d\tilde{M}_k}\frac{d\tilde{M}_k}{d\tilde{M}_1}\right)-\frac{dc}{dn}\frac{dn}{d\tilde{M}_k}\frac{d\tilde{M}_k}{d\tilde{M}_1}\\
= \;& \frac{du}{dy^*}\left(\frac{d\tilde{M}_k}{d\tilde{M}_1}-\frac{dp}{d\tilde{M}_k}\frac{d\tilde{M}_k}{d\tilde{M}_1}h^*\right)-\frac{dc}{dn}\frac{dn}{d\tilde{M}_k}\frac{d\tilde{M}_k}{d\tilde{M}_1}\\
= \;& \left(\frac{d\tilde{M}_k}{d\tilde{M}_1}\right)\left[\frac{du}{dy^*}\left(1-\frac{dp_k}{d\tilde{M}_k}h^*\right)-\frac{dc}{dn}\frac{dn}{d\tilde{M}_k}\right]
\end{aligned}$$

And taking the difference gives us:

$$\frac{d\tilde{M}_k}{d\tilde{M}_1}\left[\frac{du}{dy^*}\left(1-\frac{dp_k}{d\tilde{M}_k}h^*\right)-\frac{dc}{dn}\frac{dn}{d\tilde{M}_k}-\frac{du}{dy}\right] = \frac{d\tilde{M}_k}{d\tilde{M}_1}\left[\left(\frac{du}{dy^*}-\frac{du}{dy}\right)-\frac{du}{dy^*}\frac{dp_k}{d\tilde{M}_k}h^*-\frac{dc}{dn}\frac{dn}{d\tilde{M}_k}\right] > 0$$

Therefore we must have that there is a sorting equilibrium for all areas. ∎

Housing growth vs density

Change in density is mathematically equivalent to percentage growth in housing if areas are not changing. Since percentage growth in housing is given by:

$$growth = \frac{housing - l.housing}{l.housing}$$

$$\begin{aligned}
density &= \frac{housing}{area} \\
\Rightarrow density - l.density &= \frac{housing}{area} - \frac{l.housing}{l.area} \\
&= \frac{housing - l.housing}{area} + \frac{l.area * l.housing - area * l.housing}{l.area * area} \\
\Rightarrow (density - l.density) * \frac{area}{l.housing} &= \frac{housing - l.housing}{l.housing} + \frac{l.area * l.housing - area * l.housing}{l.area * l.housing} \\
&\rightarrow \frac{housing - l.housing}{l.housing} \text{ as l.area-area goes to 0}
\end{aligned}$$

$$\Rightarrow \frac{(density - l.density)}{l.density} \approx growth$$

Therefore, change in density is equivalent to percentage growth in housing so long as areas are sufficiently stable. This means that running the same regression on housing growth should give the same sign of coefficients, and almost the same set of coefficients if we multiply everything except $l.density$ on the right hand side by $l.density$.

3.9 Appendix: Geographical Boundaries

The ABS has freely available paper maps of statistical areas for most census years, and electronic maps for the years 1981-2011. This allowed us to check the consistency of boundaries over the period of study.

Sydney incurred very few boundary changes. The largest boundary changes affected the Sydney, South Sydney, and Leichhardt areas, but these did not affect the results since all 3 were classified as 'central'. There were only three other relevant geographic changes – the division of Warringah into Warringah and Pittwater, the merger of Colo and Windsor to become Hawkesbury, and the merger of Concord and Drummoyne to create Canada Bay. Drummoyne and Concord merged in December 2000, so the 2001 data was included twice — both for Canada Bay as a whole (in order to establish the independent lagged variables for 2006) and separately (in order to establish the dependent variable for 2001). The former data was calculated by simply adding the results (which were reported under the old LGA structure) for the Concord and Drummoyne together. Colo and Windsor merged in 1981, but since Hawkesbury was a fringe suburb these were treated as a single area for the 1976-2011 period. Pittwater was more difficult, since it separated in 1992, so there is a much larger gap between 1992 and 1996 in which housing changes may occur. Ideally, the local government areas would have been reported separately (at least as separate statistical local areas) in 1991, but sadly this was not the case. However, since both were classified as fringe suburbs, the issue became irrelevant.

It was determined that small boundary changes should not lead to two or more areas being grouped together as a single area in the data, since doing so would obscure the underlying political dynamics. Furthermore, since voters in a given period have an impact on the zoning of the area in the next period (regardless of their presence or absence in that next period), it was determined that it was better to simply keep the boundaries as they were in each period.

Brisbane, on the other hand, had a number of large and relevant boundary changes. Furthermore, it was considered that even small boundary changes may be more relevant in the case of Brisbane, since the changes in area had no impact on the underlying political dynamics, and therefore represented pure noise. Finally, the areas themselves were somewhat arbitrarily set by the ABS. As such, only very small boundary changes were ignored. While 110 areas were deemed to have stable enough boundaries to be kept as single statistical local areas, another 62 areas were amalgamated into 24 larger areas with more stable boundaries. Below is a list of those areas that were merged due to large boundary changes.

Suburbs were not amalgamated where it was determined that the boundary changes only incorporated parkland. Such areas included Taringa, Toowong and The Gap, all of which lost area to either Brookfield or Upper Brookfield. Since all of this area is Enoggera State Forest, the SLAs were not merged, but the physical areas (in sq. km) were adjusted in each period so that they matched the 2011 boundaries. Likewise Upper Mount Gravatt, Mount Gravatt, and Nathan were all kept separate (and again the physical areas adjusted as appropriate), since all boundary movements only resulted in an exchange of parkland around Griffith University campus.

Two areas offered significant difficulty. The first was the area around Eaglefarm. In this case, significant areas were relocated from Nundah and Northgate, to Eaglefarm. Most of this area appears to consist of golf courses and parkland, but some does include possible areas of low density dwellings. There is no way to remove this problem without losing large amounts of data, so no attempt has been made to correct for this issue.

The other area in question is the area surrounding Carindale. Parts of several statistical local areas were removed to create Carindale, including some suburban and some fringe areas. While suburbs have been amalgamated to reduce the impact of these boundary shifts as much as possible, Carina, and Carina Heights both lose fairly large areas, and Mansfield gains a fairly substantial area, and these have not been adjusted for.

Suburbs for Brisbane have been classified as city, suburbs or fringe based on their classifications in 1954, since in 1911 Brisbane had very few suburbs. 'City' are those areas identified as high density in 1954. 'Fringe' includes those areas not listed as suburbs in 1954, except for Chermside West, which (due to its being merged with Aspley was included as a suburb). Those classified as suburbs were also suburbs in 1954. Exceptions are Yeerongpilly, and Nudgee and Northgate, which were reclassified as fringe (due to mergers with the fringe areas Rocklea and Nudgee Beach respectively).

List of amalgamations for Brisbane

- Mount Gravatt East and Holland Park (Boundary shifts)
- Ferny Grove and Upper Kedron (Merged in 1976 data)
- Herston and Kelvin Grove (Boundary shifts)
- Moreton Island, Lytton, and Hemmant (Boundary shifts between Moreton Island and Lytton; Lytton and Hemmant merged in 1991 data)

- Rocklea and Yeerongpilly (Boundary shifts)
- Fitzgibbon and Taigum (merged in later data)
- Chandler, Capalaba West, and Burbank (Chandler and Burbank merged in 1976 data; Chandler and Capalaba West merged in 2006 data)
- Parkinson, Drewvale, Berrinba, Karawatha and Stretton (Merged in 1976 data)
- Darra, Sumner, Riverhills, Middle Park, and Westlake (all except Darra merged in 1976 data; Darra and Sumner merged in 1991 data)
- Pallara, Heathwood, Larapinta, and Willamwong (Merged in 1976 data)
- Richlands, Ellen Grove, Doolandella, and Inala (Ellen Grove and Doolandella merged in 1976; boundary shifts)
- Upper Brookfield Brookfield, Bellbowrie, Anstead, and Pullenvale (Anstead, Pullenvale, and Upper Brookfield merged in 1976 data; Brookfield and Upper Brookfield merged in 2006 data)
- Aspley and Chermside West (Boundary shifts)
- Wooloowin, Kedron (Boundary shifts)
- Lower Nudgee, Cribb Island, EagleFarm, Pinkenba, Nudgee, Nudgee Beach, and Banyo (Nudgee and Nudgee Beach merged in 1976 data; Pinkenba and Eagle Farm merged in 1991 data; other boundary shifts)
- Nundah and Northgate (Boundary shifts)
- Chermside and Geebung (Boundary shifts)
- Acacia Ridge and Sunnybank Hills (Boundary shifts)
- Wishart, Rochedale, and Eight Mile Plains (Boundary shifts)

- Ransome, Gumdale, Belmont, Makenzie, Wakerley, and Carindale (Wakerley and Ransome merged in 1976 data; Ransome and Gumdale merged in 2006 data. Other boundary shifts).
- Cannon Hill and Carina (Boundary shifts)

The electronic maps also provided more accurate estimates of the physical area of the statistical areas than the reported statistics (less rounding). These were therefore used for the study. However, since the maps were not available for 1976, these areas were drawn from the collection district file for 1976 provided by the Australian Data Association. However, this resulted in some large and unexplainable aberrations. Therefore, the results were checked using the 1981 areas in place of the 1976 areas, and they were unaffected in that case.

