

Discrete Choice, Complete Markets, and Equilibrium

Simon Mongey

Federal Reserve Bank of Minneapolis and NBER

Michael E. Waugh

Federal Reserve Bank of Minneapolis and NBER

March 2025

ABSTRACT

We study the welfare properties of discrete choice economies, and then use our results to re-evaluate objectives of applied work: how optimal policy exercises should be conducted, and how welfare effects of price changes should be measured. We characterize the allocations that emerge in general equilibrium economies populated by households with random utility preferences that make discrete consumption, employment or spatial decisions. First and second welfare theorems illustrate that in the absence of complete markets, discrete choice economies are generically inefficient. This generic inefficiency is consequential for two central objectives of applied research. First, the welfare gains of optimal policy designed to address externalities can be significantly over-stated due to market incompleteness. Second, standard formulas for the measuring welfare effects of price changes are invalid due to market incompleteness. We use our theoretical results to provide quantitative and empirical solutions to these issues.

JEL Classification: D52, E2, R13, F1

Email: michael.e.waugh@gmail.com, simonmongey@gmail.com. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System. For useful conversations and feedback we thank Manuel Amador, Andrew Atkeson, V.V. Chari, Roozbeh Hosseini, Chris Sleet, Rob Shimer, and Pierre-Olivier Weill.

In this paper, we study the efficiency properties of competitive equilibria of economies populated by households with preferences of the form

$$u^i(c_j^i) + \xi_j^i,$$

where j is a finite set of goods / locations / occupations and ξ_j^i is an idiosyncratic random variable specific to household i .¹ *Discreteness* is the restriction that households can choose only one type of good to consume, location to live in, occupation to work in, and so on.

To say these models have proven to be useful in many applications is an understatement. Discrete choice behavior in general and logit choice behavior in particular is the building block of demand analysis across many fields (for example McFadden, 1974; Anderson et al., 1992; Berry et al., 1995). In spatial settings, preferences of this form have proven useful for studying location and migration decisions (for example Kennan and Walker, 2011; Kline and Moretti, 2014; Diamond, 2016; Redding and Rossi-Hansberg, 2017). They have also been extended to micro-founded supply decisions in labor (for example Rogerson, 1988; Card et al., 2018; Berger et al., 2022). They have been used in static and dynamic contexts. As in all applied theory, key exercises in these literatures are the measurement of the welfare consequences of policy counterfactuals and measurement of welfare effects of shocks. A necessary input to such exercises is an understanding of the welfare properties of these economies.

In this paper, we (a) aim to provide a complete understanding of the welfare properties of these economies, then (b) show how these results provide a new understanding of optimal policy exercises and measurement of welfare effects of price changes. In part (a), we characterize efficient allocations in discrete choice economies, including novel expressions for efficient choice probabilities. We then show that competitive equilibrium choice probabilities are not efficient. Next, we show how complete markets restore efficient choice probabilities, revealing that the generic inefficiency in discrete choice economies is due to incomplete markets. This accommodates statements of the first and second welfare theorems, we then establish that the social welfare function used in applied work is maximized under complete markets.

In part (b), we show how understanding this generic inefficiency is key to interpreting optimal policy exercises and measurement of the welfare effects of prices changes. We describe two novel issues that arise once we recognize that markets are incomplete. First, we show that optimal policies designed to resolve other inefficiencies in the economy may inadvertently address market incompleteness, distorting the assessment of how policies improve welfare. Second, we show that welfare effects of price changes can be measured by standard formulas under complete markets, but such formulas fail in the empirically relevant case of incomplete markets. In both cases we provide concrete advice for applied work as to how to resolve these issues, informed by our theoretical results in part (a).

¹The main text focuses on the additive random utility model. In Appendix E we study the multiplicative random utility model. Our core results are not sensitive to differences between these models.

Theory. The first part of our paper describes a general choice framework. We consider a class of utility functions that allows us to connect with how these models have been used in studies of consumer demand, spatial economics, or labor supply. We don't appeal to functional forms on the distribution over the taste shocks ξ_j^i or utility function $u^i(\cdot)$. Thus, our results are not specific to, say, the popular type 1 extreme value distribution. We model the supply side of the economy as a competitive endowment economy, which we extend to a competitive production economy in our applications. This choice is deliberate and designed to abstract from inefficiencies that might arise from imperfect competition in product markets. We then proceed to characterize several allocations.

First, we characterize the allocation that is usually solved for in random utility, discrete choice models. We call this the *standard allocation*, but our characterization illustrates the following point: that ex-ante identical households (before the taste shocks are realized) ex-post value resources differently, depending upon their choice. For example, someone choosing a high priced commodity / expensive location will have a high marginal utility of consumption, compared with someone choosing a low priced commodity / cheap location. This observation suggests there are ex-ante Pareto improving trades that could be made either through a market arrangement or by a social planner. We emphasize that this is an *incomplete markets* allocation.

Second, we characterize the complete markets allocation and Pareto efficient allocations. We demonstrate the equivalence between the two. This provides a first and second welfare theorem for discrete choice economies. We then show that every Pareto efficient allocation corresponds to the maxima of a problem in which a planner maximizes a standard social welfare function under some vector of social welfare weights. These results are important contributions to the literature for two reasons. First, these results provide a benchmark for understanding welfare improving policies even in the absence of complete markets. Second, the allocations that we characterize maximize the benchmark social welfare function used by economists to evaluate welfare gains from counterfactuals in discrete choice economies, often described as “the expectation of the max” (Train, 2009; Anderson et al., 1992).² In other words, given the objective function researchers use to evaluate policies, we have characterized the allocation that maximizes that objective function.

What does “completing markets” look like? The complete markets economy allows households to trade a complete set of contingent claims on every possible realization of the taste shocks. These asset trades essentially allow the household to face one unified budget constraint rather than a budget constraint that applies, state by state as in the standard setting, or choice by choice in the discrete choice setting. The unification of the household's budget constraint allows the household to equalize the marginal value of its resources across states and the discrete

²See, for example, Busso et al. (2013, p. 902, paraphrased): “Denote the workers' social welfare as $V = E_{\xi} [\max_j \{u_j^i + \xi_j^i\}]$, where the expectation is defined over the ξ_j^i terms.”

choices induced by those states. Thus, the household is able to set consumption so that marginal rates of substitution (across the discrete commodities) equal marginal rates of transformation or relative prices. So even though households choose only one commodity to consume ex-post, households ex-ante choose consumption plans in the same way as in a model where all goods are consumed simultaneously.

A first novelty of the efficient, complete markets allocation is a fundamentally different rule describing the choice of which good or location to select. In the standard setting, the rule describing which good to choose is \max over $u^i(c_j^i) + \xi_j^i$. This is *not* the case with complete markets, any other Pareto efficient allocation or the social welfare maximizing allocation. The optimal rule under complete markets is \max over $u^i(c_j^i) + \xi_j^i - u^{i'}(c_j^i)c_j^i$. The subtraction of the additional term $u^{i'}(c_j^i)c_j^i$ is novel and represents the private cost associated with a choice. Under complete markets households can internalize this cost via their unified budget constraint. This is efficient, in that a planner also internalizes this cost, but instead via their resource constraint.

A second novelty of the complete markets allocation is the nature of contracts that are actually traded. In equilibrium, only a subset of contingent claims is valued and traded, even though we started with a complete set of contingent claims on every possible taste shock realization. The traded claims are indexed only by the choice j , not the taste shocks ξ_j^i . We term these contingent claims *Arrow vouchers*, as they pay off when a commodity choice is chosen. This result suggests that a far simpler class of contracts could accommodate an efficient allocation in a decentralized market.

Applications. The second part of our paper provides constructive proposals for how to use the insights developed thus far to do better applied welfare economics. We do not argue that complete markets are empirically relevant, however we show that the absence of such markets interferes with the evaluation of policy counterfactuals and welfare effects of price changes. First, we show how the interaction of market incompleteness with other inefficiencies (e.g. productivity spillovers) clouds the evaluation of optimal policies designed to address the latter (e.g. production subsidies). We provide a solution to this issue, showing how complete markets can be used as an analytical tool to neutralize market incompleteness and keep the focus on other distortions a researcher is studying. Second, we show how standard formulas used to measure the welfare effects of price changes do not carry over to discrete choice economies. We provide a solution to this issue, showing via analytical results how demand elasticity estimates should be incorporated in the computation of the first order welfare effects of price changes.

Optimal Policies. We consider two numerical examples designed to capture (i) industrial policy in the presence of industry-production spillovers (e.g. Electric vs Combustion car industries), (ii) spatial policy in the presence of location-production spillovers (e.g. San Francisco

(knowledge hub) vs Rest-of-US).³ In both cases a planner with access to a budget-neutral tax/-subsidy wishes to resolve the externality. However, with incomplete markets, subsidies also have the effect of reducing dispersion in prices across goods, or wages across locations. The resulting equalization of marginal utilities across choices of goods / locations improves welfare via the inefficiencies stemming from incomplete markets, not the inefficiencies stemming from the externality. In both cases, this additional benefit implies a high optimal subsidy with large welfare gains. It is immaterial whether these are viewed as welfare gains via insurance from an ex-ante perspective, or redistribution from an ex-post perspective. When computed under *complete markets*, the optimal subsidy addresses only the externality and hence has the same sign, but is much smaller and the welfare gains much smaller, too. Hence we advocate for policy analysis under complete markets as a way of neutralizing the welfare gains that come from addressing market incompleteness, which, in most applied contexts is not the focus of the exercise being considered.

Welfare Effects of Price Changes. Our second application concerns applied empirical measurement of the welfare effects of price changes. A famous result is that equivalent variation to a first order is determined by initial expenditure shares and price changes (Deaton, 1989), while elasticities show up only to a second order. We show that in our discrete choice environment under incomplete markets and arbitrary preferences and taste distributions, sufficient statistics for equivalent variation to a first order are expenditure shares, price changes *and* elasticities of demand. Heterogeneity in price sensitivity (as emphasized recently in Auer et al. 2022) is therefore a first-order determinant of welfare in discrete choice models. In contrast, when markets are complete, we show that the “standard” formula is obtained. For applied research, these results have two implications. First, we have generalized results like those in Deaton (1989) to hold in discrete choice environments, without appeals to functional forms, but under the assumption of complete markets. Second, care should be taken around the interpretation of these formulas since they depend on individual consumption either being (a) truly “a little bit of all varieties” or (b) discrete with complete markets. If one does not want to make either assumption, then a clear outcome for applied work is that demand elasticity estimates are required, and we provide the formulas to apply them.

These exercises contribute a broad lesson for applied research using discrete choice models. A criticism of our work is as follows: (i) complete markets are empirically irrelevant, (ii) hence no one would find it reasonable to implement policies that bring us closer to the efficient allocation. This coincides with a view that, because of this, no paper computes optimal policies to correct market incompleteness. Our contribution is to show that such a luxury is not afforded to a

³This second application has implications for the study of optimal spatial policy which has become central to a large and growing literature (for example Fajgelbaum and Gaubert, 2020; Diamond, 2016; Kennan and Walker, 2011; Artuç et al., 2010; Caliendo et al., 2019).

researcher. In the empirically relevant, incomplete markets environment, a researcher cannot arbitrarily choose whether the welfare gains measured from some counterfactual come from correcting a particular inefficiency that they have in mind, or inadvertently addressing market incompleteness. When choosing spatial or production subsidies measured welfare improves if the economy moves closer to the efficient allocation. Unavoidably, every “empirically relevant” exercise computing an optimal policy will have some component of welfare changes arising from market incompleteness. Our exercises illustrate how our results provide a benchmark from which the effects of market incompleteness can be neutralized, and thus aid the applied researcher in focusing on the direct effects of the policy intervention.

Literature. Several related papers connect closely with the questions we address. The basic idea behind our work dates back to Rogerson’s (1988) work on indivisible labor supply. Rogerson (1988) studies the problem where individual households can only discretely supply their labor; that is they can only work or not. It is often remembered that Rogerson introduces lotteries to convexify the problem, which leads to optimal allocations and aggregation with aggregate labor supply elasticities differing from micro elasticities. What is not often remembered is that complete markets are introduced to eliminate the risk induced by the lottery. In the discrete choice environment we consider, randomization by the taste shock naturally convexifies the problem, but the risk associated with the shock remains. In this paper, we complete markets and illustrate the properties of allocations, demonstrate the optimality of these allocations, and derive results useful for applied welfare analysis in these settings.

Ales and Sleet (2022) study the problem of a government that must raise revenue by taxing income associated with each discrete choice, and characterise the optimal tax structure. Our problem is different; in particular, to assemble welfare theorems we consider a planner that controls everything and we provide decentralization results through ex-ante trades. Ales and Sleet (2022) tackle issues of observability in a Mirrleesian approach. In contrast, we describe the first best allocation, and how policies that maximize welfare operate on the margin of correcting distortions due to market incompleteness.