Chapter 4

Defined Benefit vs Defined Contribution Pensions

4.1 The problem of pension underfunding

Recent years have seen a growth in controversy over State and local government employee pension plans. Many governments have insufficient assets set aside to meet their expected future payments, resulting in their pensions being “underfunded”. The present value of these unfunded liabilities to government employees was recently estimated at $4.4 trillion (Healey and Hess, 2012).

The underfunding problem stems from a confluence of factors including a failure of government budgeting, slower than expected market growth, and a century of positive shocks to longevity. However, it has also raised serious questions about the design of government pension systems. Specifically, there have been questions raised as to the efficiency, fairness, and long-run cost of these systems.

Government pension systems at the State and local government level are usually defined benefit pensions. Defined benefit pensions guarantee employees an ongoing pension of a fixed nominal or real amount for the entirety of their post-retirement lifetime. Government defined benefit pensions stand in stark contrast to the defined contribution systems that predominate in the private sector, which are essentially incentivized savings systems. The relative generosity of government pensions, combined with questions regarding their ongoing solvency, has led to suggestions that governments should switch to providing the defined contribution plans (Kiewiet, 2010).

We argue that in a competitive market the expected utility of remuneration should theoretically be equal regardless of the form of the pension provided. Since defined benefit pensions are less risky from the employee's perspective than defined contribution plans, then assuming that employees are risk-averse, one should expect that a move toward defined contribution plans would *increase* the total remuneration bill.

We are therefore left with two conflicting arguments. On the one hand, defined contribution plans have been proposed as a cheaper alternative to defined benefit plans. On the other hand, the relative risk model implies that defined contribution plans are more expensive for employers. There are clear implications in either case for the underfunded status of government pensions. If remuneration is indeed higher under a defined contribution system then the problem becomes one of governance rather than generosity. In fact, costs may even increase, especially in the short-to-medium term when governments would be required to simultaneously pay the existing defined benefit holders and pay higher salaries to defined contribution plan members.

The State of Nebraska offers a unique opportunity to compare pension types. Nebraska, unlike any other State, implemented a compulsory defined contribution plan for most State and county employees as long ago as 1967. We use econometric matching techniques to compare the salaries paid in Nebraskan counties with salaries paid in counties Kansas (matched on observed variables such as size, industry, GDP, etc) which offer a defined benefit plan. This allows us to compare the costs of two different types of pension plans over a period of close to 30 years. We limit ourselves to considering the pension component of the plan and do not consider health care provision, as this would add an additional level of complication to the analysis.

The first section of this paper outlines the types of pensions and their history focussing on the counties we are considering. We go on to consider the arguments for and against each type of plan. We put our discussion in the context of the literature on State finance and individual risk and return profiles. We go on to compare the total remuneration package for various types of employees in Nebraska and Kansas. We find that there is actually very little difference between compensation under the two different types of plans. Even when matching counties based on a variety of different demographic and economic characteristics, we find there is only a very small difference between the costs of the two different plans. Finally, we find that there is no consistency in the relationship between the total remuneration package for different subcategories of employees or across time.

4.2 History of pensions

4.2.1 Types of pensions

Pensions are divided into three broad categories based on the benefits they offer: defined contribution, defined benefit, and hybrid[1]. Defined contribution pensions usually require a certain proportion of the employee's salary to be placed in a separate 'fund', which is then invested in some subset of the market portfolio. Defined contribution pensions may be compulsory (as they are for all employees in Australia), or voluntary (as in the case of US 401K accounts), and contributions are usually encouraged by some form of tax incentive. Defined contribution pensions can perform some or all of the following functions:

- Assist uninformed individuals to save;
- Reduce the moral hazard created by means-tested public pensions, which may otherwise encourage individuals to reduce savings rates;
- Overcome behavioral biases in savings, such as hyperbolic discounting; and
- Provide redistribution to workers through reduced taxes.

On the other hand, defined benefit pensions consist of an annual payment from employers to employees for the entirety of their post-retirement lifespan. Payments

[1]Hybrid plans include both a defined benefit and defined contribution component, and therefore tend to provide only partial insurance — for instance, they may provide a lifetime pension that is based on the stock market at the time of retirement or only insures a small collective of people up to the point where all contributions are spent.

may or may not be indexed to inflation. These plans can perform most of the same functions as defined contribution plans, but also provide implicit longevity insurance.

From the employer's perspective, there are important differences in the timing of payments for each type of pension. Since the payments for defined benefit pensions do not fall due until the employee retires and starts to collect their pension, it is possible that at any one point the employer may not have sufficient assets to fund all of their future liabilities and remain solvent.[2] In addition, the lifespan of the employee is uncertain, so the employer must rely on actuarial assumptions (which are open to some degree of manipulation) when assessing the present value of these liabilities. As a result, employers may end up in a position where they have insufficient assets to meet the present value of their defined benefit pensions (at which point the pensions are 'underfunded').

In contrast, defined contribution plans require employers to pay the full value of their contribution to the pension fund at the time that the employee renders their services. This payment, including all interest, belongs to the employee, and therefore does not need to be accounted for by the employer (either as an asset or a liability). Consequently, defined contribution plans can never be underfunded, but their final value to the employee is subject to market and longevity risk.

4.2.2 History of pensions in Kansas and Nebraska

In the US, the public sector preceded the private sector in both the generosity and growth of pensions.[3] Major cities such as New York established plans for some public sector workers prior to 1900. These plans were generally funded entirely by worker contributions, and operated more-or-less as insurance schemes in the event of disability (Clark et al., 2003).

However, as plans were generally created and administered by local governments

[2]Theoretically, retired employees could just be given a share of the employer's revenue from the period in which they collect their pension. In practice, under the current legal framework firms and governments are required to establish contribution funds to fund their future pension payments.

[3]This is not to suggest that there were no private sector pensions. The Pennsylvania Railroad, for instance, offered a pension to its employees before the US government (US Congress Committee on Civil Service Retirement1919).

there was no consistency in the degree to which pensions were offered. The States were much later to develop plans, especially for employees that fell outside classic public service roles like teaching and public safety (Clark et al., 2003).[4] The States now supervise many of the county pension funds in the US[5].

For example, the Kansas Public Employees Retirement System (KPERS) operates the defined benefit plans for most State and County employees in Kansas. KPERS commenced in 1962, but participation by counties was elective until 1970, after which time all counties were required to use the fund. Police officers, firefighters, judges, and school teachers have plans that are administered separately, but most state and local employees are administered through a single fund.[6] The contribution rates have alternated between being set by legislation (e.g. 1962, 1970) and by the KPERS board of trustees (KPERS 2013b; 2013a).

While firms have moved away from defined benefit plans and toward defined contribution plans, local government defined benefit plans like KPERS have proven remarkably persistent (Fore, 2001). In fact, Nebraska actually *changed back* from a defined contribution plan to a defined benefit\hybrid "cash balance" plan in 2002. Only three states have ever had compulsory defined contributions plans for their general state employees: Nebraska (1967-2002), Michigan (since 1997), and Alaska (since 2005) (Munnell et al., 2014; Snell, 2010)[7]. While other states have had optional defined contribution pensions, and many have introduced hybrid programs of various descriptions, there has been only another handful of specific groups of government employees who have been enrolled in a mandatory defined contribution pension.

Nebraska was one state which enrolled its county employees in a compulsory defined contribution plan for a significant period of its history. While Nebraska also

[4]The US government, in turn, created a defined benefit plan for federal government employees in 1919.

[5]While most of the existing pension funds are supervised locally (Useem and Hess, 2001), only about 10% of employees were covered by locally operated plans for general local government employees in 1991 (Mitchell et al., 2001).