Recent work by Donald et al. (2023b) is related to this paper. They study a spatial economy, and decompose the first order effects of productivity and transportation shocks across space into terms that include spatial dispersion in marginal utility.⁴ Our paper (i) studies a general choice framework that nests spatial general equilibrium models, (ii) provides standard welfare theorems, which by their nature are global rather than local, (iii) explicitly identifies that dispersion in marginal utility is due to market incompleteness, which is shown by establishing that a complete markets decentralization of the competitive equilibrium delivers the efficient alloca-

⁴Relatedly, in a quantitative spatial framework, Davis and Gregory (2023) highlight the redistributive motives of a planner facing dispersion in marginal utility and the inability to identify these motives from location choice data.

tion, (iv) we do so without functional form assumptions on utility or distributional assumptions on taste shocks.

Small and Rosen (1981) has long served as a go to reference for understanding how to do welfare economics in discrete choice economies. However, their ‘log-sum’ formula is derived under an extra, explicit assumption that the marginal utility of consumption is constant across choices.⁵ Like our application to price changes, Bhattacharya (2021) departs from Small and Rosen (1981) by not making this additional assumption, leading to variation in marginal utility across choices. Unlike our paper, Bhattacharya (2021) restricts attention to a binary choice setting, and focuses on ex-post welfare metrics of price changes in partial equilibrium. Allen and Rehbeck (2019) similarly study welfare and price changes in a discrete choice setting, but restrict attention to the case of no intensive margin of demand. Unlike these studies, our focus is to establish welfare theorems in a much more general setting, with one application being to welfare effects of price changes.

Multiplicative case. In the main text we focus on the additive random utility model, but in the Appendix E we study the multiplicative case. We show that: (i) our key results hold on features of efficient and incomplete markets choice rules, (ii) in the special case of power utility and Fréchet distributed idiosyncratic utility, the aggregate demand system obtained under complete or incomplete markets coincides with the consumption choices of a CES representative household, (iii) the parameters of the CES representative households substantively differ across market structures. This provides the following additional result: that demand can be obtained from a CES representative household, does not go hand-in-hand with efficiency.

Outline. Section 1 lays out the general economic environment we study. Section 2 derives properties of the ‘standard’ incomplete markets allocation. Section 3 derives novel properties of the complete markets allocation and establishes the first welfare theorem. Section 4 derives properties of all Pareto efficient allocations, establishes a social welfare function and the second welfare theorem. Section 5 applies these results to revisit welfare effects of policy in environments with additional inefficiencies. Section 6 applies these results to revisit widely used formulas for welfare effects of price changes. Section 7 applies our results to clarify issues around efficiency and aggregation in discrete choice economies. Section 8 concludes.

1. Economic Environment

This section describes the economic environment.

Households. There are a finite set of N types of households, indexed by θ . Within each type, there is a continuum of ex-ante identical households with mass $\mu(\theta)$. Within each type, house-

⁵To be clear, this is not a feature of their economic environment, but an additional assumption made when developing formulas for the welfare effects of price changes.

holds' names are indexed by $i \in [0, \mu(\theta)]$. Households consume in the economy.

Goods. There are two types of goods. First, there is a single homogeneous, “outside” consumption good c . The price of this outside good is our numeraire and normalized to one. Second, there are finite “differentiated” goods with names $j = \{1, 2, \dots, J\}$; we denote quantities q_j . We restrict each household to choose only one type of the differentiated good to consume; this is the discrete choice aspect of the model.

Preferences. There are several components to preferences. First, households receive a taste shock for each differentiated good $j \in J$. Per the restriction that households must make a discrete choice across differentiated goods, only the shock corresponding to the good consumed actually enters utility.

Following the literature, we model the taste shocks as random variables that are independent and identically distributed across agents i in the population of θ -type households. Importantly, this means our results naturally cover the use of discrete choice models in dynamic settings, in which preference shocks are realized each period and expected utility is naturally evaluated before their realization.

Define a realization of taste shocks for household i as $\xi^i = (\xi_1^i, \dots, \xi_j^i, \dots, \xi_J^i)$. Associated with these J random variables are a cumulative density function $G(\xi^i; \theta)$ and a probability density function $g(\xi^i; \theta)$. At this point, we do not make any functional form assumptions on G . With that said, below we will use the canonical type 1 extreme value distribution as an example for illustrative purposes.

Household i of type θ derives the following utility, conditional on choosing good j :

$$u(c, q_j; j, \theta) + \xi_j^i. \quad (1)$$

Utility depends on quantities of consumption of the homogeneous good c and the differentiated good q_j . We assume that the utility function is well behaved in these arguments.

Our specification of utility also allows for properties of differentiated good j to matter through channels other than the direct consumption of the differentiated good. The obvious one is the additive taste shock that is specific to household i , discussed above. The second feature is that the utility function is separately indexed by j , representing the idea that the attributes of j may be more or less valuable relative to other j' products. In the case of products, different goods may have different quality or attributes. In the case of location choice, different locations may have different amenities.

Finally, we index utility by the type of the household θ . Different θ types could have different preferences over consumption of the commodities and value attributes of differentiated goods differently. For example, in the case of location, with θ representing family size, single house-

holds could have different valuations for, say, cities, than those of nuclear households.

Nested cases. This generality nests representations of canonical functional forms in discrete choice models. One example is the demand for differentiated goods as in Berry et al. (1995) or Nevo (2000). In these applications the outside good is separable, with unit demand for the differentiated good:

$$u(c, q_j; j, \theta) = \alpha_\theta c + \beta(\theta) \mathbf{X}_j.$$

Thus, our specification in (1) can entertain ideas emphasized in the demand estimation literature: heterogeneity in own-price elasticities of demand (α_θ) and heterogeneity in consumers' value of product j 's attributes or quality, \mathbf{X}_j , via the vector of parameters $\beta(\theta)$.⁶

A simpler example we use below is the following: no outside good, continuous choice in q_j , no quality differences, and common preferences across θ types:

$$u(c, q_j; j, \theta) = u(q_j).$$

When u is log and ξ^i is distributed according to a type 1 extreme value distribution, we obtain the case of Anderson et al. (1987), who establish equivalence between discrete choice demand and demand from a representative CES consumer.

Another useful example is

$$u(c, q_j; j, \theta) = u(c).$$

This is the natural case when j is a location, and thus j effects show up entirely through c via the budget constraint. This case corresponds to versions of Rosen-Roback models (Rosen, 1979; Roback, 1982) and more recently in Kline and Moretti (2014) or the quantitative spatial literature (Redding and Rossi-Hansberg 2017; Fajgelbaum and Gaubert 2020). Diamond (2016) is an example that merges an IO-like formulation (discussed above) in a spatial setting.

Endowments. Households of type θ are endowed with the commodities discussed above. All households of type θ have the same endowment. In a decentralized economy, where each differentiated good has price p_j , the value of this endowment is

$$W(\theta) := y_o(\theta) + \sum_j p_j y_j(\theta), \tag{2}$$

$y_j(\theta)$ is type- θ endowment of good j and $y_o(\theta)$ is the endowment of the outside good.

⁶Nakamura and Zerom (2010) uses a similar specification in a macro-price setting application.

2. The Standard (Incomplete Markets) Equilibrium

In this section, we describe the problems of the actors in the economy, the resource constraints, and then define an equilibrium. We call this equilibrium the “Standard (Incomplete Markets) Equilibrium.” It is standard in the sense that this is what virtually everyone computes. However, we add the incomplete markets part as a point of contrast to the complete markets equilibrium that we consider in Section 3

When characterizing the standard equilibrium, we pose the individual household’s problem in an ex-ante sense. That is, before taste shocks are realized, the household formulates a plan that maps realizations of ξ into a commodity choice and quantities. And these plans are chosen to maximize ex-ante expected utility. The typical approach formulates the problem in an ex-post sense. It starts with a realization of shocks and asks what households do. Then, given the ex-post decision rules, the researcher measures welfare by reconstructing average utility in the population. Within θ types, all households are ex-ante identical, so average utility in the population is the same as expected utility of individuals, given some decision rules.

With incomplete markets, these two approaches (ex-post and our ex-ante approach) are equivalent. The decision rules are the same, and maximized ex-ante utility corresponds to utility in the population when utility is maximized ex-post. The reason why we do everything from an ex-ante perspective is that it motivates the existence of ex-ante trades that can be made to increase utility among individuals and in the population.

Another important implication of our approach is that it maps into the wide application of dynamic discrete choice models. In these settings, the shocks are realized each period, and hence ex-ante expected utility is the obvious criteria for individual maximization.

The Household’s Problem. Household i of type θ has expected utility $V^i(\theta)$, given by

$$V^i(\theta) = \int_{\xi} \sum_j x_j^i(\xi, \theta) \left\{ u[c^i(\xi, \theta), q_j^i(\xi, \theta); j, \theta] + \xi_j^i \right\} g(\xi; \theta) d\xi. \quad (3)$$

We introduce the following notation: $x_j^i(\xi, \theta)$ is an indicator function that maps ξ into a one if j is chosen and zero otherwise, and $c^i(\xi, \theta)$ and $q_j^i(\xi, \theta)$ are functions that map ξ into the quantities consumed.

The inside summation in (3) says that for a given vector of taste shocks ξ , utility depends on whatever good is consumed, how much of it is consumed, plus the shock associated with the good chosen. The outside integral integrates over all possible realizations of ξ with density $g(\xi, \theta)$. This defines expected utility for household i of type θ .

The household chooses quantities and indicator functions for all possible realizations of the

taste shocks to maximize (3), subject to the household's budget constraint:

$$\max_{c_j^i(\boldsymbol{\xi}, \theta), q_j^i(\boldsymbol{\xi}, \theta), x_j^i(\boldsymbol{\xi}, \theta)} \int_{\boldsymbol{\xi}} \sum_j x_j^i(\boldsymbol{\xi}, \theta) \left\{ u[c^i(\boldsymbol{\xi}, \theta), q_j^i(\boldsymbol{\xi}, \theta); j, \theta] + \xi_j^i \right\} g(\boldsymbol{\xi}; \theta) d\boldsymbol{\xi}, \quad (4)$$

$$\text{subject to } [\lambda^i(\boldsymbol{\xi}, \theta)] : \sum_j x_j^i(\boldsymbol{\xi}, \theta) [c^i(\boldsymbol{\xi}, \theta) + p_j q_j^i(\boldsymbol{\xi}, \theta)] \leq W(\theta), \quad \text{for all } \boldsymbol{\xi}. \quad (5)$$

In (5), $\lambda^i(\boldsymbol{\xi}, \theta)$ are multipliers on household budget constraints, with one budget constraint for every realization of the vector of shocks. As we discuss more below, the key issue in this problem is the budget constraints. The constraint (5) must hold shock by shock, and thus a household's valuation of its resources will vary across shocks, with the relative scarcity of resources reflected in the multipliers $\lambda^i(\boldsymbol{\xi}, \theta)$. Not stated in this problem is the implicit constraint which makes this a discrete choice problem: $x_j^i(\boldsymbol{\xi}, \theta)$ can equal one for only one j .

Ex-ante perspective. We take the ex-ante perspective before shocks are realized for two reasons. First, because the objective function that arises from this perspective is consistent with the objective function maximized analytically or numerically in applied optimal policy exercises. Second, discrete choice models have proliferated in their use in dynamic settings. In these settings, shocks are explicitly treated as uninsurable.

Resource constraints. Finally there are the following resource constraints. Product demand cannot exceed product supply, requiring for all $j = 1, \dots, J$ that

$$\sum_{\theta} \int_0^{\mu(\theta)} y_j(\theta) di \geq \sum_{\theta} \int_0^{\mu(\theta)} \left\langle \int_{\boldsymbol{\xi}} x_j^i(\boldsymbol{\xi}, \theta) q_j^i(\boldsymbol{\xi}, \theta) g(\boldsymbol{\xi}; \theta) d\boldsymbol{\xi} \right\rangle di. \quad (6)$$

The right-hand side is aggregate demand, which integrates over household types θ , then households within each type i . By the law of large numbers, $\langle \cdot \rangle$ is total demand for good j by households of type θ . The left-hand side is aggregate supply, which is the endowments each i, θ household is endowed with. For the outside good, a similar condition must hold:

$$\sum_{\theta} \int_0^{\mu(\theta)} y_o(\theta) di \geq \sum_{\theta} \int_0^{\mu(\theta)} \left\langle \int_{\boldsymbol{\xi}} c^i(\boldsymbol{\xi}, \theta) g(\boldsymbol{\xi}, \theta) d\boldsymbol{\xi} \right\rangle di. \quad (7)$$

Definition 1 (The Standard (Incomplete Markets) Equilibrium.) *An equilibrium consists of an allocation for each i and θ type $c^i(\boldsymbol{\xi}, \theta)$, $q_j^i(\boldsymbol{\xi}, \theta)$, $x_j^i(\boldsymbol{\xi}, \theta)$ and prices p_j , such that*

- i. allocations $(c^i(\boldsymbol{\xi}, \theta), q_j^i(\boldsymbol{\xi}, \theta), \text{ and } x_j^i(\boldsymbol{\xi}, \theta))$ satisfy the household's problem in (4);*
- ii. resource constraints (6), (7) are satisfied.*

2.1. Properties of the Household's Problem

In this section, we characterize properties of the allocations that satisfy the household problem. Yes, this problem has been solved many times. But we do so in a different way which facilitates the rest of the analysis. Appendix A details our approach; below, we state the main results.