[6]School employees are also covered by KPERS, although these were generally not included until after 1970, and after 1988 were subject to a different contribution rate from other local government employees.

[7]DC has also had a compulsory defined benefit plan since 1987.

has defined benefit plans for judges, state police, and teachers, the Nebraska Public Employees Retirement System (NPERS) was established in 1965 as a defined contribution plan, and counties who did not yet have pension systems were allowed to join. In 1973, those counties that were covered by the former Retirement System for Nebraska Counties were brought into the plan, and all counties except Douglas and Lancaster are now required to belong to the system (Kramer and Partner, 2013). Nebraska's defined contribution plan was both compulsory and in effect for a significant stretch of time, until a defined benefit portion was added in 2001. NPERS requires county employees to contribute 4% of their salary to their defined contribution funds while Nebraskan employers contribute an additional 150% of Member contributions (Buck, 2000).

NPERS and KPERS are similar in terms of the timing of their implementation, and the types of employees which they cover. On the other hand, they are very different in terms of the types of benefits they provide. In the next section, we consider whether one or the other type is likely to be more costly to the county members.

4.3 Pension cost

While most pension funds remained well funded through the 1990s Fore (2001), the years 2000 since have seen a growth in the number of underfunded pensions. Many state and local governments ended up in 'crisis' after 2007, and underfunded local government pension liabilities increased to $4.4 trillion nationwide (Healey and Hess, 2012). In some cases, underfunding has resulted in cuts to local government services, as more current funding is diverted to pay for the accumulated liabilities Kiewiet and McCubbins (2014). In addition, crises in bankrupt cities such as Detroit have actual resulted in some of these liabilities being reduced by the courts, shifting the cost back on to the employees who were originally promised the entitlement Kiewiet and McCubbins (2014).

One proposed solution to the underfunding problem is to switch to defined contribution plans (Kiewiet, 2010). A move to a defined contribution plan clearly eradicates the problem of underfunded pension plans, as defined contribution plans are always funded by definition. Defined contribution plans are now favored across the world: they account for most pension assets globally, and are used by both the US and international private sectors and foreign governments (Broadbent et al., 2006).

However, it is not clear whether a shift between pension types would result in lower costs. In addition to the short run costs of such a move[8], it may be that a move to defined contribution plans actually increase costs, if employees need to be compensated for taking on more retirement risk.

4.3.1 Individual risks and returns

On a purely theoretical level, there is a strong argument that a defined contribution plan should result in higher total remuneration for employees than a defined benefit plan, due to its relatively higher risk profile. In its simplest terms, one can model the remuneration of employees under each system as follows.

Employees work for T years, and have a retirement period of R years, which is uncertain. For simplicity, we assume that in the case of the defined contribution plan, the employer pays a salary and then places an additional component into an asset fund that has an expected return of α. Then, if the utility function is separable:

$$U_{DC} = u\left(\sum_{1}^{T} \delta^t salary_{DC}\right) + E_R\left[\sum_{T+1}^{R} \delta^t u\left(E\left[withdrawals\right]\right)\right]$$

where *withdrawals* represent the optimal withdrawals in each state of the world given the risk attached to lifespan R, the realized market return, and the original payments made by the employer.

In the case of the defined benefit plan, the employer again pays a salary, and at the end of the employee's working career pays them a fixed pension for the remainder

[8] In the short term, a move to a defined contribution plan does not absolve the government of its liabilities for existing underfunded plans, and adds an immediate cost as governments cannot delay contributions to defined contribution plans in periods where the budget is tight. Thus it brings forward the cost of the excess pension liabilities that the government already cannot afford to fund!

of their natural life.

$$U_{DB} = u\left(\sum_{1}^{T} \delta^t salary_{DB}\right) + E_R\left[\sum_{T+1}^{R} \delta^t u\left(pension\right)\right]$$

From these two equations, we can see that under the defined contribution scenario the employee faces two additional sources of risk relative to the defined benefit scenario. The primary source of risk is market risk. If the employee is risk-averse, then to keep expected utility from remuneration constant we must have that:

$$E\left[withdrawals\right] > pension$$

The second source of risk is longevity risk. The employee may die prior to obtaining all their benefits under either system[9], but only in the defined contribution case is it possible to run out of money. If risk-averse, an employee with a defined contribution must either purchase a private market annuity[10] or save more and spend less in order to smooth consumption. Therefore, the sum of the expected withdrawals needs to be higher than the defined benefit pension in order to keep utility constant.

There are other important, and more subtle risks associated with each type of pension. The most important with respect to defined benefit pensions are the possibility of employee termination, uncertainty regarding the wage trajectory (Bodie et al., 1988) and, as more recent events have demonstrated, employer bankruptcy. While we argue that the third is substantially reduced for local governments (though certainly not eliminated as Detroit proceedings have shown) the first two remain possibilities. In fact, if labor markets are sufficiently non-competitive for local government employment employees, their uncertainty regarding termination and wage trajectories should decline, increasing the likelihood that the defined contribution

[9] Life insurance may be offered under either system, and may or may not include the present value of any remaining contributions. However, for the purposes of this paper, we do not consider the complicated process of determining inheritance value.

[10] Private market annuities still leave employees open to market risk at the time of retirement. In addition, it is not necessarily that case that the gains from pooling in a private market would be easily transferred back to the original employer.

plan should be cheaper.

Based on the difference in employee risk we would expect that lifetime remuneration of employees would be higher under defined contribution pensions relative to defined benefit pensions. Empirical work supports this prediction in the private sector: individual accumulation of retirement wealth is higher under defined contribution than under defined benefit plans, but has a greater spread (see, for instance, Samwick and Skinner (2004) and Poterba et al. (2007)). While this does not tell us the comparative cost to the *employer* over time (the result may be driven by higher-than-average market returns), it does suggest that the employee achieves more average compensation under the defined contribution plan as predicted.

Overall, while it is unclear exactly how the risks play out for each employee, our intuition remains the same. If the employee expects to take on more risk under one type of pension than another, they should be compensated by increased remuneration.

4.3.2 Government finances, the holdup problem, and governance

One reason that the risk and return argument may fail is that it does not take into account public sector bargaining. Public sector pensions tend to be more generous than private sector plans across the board (Poterba et al., 2007). This has been attributed to a number of factors, the most commonly cited being the potential holdup problem (Ippolito, 1985).

From a theoretical perspective uncompetitive labor markets should not necessarily alter the relative costs of different types of pensions. While labor unions can theoretically increase employee compensation relative to a free labor market, there is no reason this should only be true in a defined benefit framework. Uncompetitive labor markets can raise the present value of expected remuneration under either pension system. Furthermore, it may simply be that the more generous retirement plans are offset by lower wages during one's working life Hustead and Mitchell (2001).

It could, however, be argued that if public sector workers extract more generous defined benefit pensions due to labor union influence, then local governments should switch to defined contribution pensions. There is indeed evidence of a negative correlation between funding status and union activity (Mitchell and Smith, 1991).

While the net outcome is unclear, government employees appear to demonstrate a preference for defined benefit plans as opposed to defined contribution plans. When Michigan correctional workers were given the opportunity to voluntarily switch from defined benefit to defined contribution plans, only 1.6 percent opted to make the switch (Papke, 2004). Even among those men with only one year of tenure (who had no vested retirement benefits under the defined benefits system as yet), the switching rate was only 6.1percent. This is so incredibly low, that it may simply have been that the only people who switched were those who were seeking alternative employment.

The relative generosity of government pensions may, however, have little to do with their structure. Public sector pensions are more generous than private sector pensions regardless of design (Poterba et al., 2007; Beshears et al., 2010). Administration costs for defined contribution plans are also higher Munnell et al. (2008). Indeed, optimal private sector and public sector pensions need not be the same, as the two types of employers face very different constraints. Private sector employers face significant uncertainty at the firm level. For example, future revenues generated by current employees may be insufficient to fund their remuneration, capital market imperfections may limit financing, demand shocks may occur, and so forth. Firms who provide defined benefit pensions are essentially making employees implicit bondholders of the firm Cooper and Ross (2001), which may not be desirable to *either* the firm (who will then have reduced access to potentially cheaper forms of financing) *or* employees (who have to bear too much firm-level risk).

In contrast, short term solvency is less of an issue in the public sector (Anderson, 2009). Local governments have longer timelines and more certain sources of revenue, through their ongoing ability to tax their resident populations. Even in the event of a short-term shock to finance markets, the fact that in the longer term there is less variability to revenues will mean that, in most cases, governments can easily borrow money to smooth expenditure over time. A stable employer with easy access to finance can exploit pooling across employees to reduce aggregate longevity risk and thereby reduce the gross insurance requirement across all employees.

It could instead be argued that underfunding is less the result of over-generosity, and more the result of poor governance. There is, for example, no evidence of correlation between a State's per capita income (or their ability to tax) and whether and by how much they are underfunded (Fore, 2001). The calculation of present liabilities is incredibly sensitive to small variations in actuarial assumptions (Hustead, 2003). Governments may therefore choose to use an inappropriate discount rate for their asset returns, or fail to accurately account for accumulation of expected future liabilities (Novy-Marx and Rauh, 2009).

Even a move to a defined contribution plan would not necessarily eradicate all governance issues. In the case of Alaska, pension underfunding increased from \$5.7 billion when the change to defined contribution occurred in 2005, to \$12.4 billion in 2013. Local government pensions have also historically been subject to the same governance issues as social security, including the tendency for governments to borrow from their own pension funds to fund other activities, by having their pension funds purchase their own bonds (Clark et al., 2003). This kind of "double dipping" is an additional and often hidden form of underfunding, not included in the underfunding statistics. A move to a defined contribution plan will not change this dynamic, unless the borrowing-from-oneself option is eliminated.

Overall, it is unclear whether the generosity of government pensions is the result of a holdup problem or governance issues more generally, or merely an alternative to paying a larger wage. We also do not know how much of the problem would potentially be mitigated by a change in the design of pension schemes, or how much any savings would be offset if employees demand compensation for a shift in retirement risk. In the remainder of the paper, we examine how the adoption of the different types of pension schemes actually affected gross compensation across similar counties in Nebraska and Kansas.

4.4 Data and econometric method

The State financing literature makes clear that pension liabilities are a very important component of government budgets, and yet it says little about which type of plan is of lower cost in the long run. On the other hand, while the literature on risks and returns clearly outlines some of the plan tradeoffs from the individual's perspective, it gives no answer to the problem of underfunded government pensions. This leaves open the question of whether underfunding is the result of using a more expensive type of plan or an issue of governance and accountability.

Rather than comparing ex-post *returns* from the different pension plans we compare the ex ante *cost* to local county governments. This approach allows us to capture the average expectation of the employees, eliminating individual variations. We use nearest-neighbor propensity score matching (Rosenbaum and Rubin, 1983) to compare counties in Nebraska (defined contribution scheme) and Kansas (defined benefit scheme[11]) to test whether in fact there is a difference in total remuneration. We calculate the propensity scores using a logit regression, based on a number of important county-level characteristics.[12]

The advantage of this approach is that we can leverage the variation in county salaries in a way that we could not at the state level. Since so few states have had defined contribution plans for any substantial period, there is not enough variation in the remuneration of state level employees to separate the effect of the plan from the other characteristics of the state, such as degree of urbanization, population, industry, and GDP. However, since individual counties set their own salaries we can match counties across states and use those to control for county characteristics that are unrelated to the retirement plan. In addition, while there may be endogeneity in the decision to move to a defined contribution plan at the State level, the endogeneity does not exist at the county level as the county is not responsible for the choice of

[11] While states other than Kansas had defined contribution plans over the period we focus on Kansas as the plan is centralized for the whole state, and among the states the Kansas' and Nebraskan economies are relatively similar in size and structure.

[12] We considered using a geographic regression discontinuity design, but this would have required more dense geographical data, especially close to the Kansas/Nebraska boundary.

pension plan.

We obtained data on the number of full-time equivalent employees and total payroll data at the county level for all counties since the 1950s from the Census of Governments. Employee data includes only employees of the county governments, and dependent agencies. Independent agencies, such as school districts, are not included.[13]

We use both the average wage full-time equivalent and the average total remuneration (adjusted to include the employer compensation to the retirement fund) as output variables. We use total remuneration as we have argued in the previous section that there is essentially no difference between the cost to the employer of paying for a wage now and a pension later. The Nebraskan employer contribution rate was 150 percent of employee contributions (Committee, 2010). Since employees were required to contribute 4% of their salaries, we therefore used 6 percent as the mandated contribution figure. The Kansas historical contribution rates were obtained from KPERS annual reports (KPERS 1973-2000). Average remuneration for Kansas and Nebraska are shown in figure

Since employee data is based on full-time equivalents, we include the proportion of part-time employees in our match. In 1985 total number of employees was reported instead of the number of full-time equivalents, so we drop this year entirely from our data set. We adjust the payroll data for inflation using the implicit price deflator obtained from the OECD (2010). To limit the effects of long-run trends in national real wage growth, we manipulate our propensity scores slightly so as to ensure that matching occurs only within the same year.

The county Census of Governments also provided data on counties' annual revenue, expenses, and population. These were included in the basis for the match, in order to control for the size and scope of the county government. We also match on area[14], and 1989 median household income data, also obtained from the United States Census Bureau.

[13]No data was available for the year 1996, so this was omitted.

[14]We use the area provided for the year 1990.

Figure 4.1: Average remuneration for counties in Kansas and Nebraska

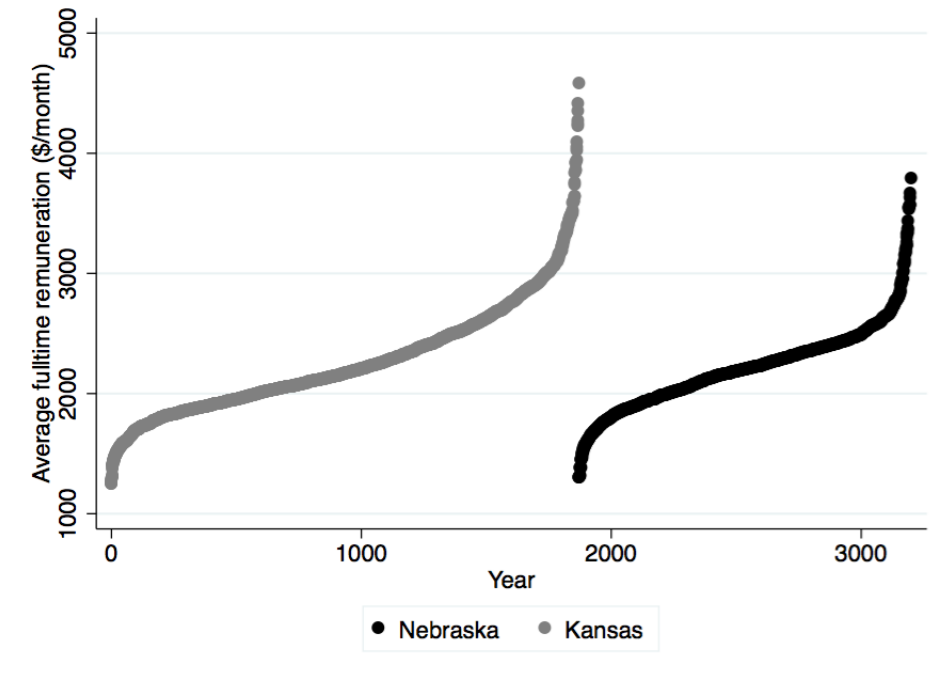

In addition to considering the effect of the choice of plan for county employees as a whole, we also consider its effect by employee types, based on the classifications in the Census of Governments. By looking at subsets of employees we can determine which groups, if any, are benefitting most from the choice of pension structure. We control for the demand for each type of employee by matching on the total number of each type employed by neighboring counties. In addition, we match based on the total number of similar employees (based on the Standard Industrial Classifications) in the private sector across the county and its neighbors for the year 1990, again obtained from the United States Census Bureau.

4.5 Results

We find that there is no appreciable difference between the costs of the two different types of pension plan. After using a logit regression to calculate propensity scores, we then determined a nearest neighbor match for each county in Nebraska[15] Figures 4.2

[15] Kansas has more counties than Nebraska.