Consumption policy. First, consumption allocations must satisfy the properties

$$u_c [c(j, \theta), q_j(\theta); j, \theta] = \lambda_j(\theta) \quad \text{and} \quad u_q [c(j, \theta), q_j(\theta); j, \theta] = \lambda_j(\theta) p_j, \quad (8)$$

where u_c is the marginal utility of outside good consumption and u_q is the marginal utility of the differentiated good. Households equate marginal utility of consumption to marginal costs, measured by the multiplier on the budget constraint times the price.

Notice that we purposely dropped the indexing by the taste shock in several ways. This is a result, not a typo or an attempt to save on notation. First, $c(j, \theta)$ no longer depends upon the taste shock but now inherits dependence upon the choice of the differentiated good j . Similarly, $q_j(\theta)$ does not vary with the taste shock. The multipliers that were indexed by events ξ are now re-indexed by the choice, not the shock. All of these arguments follow from the observation that the taste shock does not affect marginal conditions and the taste shock affects the multiplier only through the choice j , not the shock per se. Finally, these same arguments — along with identical resources and preferences for all $i \in \theta$ — imply the household's identity does not matter.

Discrete choice rule. The second aspect of the solution is the discrete choice. Optimal $x_j^i(\xi, \theta)$ takes the form

$$x_j^i(\xi, \theta) = \begin{cases} 1, & \text{if } u [c(j, \theta), q_j(\theta); j, \theta] + \xi_j^i \geq \max_{j'} \left\{ u [c(j', \theta), q_{j'}(\theta); j', \theta] + \xi_{j'}^i \right\} \\ 0, & \text{otherwise.} \end{cases} \quad (9)$$

This is essentially the starting point from the literature. But we did not start here. As we discuss in Appendix A, we first form the Lagrangian from (4), then make a formal variational argument to obtain (9). Thus, our ex-ante approach delivers the same choice rule that would come from an approach that starts with realized shocks.

Discrete choice probability. We can integrate (9) across the taste shocks and arrive at standard discrete choice results. If the shocks are distributed type 1 extreme value with shape parameter η_θ , the probability a household of type θ chooses choice j is

$$\rho_j(\theta) = \exp \left(\frac{u [c(j, \theta), q_j(\theta); j, \theta]}{\eta_\theta} \right) / \sum_{j'} \exp \left(\frac{u [c(j', \theta), q_{j'}(\theta); j', \theta]}{\eta_\theta} \right). \quad (10)$$

This is exactly what one would expect if, for example, one were flipping through Train's (2009) textbook presentation. With a continuum of $i \in \theta$, the law of large numbers applies, and $\rho_j(\theta)$ is also the measure of type- θ households consuming j . Then type- θ demand in (6) becomes $\mu(\theta)\rho_j(\theta)q_j(\theta)$.

All of this is natural. The conditions in (8) equate ratios of marginal utility between the outside and differentiated good with relative prices. Condition (9) says that given these optimal consumption plans, do the best possible given the shocks that were realized. Risk associated with the event ξ does not appear to directly affect consumption. These are the arguments one might put forward to argue that nothing is wrong with this economy.

Except for one thing. The conditions in (8) do *not* deliver a "risk sharing"-like condition where the shadow value of resources (the multipliers) is equated *across* events. In allocations where risk is shared (see, e.g., Townsend 1994 or Backus and Smith 1993), there is a condition like

$$u_q[c(j, \theta), q_j(\theta); j, \theta]/p_j = \lambda(\theta) \quad \forall j. \quad (11)$$

Here, the j index is absent from the multiplier: across all states of nature (the taste shocks directly and the choices they induce) consumption is such that resources are valued equally. This is *not* occurring in the standard allocation when one inspects (8).

The problem in (8) is that the events ξ indirectly show up in the choice through the budget constraint. Resources are more valuable when ξ leads to a choice of a high price good (expensive city, or low paying job). In these states, the multiplier $\lambda_j(\theta)$ will be high. In contrast, in the event that ξ leads to the choice of a low price good (cheap city, or high paying job), the multiplier $\lambda_j(\theta)$ is low, resources are abundant, and marginal utility is low. A better arrangement would have a bit more resources available to the household that likes the high price commodity and a bit less available to the household that likes the low price commodity, up to the point that marginal utility across choices is equated. However, to paraphrase Rogerson (1988, p. 11), *making choices discrete creates a barrier to trade*.

What trading opportunities and market arrangements can achieve a condition like (11)? How is commodity choice affected? How do these market arrangements line up with Pareto efficient allocations? Do Pareto efficient allocations line up with social welfare maximizing allocations? These are the questions that we answer next.

3. Complete Markets Equilibrium and First Welfare Theorem

In this section we allow agents to trade contingent claims that pay out given the realizations of the taste shocks. This is our complete markets equilibrium. We argue this equilibrium is Pareto efficient and then show how it differs from the standard equilibrium. Put simply, if a researcher

assesses welfare in a discrete choice model, this is the benchmark allocation that maximizes it.

3.1. Insurance Contracts

Households purchase $a^i(\boldsymbol{\xi}, \theta)$ many claims at the state price $\varphi(\boldsymbol{\xi}, \theta)$. These claims pay out $a^i(\boldsymbol{\xi}, \theta)$ upon the realization of the taste shock vector $\boldsymbol{\xi}$ and zero otherwise. There are no restrictions on short-selling: $a^i(\boldsymbol{\xi}, \theta) \in (-\infty, \infty)$.