Figure 4.2: Ratio of defined benefit to defined contribution wages in matched counties : by year*

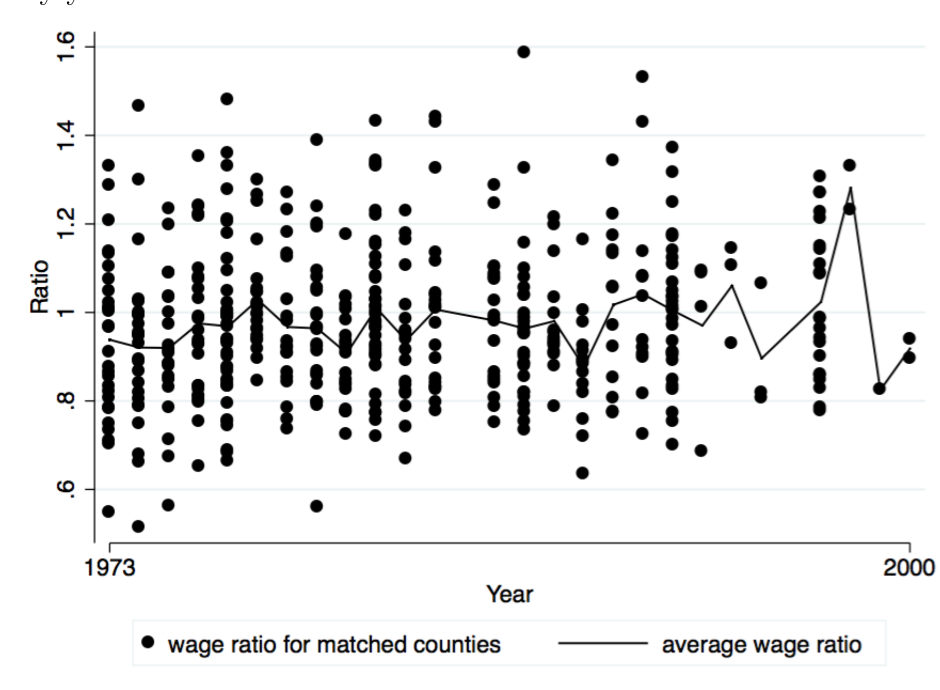

** Each dot point represents the ratio between the average full-time equivalent wage in nearest-neighbor matched counties in Kansas (numerator) and Nebraska (denominator). The average wage ratio takes the annual average across all matched counties.*

and 4.3 graph the ratios of the average full-time equivalent wage and remuneration (limited to the subset of occupations covered by NPERS county plan) for the matched counties. In each case the numerator is the Kansas (defined benefit) county, and the denominator is the Nebraskan (defined contribution) county. It is clear from the graphs that the ratios are concentrated around 1. This pattern suggests that the choice of plan type is not a big contributor to the average cost per full-time equivalent employee.

Table 4.2 further supports our findings. Once again, we use a logit regression to determine propensity scores. We then match counties with their nearest neighbor based on propensity scores, and using the defined contribution plan to represent the treatment, determine the average effect of treatment on the treated. The matching process considerably improves the balance statistic, as demonstrated in table 4.1.

The defined contribution plan is significantly cheaper (the average treatment effect

Figure 4.3: Ratio of defined benefit to defined contribution remuneration in matched counties : by year*

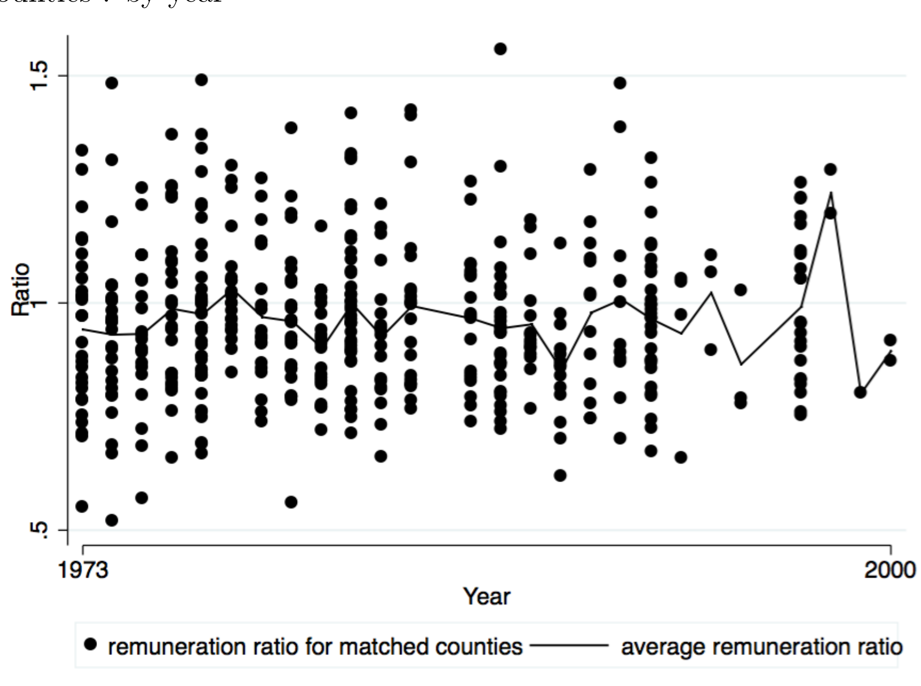

** Each dot point represents the ratio between the average full-time equivalent remuneration (including gross payroll and employer contributions to the employee's retirement fund) in nearest-neighbor matched counties in Kansas (numerator) and Nebraska (denominator). The average remuneration ratio takes the annual average across all matched counties.*

Table 4.1: Balance and balance improvement for matching parameters

	Matched	Mean Treated	Mean Control	%bias	%reduct |bias|	t	p>|t|
Percent employees fulltime	N	91.805	92.597	-12.9		-3.65	0.000
	Y	93.471	92.757	11.6	9.9	1.97	0.049
Med household income '89	N	22490	23573	-29.1		-8.00	0.000
	Y	22512	22703	-5.1	82.3	-0.85	0.394
Total gov revenue	N	3720	10138	-33.3		-8.54	0.000
	Y	3493.2	4503.9	-5.2	84.3	-1.10	0.272
Total gov expenditure	N	3592.3	9895.8	-34.2		-8.78	0.000
	Y	3362	4392.5	-5.6	83.7	-1.13	0.260
County population	N	13795	30692	-35.3		-9.17	0.000
	Y	12800	14218	-3.0	91.6	-0.87	0.386
County population density e	N	26.54	56.834	-25.8		-6.77	0.000
	Y	23.452	24.815	-1.2	95.5	-0.28	0.778

Note: Some subsets of employees do not have data available for all periods. Some subsets of employees were not hired by all local governments in all periods. The number of observations varies with the availability of data and the number of employees hired.

on the treated is negative and significant), but the difference represents only a very small proportion of the real wage (less than 5 per cent of the constant). Once the annual employer contribution to the fund is accounted for, the benefit drops even further, to below 2.5 per cent of the constant. The result is the same for both the average full-time wage across all sectors[16] (the first two columns) and the average full-time equivalent in sectors definitely covered by the considered plans (second two columns). This result is robust to the type of match used (nearest-neighbor, k-neighbor, kernel, and with or without replacement).

Table 4.2: Impact of defined benefit contribution on wages and remuneration

	(1) Fulltime wage	(2) Fulltime remuneration	(3) Wage	(4) Remuneration
Impact of defined contribution	-88.23***	-56.00***	-90.47***	-59.58***
fund ($/month)	(14.54)	(14.87)	(13.71)	(14.08)
Constant	2163.1***	2255.4***	2127.7***	2219.0***
	(9.361)	(9.574)	(8.831)	(9.063)
Observations	3146	3146	3145	3145

Note: Some subsets of employees do not have data available for all periods. Some subsets of employees were not hired by all local governments in all periods. The number of observations varies with the availability of data and the number of employees hired.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Our regression does not take into account underfunding. KPERS local plan was 97.56 per cent funded as at the year 2000, with a net unfunded liability of $36 million KPERS (2000). If we assume that there this cost should have been spread equally spread across counties and years in the previous two decades, we would need to increase wages under the defined benefit plan by an average of around 0.15 per cent.

Even though the aggregate cost of the defined benefit contribution is smaller, table 4.3 indicates that the defined contribution plan is significantly more expensive for

[16]Teachers are not included in the data for Nebraska or Kansas because they are employed by school districts rather than counties. The main concerns here are police and fire employees, which in Kansas are covered by a different plan with a slightly different contribution rate.

some sub-sectors of employment.[17] While the results for each sector use a different set of variables to match counties (to include regional public and private sector variables that are relevant to the county), the results are also robust to limiting the matching variables to the same subset used in table 1. In all cases, neither size nor significance changes to any notable degree. Once again the results are also robust to the type of matching framework used, and to use of least squares.

In general there is a degree of variation in the impact of the type of plan depending on the type of employment. The results for the remaining employment sectors are presented in appendix 4.7.