Competitive firms provide these insurance contracts to households. Insurance firms determine the state prices; otherwise, they are simply veils. These firms operate by selling insurance contracts $a^i(\boldsymbol{\xi}, \theta)$ minus the cost. The cost is the probability that event $\boldsymbol{\xi}$ occurs times the quantity of insurance provided. The PDF $g(\boldsymbol{\xi}; \theta)$ determines the probability of the event $\boldsymbol{\xi}$ occurring. Competitive pricing in the market for each contract yields zero profits and implies that the household faces actuarially fair prices:

$$\varphi(\boldsymbol{\xi}, \theta) = g(\boldsymbol{\xi}; \theta). \quad (12)$$

At this point, we emphasize how rich this contract structure is. Ex-ante, households buy and sell claims against every possible realization. Ex-post, paying claims requires identifying households $\boldsymbol{\xi}$ -by- $\boldsymbol{\xi}$. Below we show that the complete markets equilibrium can be supported by a much simpler set of securities.

3.2. The Household Problem with Complete Markets

The household's problem with complete markets is

$$\max_{a^i(\boldsymbol{\xi}, \theta), c^i(\boldsymbol{\xi}, \theta), q_j^i(\boldsymbol{\xi}, \theta), x_j^i(\boldsymbol{\xi}, \theta)} \int_{\boldsymbol{\xi}} \sum_j x_j^i(\boldsymbol{\xi}, \theta) \left\{ u[c^i(\boldsymbol{\xi}, \theta), q_j^i(\boldsymbol{\xi}, \theta); j, \theta] + \xi_j^i \right\} g(\boldsymbol{\xi}; \theta) d\boldsymbol{\xi}, \quad (13)$$

$$\text{subject to } \sum_j x_j^i(\boldsymbol{\xi}, \theta) [c^i(\boldsymbol{\xi}, \theta) + p_j q_j^i(\boldsymbol{\xi}, \theta)] \leq W(\theta) + a^i(\boldsymbol{\xi}, \theta) \quad \forall \boldsymbol{\xi}, \quad (14)$$

$$\int_{\boldsymbol{\xi}} \varphi(\boldsymbol{\xi}, \theta) a^i(\boldsymbol{\xi}, \theta) d\boldsymbol{\xi} = 0. \quad (15)$$

A household chooses contingent claims, consumption quantities, and commodity choices for every possible state, subject to two constraints. The first constraint is that consumption equals labor income and amount of insurance purchased, state by state. The second constraint says the household's net asset position must be zero.

To illustrate the key property of this problem, substitute all budget constraints (14) into the

constraint for assets (15). The resulting problem is

$$\max_{a^i(\boldsymbol{\xi}, \theta), c^i(\boldsymbol{\xi}, \theta), q_j^i(\boldsymbol{\xi}, \theta), x_j^i(\boldsymbol{\xi}, \theta)} \int_{\boldsymbol{\xi}} \sum_j x_j^i(\boldsymbol{\xi}, \theta) \left\{ u[c^i(\boldsymbol{\xi}, \theta), q_j^i(\boldsymbol{\xi}, \theta); j, \theta] + \xi_j^i \right\} g(\boldsymbol{\xi}; \theta) d\boldsymbol{\xi}, \quad (16)$$

$$\text{subject to } [\lambda^i(\theta)]: \int_{\boldsymbol{\xi}} \varphi(\boldsymbol{\xi}, \theta) \left\{ W(\theta) - \sum_j x_j^i(\boldsymbol{\xi}, \theta) [c^i(\boldsymbol{\xi}, \theta) + p_j q_j^i(\boldsymbol{\xi}, \theta)] \right\} d\boldsymbol{\xi} = 0. \quad (17)$$

The constraint (17) is the distinguishing feature between the complete markets problem and the problem in (4). Here, (17) allows the household to consolidate all possible outcomes of $\boldsymbol{\xi}$ before their realization.⁷ In contrast, in the standard equilibrium, households face different constraints for each outcome $\boldsymbol{\xi}$ and are unable to ex-ante trade across these different outcomes.

Definition 2 defines the complete markets equilibrium.

Definition 2 (The Complete Markets Equilibrium.) *A complete markets equilibrium consists of allocations $c^i(\boldsymbol{\xi}, \theta)$, $q_j^i(\boldsymbol{\xi}, \theta)$, $a^i(\boldsymbol{\xi}, \theta)$, $x_j^i(\boldsymbol{\xi}, \theta)$ for each i and type θ ; goods prices p_j and state prices $\varphi(\boldsymbol{\xi}, \theta)$ such that*

- i. the consumption allocations ($c^i(\boldsymbol{\xi}, \theta)$, $q_j^i(\boldsymbol{\xi}, \theta)$, and $x_j^i(\boldsymbol{\xi}, \theta)$) and asset positions $a^i(\boldsymbol{\xi}, \theta)$ satisfy the household's problem in (16);*
- ii. state prices satisfy (12);*
- iii. goods and asset markets clear.*

3.3. Properties of the Household's Problem with Complete Markets

In this section, we characterize properties of the allocations that satisfy the household problem with complete markets. Appendix A details our approach; below, we state the main results.

Consumption policy. Consumption of the outside and differentiated good must satisfy:

$$u_c[c(j, \theta), q_j(\theta); j, \theta] = \lambda(\theta) \quad \text{and} \quad u_q[c(j, \theta), q_j(\theta); j, \theta] = \lambda(\theta)p_j. \quad (18)$$

Again, re-indexing is a result, not a shortcut. Consumption $c^i(\boldsymbol{\xi}, \theta)$, $q_j^i(\boldsymbol{\xi}, \theta)$ depends on θ -dependent preferences and endowments (we can drop i) and the good chosen under ξ^i , but not the taste shock itself.

The most important thing to notice is that the multiplier $\lambda(\theta)$ does not depend upon the real-

⁷As an analog to dynamic models, (17) can be interpreted as a household's "lifetime" budget constraint, and thus the ability to trade insurance allows all these different outcomes to be combined in a "date 0" way.

ization of the taste shock nor the choice. With complete markets there is one constraint (not shock-specific constraints as in the standard setting), as the household is now able to shift resources across states of nature.

Because these first order conditions are equated with the same multiplier, we now have a standard risk sharing-type result. Compare a θ -type household consuming j with one consuming alternative good j' :

$$\frac{u_c[c(j, \theta), q_j(\theta); j, \theta]}{u_c[c(j', \theta), q_{j'}(\theta); j', \theta]} = 1 \quad \text{and} \quad \frac{u_q[c(j, \theta), q_j(\theta); j, \theta]}{u_q[c(j', \theta), q_{j'}(\theta); j', \theta]} = \frac{p_j}{p_{j'}}. \quad (19)$$

The ratios of marginal utility across goods should equal their relative price. Marginal utility of the homogeneous outside good is the same *independent* of the identity of good j chosen. Marginal utility of the differentiated good differs only to the extent that relative prices are different.

The conditions in (19) also imply that households consume *as if* the consumption choice of the differentiated good were continuous. That is, marginal rates of substitution across the discrete commodities equal relative prices. Even though households choose only one commodity to consume, households, ex-ante choose consumption plans in the same way that would arise in a model where all goods are consumed simultaneously. In Section 6 we will show that this implies that standard formulas for welfare effects of price changes used in applied work only hold in the case of complete markets.

Discrete choice rule. When households can trade all these contingent claims, which good do they actually consume? Interestingly, the commodity choice rule takes a unique form with

$$x_j^i(\boldsymbol{\xi}, \theta) = \begin{cases} 1, & \text{if } u[c(j, \theta), q_j(\theta); j, \theta] + \xi_j^i - \lambda(\theta) [c(j, \theta) + p_j q_j(\theta)] \geq \\ & \max_{j'} \left\{ u[c(j', \theta), q_{j'}(\theta); j', \theta] + \xi_{j'}^i - \lambda(\theta) [c(j', \theta) + p_{j'} q_{j'}(\theta)] \right\} \\ 0, & \text{otherwise.} \end{cases} \quad (20)$$

The key novelty is the subtracted $\lambda(\theta)[\cdot]$ term, which is *not* in the rule characterizing the incomplete markets allocation (equation 9). With incomplete markets, the effects of choosing j are siloed into the good- j budget constraint. With complete markets, choosing j imposes a cost across all other states of nature, by removing resources from the aggregated budget constraint (17). This cost is the shadow value of resources $\lambda(\theta)$, times the additional burden $c(j, \theta) + p_j q_j(\theta)$. Efficient choice probabilities now account for this cost.

What are the claims that are actually traded? Here we find that a only a subset of contingent

claims are traded in equilibrium. Let $\Xi_j(\theta)$ be the set of all ξ for which a household of type θ chooses j . If $\xi \in \Xi_j(\theta)$, then the household's asset position is

$$a^i(\xi, \theta) = W(\theta) - [c^i(\xi, \theta) + p_j q_j^i(\xi, \theta)] = W(\theta) - [c(j, \theta) + p_j q_j(\theta)]. \quad (21)$$

Since the right-hand side is independent of ξ directly, then so is $a^i(\xi, \theta)$, which we denote $a(j, \theta)$.

This tells us that the only claims that are actually traded are those that depend upon the choice j , not the taste shocks per se. In other words, the same allocation would have been obtained with a set of securities that paid off conditional on choices. We term these contingent claims “*Arrow vouchers*,” as they pay off in states of nature that coincide with the choice of a commodity — hence our voucher terminology.

We find this result striking. First, it suggests that policies that give subsidies based on realized choices may go in the direction of completing markets and hence improving welfare. Take, for example, a simple policy of a subsidy to individuals who choose to live in San Francisco. The complete markets allocation says that a component of this is welfare improving, as it provides insurance against a high preference for living in an expensive city. Second, this result shows that these trades really are about insurance conditional on type θ , not redistribution across type θ 's. Observe that trade could be restricted to be *within* household types (or just assume that there was only one θ -type in the economy). In this case, identical households will want to trade against the potential ex-post different choices that the taste shocks induce.

Returning to the choice rule in (20), we can connect the Arrow vouchers with choices. Substituting $a(j, \theta)$ equation (21) into the choice rule:

$$x_j^i(\xi, \theta) = \begin{cases} 1, & \text{if } u[c(j, \theta), q_j(\theta); j, \theta] + \xi_j^i - \lambda(\theta)a(j, \theta) \geq \\ & \max_{j'} \left\{ u[c(j', \theta), q_{j'}(\theta); j', \theta] + \xi_{j'}^i - \lambda(\theta)a(j', \theta) \right\} \\ 0, & \text{otherwise.} \end{cases} \quad (22)$$

The extra term, which penalizes the choice of j relative to the incomplete markets choice rule, is the asset position required to optimally smooth consumption across choices $a(j, \theta)$, evaluated at its marginal value of resources $\lambda(\theta)$.

Discrete choice probability. If we impose the Type I extreme value distributional assumption and integrate (22), we obtain the *efficient choice probabilities*:

$$\rho_j(\theta) = \exp\left(\frac{u[c(j, \theta), q_j(\theta), j, \theta] - \lambda(\theta)a(j, \theta)}{\eta_\theta}\right) / \sum_{j'} \exp\left(\frac{u[c(j', \theta), q_{j'}(\theta), j', \theta] - \lambda(\theta)a(j', \theta)}{\eta_\theta}\right). \quad (23)$$

Equation (23) has two meanings. One is the probability that a θ type agent chooses j . With complete markets, the choice probabilities are also the state price for the Arrow voucher associated with choice j . In other words, the choice probabilities are also the price of purchasing insurance against events that lead to j being chosen.

At this point we have demonstrated several results. First, the complete markets allocation leads to a consolidated budget constraint across states of nature, and thus optimal allocations have a risk sharing condition with marginal utility *across* goods equaling their relative price. Second, complete markets influence choice probabilities by forcing the household to internalize how the cost of one choice versus others varies across states of nature.⁸ Finally, the set of contingent claims needed to implement complete markets needs to span only the choices, not all taste shocks. Proposition 1 summarizes these results.

Proposition 1 (Complete Markets Allocations) *The following conditions characterize the complete markets allocations:*

1. *Consumption allocations must satisfy*

$$u_c[c(\theta, j), q_j(\theta); j, \theta] = \lambda(\theta) \quad \text{and} \quad u_{q_j}[c(\theta, j), q_j(\theta); j, \theta] = \lambda(\theta)p_j.$$

2. *The commodity choice rule is*

$$x_j^i(\xi, \theta) = \begin{cases} 1, & \text{if } u[c(j, \theta), q_j(\theta); j, \theta] + \xi_j^i - \lambda(\theta)a(j, \theta) \geq \\ & \max_{j'} \left\{ u[c(j', \theta), q_{j'}(\theta); j', \theta] + \xi_{j'}^i - \lambda(\theta)a(j', \theta) \right\} \\ 0, & \text{otherwise.} \end{cases}$$

3. *“Arrow Vouchers.” Asset positions are given by*

$$a(j, \theta) = W(\theta) - [c(j, \theta) + p_j q_j(\theta)],$$

which are contingent only on the choice j , not the taste shock ξ . The state prices for the Arrow

⁸Note that this second observation implies that the government cannot implement the efficient allocation simply by offering individuals a transfer conditional on their choice. That is, if we took the Arrow security pay-offs and relabelled them lump-sum transfers—which by definition will also sum to zero—then individuals would not choose the same allocation. Why is this? With lump-sum transfers, there is an aggregate budget constraint of the government that is not internalized by individuals. This budget constraint links how discrete choices impact the resources available across payments. Under complete markets that budget constraint appears in the individual’s problem and the effect of choices on state prices is understood by the household.