Table 4.3: Impact of defined contribution pensions on remuneration in different sectors

	(1)	(2)	(3)	(4)	(5)
	Finance	Environmental	Welfare	Sanitation	Admin
Impact of defined contribution	62.64***	-302.0***	270.2***	747.5***	684.0***
fund ($/month)	(14.70)	(32.92)	(30.72)	(107.7)	(136.0)
Constant	2030.6***	2303.1***	1968.7***	2363.5***	2153.7***
	(9.410)	(21.93)	(24.47)	(21.71)	(27.42)
Observations	3099	1437	1301	1082	1082

Note: Some subsets of employees do not have data available for all periods. Some subsets of employees were not hired by all local governments in all periods. The number of observations varies with the availability of data and the number of employees hired.
** $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$*

Finally, we note that the data covers two periods with very different average market returns. Market returns were relatively flat in the 1970s and early 1980s,

[17] Not all local governments hire all kinds of employees. Further, there are missing data for some counties and some sub-sectors Restricting the subset of counties to include only those that report positive wages for all types of employees presented in table 4.3 leaves only 162 data points. As a result, the number of observations is not the same across all sectors.

but grew rapidly from around 1987 onward. In table 4.4 we split our sample into pre-1987 and post-1987 to determine whether this has any effect on the relationship between the two types of plan. We find that prior to 1987 defined contribution funds were if anything slightly more expensive than defined benefit funds, whereas from 1987-2000 they were more cheaper. We find a similar pattern for most sub-categories of employees, as we note in appendix 4.7.

Table 4.4: Impact of defined contribution pensions on wages and remuneration pre-1987 and post-1987

	(1) FTWage<87	(2) FTRemun<87	(3) FTWage>=87	(4) FTRemun>=87
Impact of defined contribution	11.40	15.87	-202.9***	-136.7***
fund ($/month)	(15.51)	(16.23)	(24.55)	(25.41)
Constant	2021.2***	2138.6***	2335.0***	2396.7***
	(10.13)	(10.60)	(15.51)	(16.06)
Observations	1759	1759	1387	1387

Note: Some subsets of employees do not have data available for all periods. Some subsets of employees were not hired by all local governments in all periods. The number of observations varies with the availability of data and the number of employees hired.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

4.6 Conclusion and further questions

Our findings suggest that to the degree that there is an additional risk factor associated with provision of defined benefit vs defined contribution plans, this has at best a small impact on overall government budgets. On the other hand, there is little evidence to suggest that defined benefit pensions are tied to a significant degree of extraction. A wholesale move to defined contribution plans is unlikely to lower local government pension costs or reduce the burden on taxpayers.

There is evidence of variation in the optimal plan between employee types. This may be due to the ability of various types of employees to bargain with their employers, due to either the strength of their unions or their outside options. Nevertheless, the savings from one plan or another are relatively small, and given the number of employees involved it seems likely that it would be more costly to run multiple plans than the savings would compensate for.

Finally, we see an increase in the relative cost of the defined benefit plan after 1987, when average market returns were higher. This finding is interesting in a number of ways. For one, it appears that employees slightly prefer the defined contribution plan when market returns are higher, but that they are potentially not adjusting sufficiently for market risk. On the other hand, the increase in market returns does not result in a 100% offsetting decline in county employer contributions. The latter finding is encouraging from a policy perspective, since it suggests some degree of long run smoothing.

This paper does not consider health care costs. These are likely to be a significant contributor to funding issues, and should be considered in further research. We also do not consider how aggregate shocks may lead to unanticipated future costs, and the way in which governments would be best placed to deal with them. We believe this is another important avenue for others to pursue.

4.7 Appendix

Additional sub-sectors

There are a number of additional sub-sectors that we did not consider in the results section. We present the results for the remainder of these in Tables 4.5 and 4.6 below. We note that there is insufficient data to measure the effects for sewerage workers.

Table 4.5: Impact of defined contribution pensions on remuneration in different sectors 1

	(1) Highways	(2) Parks	(3) Housing	(4) Libraries	(5) Waste
Impact of defined contribution fund ($/month)	6.823	-122.3	-353.9	-427.3***	711.8***
	(18.61)	(116.8)	(300.2)	(118.8)	(120.4)
Constant	2392.9***	2224.3***	3323.3***	2148.5***	2314.1***
	(11.91)	(32.47)	(191.8)	(87.22)	(22.62)
Observations	3078	401	49	141	963

Note: Some subsets of employees do not have data available for all periods. Some subsets of employees were not hired by all local governments in all periods. The number of observations varies with the availability of data and the number of employees hired.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4.6: Impact of defined contribution pensions on remuneration in different sectors 2

	(1) All Health	(2) Hospitals	(3) Health	(4) Other
Impact of defined contribution fund ($/month)	-83.98**	124.5***	-222.4***	-53.32**
	(32.65)	(46.29)	(38.26)	(24.37)
Constant	2300.2***	2281.8***	2325.5***	2125.1***
	(16.03)	(17.76)	(24.35)	(15.56)
Observations	2026	1671	773	2740

Note: Some subsets of employees do not have data available for all periods. Some subsets of employees were not hired by all local governments in all periods. The number of observations varies with the availability of data and the number of employees hired.

** $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$*

Fire and police

While there are no teachers in the data, there were two other groups who were covered by plans other than the standard county employee plans. The first group includes public safety employees such as police and fire, who were covered by the standard defined contribution plan in Nebraska but were covered by a separate defined benefit Kansas from other county workers[18]. In table 4.7 we show that there is no statistical difference between the wages paid to police and fire workers under either system. Once we add employer contributions to the police and fire plan for Kansas we find that wages are significantly higher for police, but not statistically different (and signed

[18] It should be noted that in Nebraska State troopers were administered separately.

negative) for firefighters (see table 4.8).

Table 4.7: Impact of defined contribution pensions on police and fire wages

	(1) Total Police	(2) Police Officers	(3) Total Fire	(4) Firefighters
Impact of defined contribution fund ($/month)	29.59	29.59	679.8	537.3
	(24.34)	(24.34)	(584.4)	(589.0)
Constant	2610.6***	2610.6***	2392.7***	2535.2***
	(15.52)	(15.52)	(60.59)	(61.41)
Observations	2574	2574	279	276

Note: Some subsets of employees do not have data available for all periods. Some subsets of employees were not hired by all local governments in all periods. The number of observations varies with the availability of data and the number of employees hired.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4.8: Impact of defined contribution pensions on police and fire remuneration

	(1) Police	(2) Officers	(3) Total Fire	(4) Firefighters
Impact of defined contribution	-202.6***	-202.6***	521.4	360.5
fund ($/month)	(26.55)	(26.55)	(660.7)	(661.4)
Constant	3001.2***	3001.2***	2735.5***	2896.4***
	(16.93)	(16.93)	(68.51)	(68.95)
Observations	2574	2574	279	276

Note: Some subsets of employees do not have data available for all periods. Some subsets of employees were not hired by all local governments in all periods. The number of observations varies with the availability of data and the number of employees hired.

** $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$*

Judicial and legal

The second category of employees that was not included in the standard county plans is judges. The standard county plan does not cover judges in either state, but both states offer judges a defined benefit plan. In table 4.9 we show that judicial and legal wages are higher in Nebraska. However, while the contribution rates Kansas are available, the contributions for Nebraska are not paid as a proportion of the wage, so we cannot determine if the same result applies to total remuneration.

Table 4.9: Impact of defined benefit pensions on judicial and legal wages

	(1) Judicial
Impact of Kansas fund (\$/month)	-226.7***
	(64.66)
Constant	2564.6***
	(32.80)
Observations	587

Note: Some subsets of employees do not have data available for all periods. Some subsets of employees were not hired by all local governments in all periods. The number of observations varies with the availability of data and the number of employees hired.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Pre and post 1980

The Reagan era is associated with rapid growth in private market expected returns. Were this to feed into expectations, we would expect the cost of defined contribution plans to drop relative to defined benefit plans. In the results section we noted that there was a clear switch between the remuneration under each type of plan as returns rose. In tables 4.10 and 4.11 we document a similar pattern in many of the employment sub-sectors.

Table 4.10: Impact of defined contribution pension on total remuneration pre-1987

	(1) Finance	(2) Environmental	(3) Welfare	(4) Health	(5) Admin
Impact of defined contribution fund ($/month)	102.7***	-237.9***	304.1***	-122.4***	629.8***
	(16.56)	(50.66)	(33.41)	(39.57)	(116.9)
Constant	1936.9***	2280.4***	1942.1***	2207.1***	1952.6***
	(10.80)	(34.43)	(28.37)	(20.01)	(25.25)
Observations	1757	656	814	1095	450

Note: Some subsets of employees do not have data available for all periods. Some subsets of employees were not hired by all local governments in all periods. The number of observations varies with the availability of data and the number of employees hired.

** $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$*

Table 4.11: Impact of defined contribution pensions on total remuneration post 1987

	(1) Finance	(2) Environmental	(3) Welfare	(4) Health	(5) Admin
Impact of defined contribution	24.53	-357.4***	228.1***	-11.87	775.5***
fund ($/month)	(25.15)	(43.08)	(61.16)	(52.68)	(222.0)
Constant	2145.7***	2321.1***	1993.2***	2405.2***	2295.3***
	(15.67)	(28.21)	(42.85)	(24.90)	(42.35)
Observations	1342	781	487	931	632

Note: Some subsets of employees do not have data available for all periods. Some subsets of employees were not hired by all local governments in all periods. The number of observations varies with the availability of data and the number of employees hired. ** $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$*

Least squares

We present the results from running the same regression in a least squares in table 4.12. [19] There are no discernible cost differences between the two pension funds in this framework.

[19] We do not assume linearity here, we simply include this as a robustness check to ensure that a least squares framework does not produce the opposite result from the matching methodology we have used.

Table 4.12: The impact of defined contribution pensions on wages and remuneration, calculated using an ordinary least squares framework

	(1) Fulltime Wage	(2) Fulltime Remuneration	(3) Wage	(4) Remuneration
Impact of defined contribution fund ($/month)	-7.279	26.78	-20.38	11.96
	(27.56)	(28.86)	(26.39)	(27.64)
Median household income 1989 ($)	0.0332***	0.0349***	0.0292***	0.0307***
	(0.00526)	(0.00549)	(0.00518)	(0.00541)
Total revenue ($)	0.000869	0.00115	0.000754	0.00103
	(0.00181)	(0.00190)	(0.00182)	(0.00190)
Total expenditure ($)	0.00247	0.00186	0.00247	0.00189
	(0.00281)	(0.00287)	(0.00298)	(0.00303)
Population	0.000976**	0.00113***	0.000695	0.000832*
	(0.000411)	(0.000429)	(0.000448)	(0.000461)
Population density	0.138	0.145	0.159	0.167
	(0.138)	(0.143)	(0.176)	(0.181)
Percentage of employees working fulltime			-0.226	-0.250
			(1.960)	(2.060)
Observations	3146	3146	3145	3145

Note: Some subsets of employees do not have data available for all periods. Some subsets of employees were not hired by all local governments in all periods. The number of observations varies with the availability of data and the number of employees hired.

$* \ p < 0.10$, $** \ p < 0.05$, $*** \ p < 0.01$

Bibliography

Allen, R. C. (1983). Collective invention. *Journal of Economic Behavior & Organization 4*(1), 1–24.

Anderson, G. (2009). *The future of public employee retirement systems.* Oxford University Press.

Baland, J.-M. and D. Ray (1991). Why does asset inequality affect unemployment? a study of the demand composition problem. *Journal of Development Economics 35*(1), 69–92.

Beshears, J., J. J. Choi, D. Laibson, and B. C. Madrian (2010). Defined contribution savings plans in the public sector: Lessons from behavioral economics. NBER State and Local Pensions Conference.

Bodie, Z., A. J. Marcus, and R. C. Merton (1988). Defined benefit versus defined contribution pension plans: What are the real trade-offs? In *Pensions in the US Economy*, pp. 139–162. University of Chicago Press.

Broadbent, J., M. Palumbo, and E. Woodman (2006). The shift from defined benefit to defined contribution pension plans—implications for asset allocation and risk management. *Reserve Bank of Australia, Board of Governors of the Federal Reserve System and Bank of Canada.*

Brown, D. J. (1991). Equilibrium analysis with nonconvex technologies. *Handbook of mathematical economics 4*, 1963–995.

Brown, D. J. and G. Heal (1979). Equity, efficiency and increasing returns. *The Review of Economic Studies 46*(4), 571–585.

Buck (2000, August). Benefit review study of the nebraska retirement system. Report, 1200 17th Street Suite 1200 Denver, CO 80202.

Burkhauser, R. V., B. A. Butrica, and M. J. Wasylenko (1995). Mobility patterns of older homeowners are older homeowners trapped in distressed neighborhoods? *Research on Aging 17*(4), 363–384.

Chipman, J. S. (1970). External economies of scale and competitive equilibrium. *The Quarterly Journal of Economics 84*(3), 347–385.

Clark, R. L., L. A. Craig, and J. W. Wilson (2003). *A history of public sector pensions in the United States*. Univ of Pennsylvania Press.

Committee, N. L. R. S. (2010). *Addendum to Legislative Resolution 542 Interim Study Report: Public Pension Obligations.*

Cooley, T. F. and C. LaCivita (1982). A theory of growth controls. *Journal of Urban Economics 12*(2), 129 – 145.

Cooper, R. W. and T. W. Ross (2001). Pensions: theories of underfunding. *Labour Economics 8*(6), 667–689.

Department, K. L. R. (2013a). Kansas legislator briefing book.

Department, K. L. R. (2013b, January). Review of kpers history for the house committee on pensions and benefits.

Dubin, J. A., D. R. Kiewiet, and C. Noussair (1992). Voting on growth control measures: Preferences and strategies*. *Economics & Politics 4*(2), 191–213.

Ellickson, R. (1977). Suburban growth controls: An economic and legal analysis. *The Yale Law Journal 86*(3), 385–511.

Fischel, W. (1980). Zoning and the exercise of monopoly power: A reevaluation. *Journal of Urban Economics 8*(3), 283–293.

Fischel, W. (2001). Homevoters, municipal corporate governance, and the benefit view of the property tax. *National Tax Journal 54*(1), 157–174.

Foellmi, R. and J. Zweimüller (2006). Income distribution and demand-induced innovations. *The Review of Economic Studies 73*(4), 941–960.

Fore, D. (2001). Going private in the public sector. *Pensions in the public sector*, 267–287.

Glaeser, E. L., J. Gyourko, and R. Saks (2005a, February). Why have housing prices gone up? Working Paper 11129, National Bureau of Economic Research.

Glaeser, E. L., J. Gyourko, and R. Saks (2005b). Why is manhattan so expensive? regulation and the rise in housing prices. *Journal of Law and Economics 48*(2), pp. 331–369.

Glaeser, E. L., J. Schuetz, and B. Ward (2006). Regulation and the rise of housing prices in greater boston. *Cambridge: Rappaport Institute for Greater Boston, Harvard University and Boston: Pioneer Institute for Public Policy Research*.

Gottlieb, J. and E. Glaeser (2008). The economics of place-making policies. *Brookings Papers on Economic Activity 2008*(1), 155–239.

Green, B. and F. Shaheen (2014). Economic inequality and house prices in the uk. NEF Working paper.

Green, H. J. (1961). Direct additivity and consumers' behaviour. *Oxford Economic Papers*, 132–136.

Gyourko, J. (2009). Housing supply. *Annual Review of Economics 1*(1), 295–318.

Hamilton, B. W. (1978). Zoning and the exercise of monopoly power. *Journal of Urban Economics 5*(1), 116–130.

Hamnett, S. and R. Freestone (1999). *The Australian Metropolis: A Planning History*. Allen & Unwin Australia.

Healey, T. J. and C. Hess (2012). Underfunded public pensions in the united states. *Harvard Kennedy School M-RCBG Faculty Working Paper No 8*, 13.

Hustead, E. (2003). Determining the cost of public pension plans. *Pensions in the Public sector*, 218–240.

Hustead, E. C. and O. S. Mitchell (2001). *Public sector pension plans.* Philadelphia: University of Pennsylvania Press.

Ippolito, R. A. (1985). The economic function of underfunded pension plans. *Journal of Law and Economics 28*(3), pp. 611–651.

Kiewiet, D. R. and M. D. McCubbins (2014). State and local government finance: The new fiscal ice age. *Annual Review of Political Science 17*, 105–122.

Kiewiet, R. (2010). The day after tomorrow: The politics of public employee benefits. *California Journal of Politics and Policy 2*(3).

KPERS (1973). Comprehensive annual financial report of the kansas public employees retirement system fiscal year ended 1973.

KPERS (1974). Comprehensive annual financial report of the kansas public employees retirement system fiscal year ended 1974.

KPERS (1975). Comprehensive annual financial report of the kansas public employees retirement system fiscal year ended 1975.

KPERS (1976). Comprehensive annual financial report of the kansas public employees retirement system fiscal year ended 1976.

KPERS (1977). Comprehensive annual financial report of the kansas public employees retirement system fiscal year ended 1977.

KPERS (1978). Comprehensive annual financial report of the kansas public employees retirement system fiscal year ended 1978.

KPERS (1979). Comprehensive annual financial report of the kansas public employees retirement system fiscal year ended 1979.

KPERS (1980). Comprehensive annual financial report of the kansas public employees retirement system fiscal year ended 1980.

KPERS (1981). Comprehensive annual financial report of the kansas public employees retirement system fiscal year ended 1981.

KPERS (1982). Comprehensive annual financial report of the kansas public employees retirement system fiscal year ended 1982.

KPERS (1983). Comprehensive annual financial report of the kansas public employees retirement system fiscal year ended 1983.

KPERS (1984). Comprehensive annual financial report of the kansas public employees retirement system fiscal year ended 1984.

KPERS (1985). Comprehensive annual financial report of the kansas public employees retirement system fiscal year ended 1985.

KPERS (1986). Comprehensive annual financial report of the kansas public employees retirement system fiscal year ended 1986.

KPERS (1987). Comprehensive annual financial report of the kansas public employees retirement system fiscal year ended 1987.

KPERS (1988). Comprehensive annual financial report of the kansas public employees retirement system fiscal year ended 1988.

KPERS (1989). Comprehensive annual financial report of the kansas public employees retirement system fiscal year ended 1989.

KPERS (1990). Comprehensive annual financial report of the kansas public employees retirement system fiscal year ended 1990.

KPERS (1991). Comprehensive annual financial report of the kansas public employees retirement system fiscal year ended 1991.

KPERS (1992). Comprehensive annual financial report of the kansas public employees retirement system fiscal year ended 1992.

KPERS (1993). Comprehensive annual financial report of the kansas public employees retirement system fiscal year ended 1993.

KPERS (1994). Comprehensive annual financial report of the kansas public employees retirement system fiscal year ended 1994.

KPERS (1995). Comprehensive annual financial report of the kansas public employees retirement system fiscal year ended 1995.

KPERS (1996). Comprehensive annual financial report of the kansas public employees retirement system fiscal year ended 1996.

KPERS (1997). Comprehensive annual financial report of the kansas public employees retirement system fiscal year ended 1997.

KPERS (1998). Comprehensive annual financial report of the kansas public employees retirement system fiscal year ended 1998.

KPERS (1999). Comprehensive annual financial report of the kansas public employees retirement system fiscal year ended 1999.

KPERS (2000). Comprehensive annual financial report of the kansas public employees retirement system fiscal year ended 2000.

Kramer, D. J. and B. H. L. Partner (2013). Nebraska's public pension system. *The Platte Institute for Economic Research*, 5.

Krugman, P. (1991). Increasing returns and economic geography. *The Journal of Political Economy 99*(3), 483–499.

Krugman, P. R. (1979). Increasing returns, monopolistic competition, and international trade. *Journal of international Economics 9*(4), 469–479.

Levine, N. (1999). The effects of local growth controls on regional housing production and population redistribution in california. *Urban Studies 36*(12), 2047–2068.

Matlack, J. L. and J. L. Vigdor (2008). Do rising tides lift all prices? income inequality and housing affordability. *Journal of Housing Economics 17*(3), 212–224.

Matsuyama, K. (1991). Increasing returns, industrialization, and indeterminacy of equilibrium. *The Quarterly Journal of Economics 106*(2), 617–650.

Matsuyama, K. (2002). The rise of mass consumption societies. *Journal of Political Economy 110*(5), 1035–1070.

Mitchell, O. S., D. McCarthy, S. C. Wisniewski, and P. Zorn (2001). Developments in state and local pension plans. *Pensions in the Public Sector*, 11–40.

Mitchell, O. S. and R. S. Smith (1991). Pension funding in the public sector. Technical report, National Bureau of Economic Research.

Mohr, R. D. (2002). Technical change, external economies, and the porter hypothesis. *Journal of Environmental economics and management 43*(1), 158–168.

Munnell, A. H., J.-P. Aubry, M. Cafarelli, et al. (2014). Defined contribution plans in the public sector: An update. Technical report.

Munnell, A. H., A. Golub-Sass, K. Haverstick, M. Soto, and G. Wiles (2008). Why have some states introduced defined contribution plans. *State and Local Pension Plans Brief 3*.

Murphy, K. M., A. Shleifer, and R. Vishny (1989). Income distribution, market size, and industrialization. *The Quarterly Journal of Economics 104*(3), 537–564.

Novy-Marx, R. and J. D. Rauh (2009). The liabilities and risks of state-sponsored pension plans. *The Journal of Economic Perspectives*, 191–210.

OECD (2010). Main economic indicators - complete database.

of Queensland, E. C. (2012, November). Details of polling at brisbane city council quadrennial elections.

of Statistics, A. B. (2010, September). Measures of australia's progress, 2010 cat. no. 1370.0.

of Statistics, A. B. (2015, March). Regional population growth, australia, 2013-14 cat. no. 3218.0. Catalogue.

Ortalo-Magné, F. and A. Prat (2007). The political economy of housing supply: homeowners, workers, and voters. *STICERD–Theoretical Economics Paper Series 514*.

Ortalo-Magné, F. and S. Rady (2002). Tenure choice and the riskiness of non-housing consumption. *Journal of Housing Economics 11*(3), 266–279.

Panel, I. L. G. R. (2013). Final report of the independent local government review panel. http://www.localgovernmentreview.nsw.gov.au/documents/LGR/Revitalising

Papke, L. E. (2004). Pension plan choice in the public sector: The case of michigan state employees. *National Tax Journal*, 329–339.

Piketty, T. and E. Saez (2006). The evolution of top incomes: a historical and international perspective. Technical report, National Bureau of Economic Research.

Poterba, J., J. Rauh, S. Venti, and D. Wise (2007). Defined contribution plans, defined benefit plans, and the accumulation of retirement wealth. *Journal of Public Economics 91*(10), 2062–2086.

Quigley, J. and L. Rosenthal (2005). The effects of land use regulation on the price of housing: What do we know? what can we learn? *Cityscape*, 69–137.

Quigley, J. M. and S. Raphael (2004). Is housing unaffordable? why isn't it more affordable? *The Journal of Economic Perspectives 18*(1), 191–214.

Quigley, J. M. and S. Raphael (2005). Regulation and the high cost of housing in california. *The American Economic Review 95*(2), pp. 323–328.

Romer, P. M. (1986). Increasing returns and long-run growth. *The Journal of Political Economy*, 1002–1037.

Rosenbaum, P. R. and D. B. Rubin (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika 70*(1), 41–55.

Saez, E. (2009). Striking it richer: The evolution of top incomes in the united states (update with 2007 estimates).

Saiz, A. (2010). The geographic determinants of housing supply. *The Quarterly Journal of Economics 125*(3), 1253–1296.

Samwick, A. A. and J. Skinner (2004). How will 401 (k) pension plans affect retirement income? *American Economic Review*, 329–343.

Sen, A. (1973). *On economic inequality.* Oxford: Clarendon Press.

Service, U. S. C. S. C. and C. o. Retrenchment (1919). *Retirement of Civil Service Employees: Hearings ... on S. 1699, a Bill for the Retirement of Employees in the Classified Civil Service, and for Other Purposes ...* U.S. Government Printing Office.

Snell, R. (2010). State defined contribution and hybrid pension plans. In *National Conference of State Legislatures, June.*

Suzuki, T. (1995). Nonconvex production economies. *Journal of Economic Theory 66*(1), 158–177.

Suzuki, T. (1996). Intertemporal general equilibrium model with external increasing returns. *Journal of Economic Theory 69*(1), 117–133.

Suzuki, T. (2009). *General Equilibrium Analysis of Production and Increasing Returns.* Series on Mathematical Economics and Game Theory, 4. World Scientific Publishing Company, Incorporated.

Thorson, J. (1996). An examination of the monopoly zoning hypothesis. *Land Economics*, 43–55.

Useem, M. and D. Hess (2001). Governance and investments of public pensions. *Pensions in the Public sector*, 132–152.

Verba, S., K. L. Schlozman, H. E. Brady, and H. E. Brady (1995). *Voice and equality: Civic voluntarism in American politics*, Volume 4. Cambridge Univ Press.

White, M. (1975). Fiscal zoning in fragmented metropolitan areas. *Fiscal zoning and land use controls*, 31–100.

Wolff, E. N. (2010). Recent trends in household wealth in the united states: Rising debt and the middle-class squeeze-an update to 2007.