Vouchers are the choice probabilities, which with the type 1 extreme value assumption are

$$\exp\left(\frac{u[c(j, \theta), q_j(\theta), j, \theta] - \lambda(\theta)a(j, \theta)}{\eta_\theta}\right) / \sum_{j'} \exp\left(\frac{u[c(j', \theta), q_{j'}(\theta), j', \theta] - \lambda(\theta)a(j', \theta)}{\eta_\theta}\right).$$

3.4. The Complete Markets Equilibrium Is Pareto Efficient

It turns out that the complete markets equilibrium is a Pareto efficient allocation, and thus we have a first fundamental theorem of welfare economics in this environment. Proposition 2 states the result.

Proposition 2 (The First Welfare Theorem.) *The complete markets equilibrium is a Pareto efficient allocation.*

The argument behind this statement is standard, relying on the idea that there can't be a feasible allocation that is preferred. A key step in the argument is the equivalence of the household budget constraint under competitive, actuarially fair prices for Arrow securities, and the resource constraint. Appendix A details the entire argument.

The importance of this result for applied work is not that we think that opening complete markets is a viable policy. Instead it firmly benchmarks an efficient allocation, and its properties. Hence the effectiveness of other viable policies may be interpreted in terms of how they move the economy toward or away from the properties of the complete markets benchmark. We put this into practice in Section 5.

4. Pareto Efficient Allocations and Second Welfare Theorem

We have shown that the complete markets equilibrium is a Pareto efficient allocation. We now characterize *all* Pareto efficient allocations. This provides the foundation for the second welfare theorem: any Pareto efficient allocation can be decentralized as a complete markets equilibrium allocation with the appropriate ex-ante transfers. Appendix C provides all derivations.

We set up the Pareto problem as follows. Fix one household with name (i, θ) . The planner chooses allocations of consumption and commodity choice rules to maximize (i, θ) utility subject to (i) resource constraints, and (ii) making all other (k, θ') households no worse than some given level. Below are the details.

The objective function of the planner is to maximize expected utility for (i, θ) household:

$$\max V^i(\theta) = \int_{\xi} \sum_j x_j^i(\xi, \theta) \left\{ u[c^i(\xi, \theta), q_j^i(\xi, \theta); j, \theta] + \xi_j^i \right\} g(\xi; \theta) d\xi, \quad (24)$$

where the choice variables are the household's consumption and commodity choices and consumption and commodity choices for *all* other agents in the economy which are indexed by (k, θ') . The first constraints are the resource constraints:

$$[\Lambda_o]: \sum_{\theta'} \int_0^{\mu(\theta')} y_o^k(\theta') dk \geq \sum_{\theta'} \int_{\xi} \int_k x_j^k(\xi, \theta') c_j^k(\xi, \theta') dk g(\xi, \theta') d\xi, \quad (25)$$

$$[\Lambda_j]: \sum_{\theta'} \int_0^{\mu(\theta')} y_j^k(\theta') dk \geq \sum_{\theta'} \int_{\xi} \int_k x_j^k(\xi, \theta') q_j^k(\xi, \theta') dk g(\xi, \theta') d\xi \quad \forall j, \quad (26)$$

where the new multipliers Λ_j and Λ_o capture the shadow value of each good. The next constraint is the Pareto constraint:

$$[\Upsilon^k(\theta')]: V^k(\theta') \leq \int_{\xi} \sum_j x_j^k(\xi, \theta') \left\{ u[c^k(\xi, \theta'), q_j^k(\xi, \theta'); j, \theta'] + \xi_j^k \right\} g(\xi, \theta') d\xi \quad \forall k, \theta' \neq i, \theta. \quad (27)$$

Any allocation must deliver utility level $V^k(\theta')$ or better for every (k, θ') household. Associated with this constraint is the Lagrange multiplier $\Upsilon^k(\theta')$. Any Pareto efficient allocation satisfies this problem.

4.1. Properties of Pareto Efficient Allocations

The characterization of the Pareto problem takes the following form. For household (i, θ) , the marginal conditions are as follows (where, again, collapsing notation is a result):

$$u_c[c^i(j, \theta), q_j^i(\theta); j, \theta] = \Lambda_o \quad \text{and} \quad u_{q_j}[c^i(j, \theta), q_j^i(\theta); j, \theta] = \Lambda_j. \quad (28)$$

Then for all other (k, θ') households, they are

$$\Upsilon^k(\theta') u_c[c^k(j, \theta'), q_j^k(\theta'); j, \theta'] = \Lambda_o \quad \text{and} \quad \Upsilon^k(\theta') u_{q_j}[c^k(j, \theta'), q_j^k(\theta'); j, \theta] = \Lambda_j. \quad (29)$$

These conditions mimic what is occurring in the complete markets equilibrium. Adjusted, marginal utility equals the shadow cost of the commodity, which is the multipliers on the resource constraint. The multipliers $\Upsilon^k(\theta')$ on the Pareto constraints reflect the social value of each individual to the planner.

The next step is the optimal choice rule describing which commodity is to be consumed. The

choice rule is given by

$$x_j^k(\boldsymbol{\xi}, \theta') = \begin{cases} 1, & \text{if } u[c^k(j, \theta'), q_j^k(\theta'); j, \theta'] + \xi_j^k - \frac{1}{\Upsilon^k(\theta')} \left[\Lambda_o c^k(j, \theta') + \Lambda_j q_j^k(\theta') \right] \geq \\ & \max_{j'} \left\{ u[c^k(j', \theta'), q_{j'}^k(\theta'); j', \theta'] + \xi_{j'}^k - \frac{1}{\Upsilon^k(\theta')} \left[\Lambda_o c^k(j', \theta') + \Lambda_{j'} q_{j'}^k(\theta') \right] \right\} \\ 0, & \text{otherwise.} \end{cases} \quad (30)$$

This holds for all (k, θ') households and for (i, θ) by setting the Υ term to one.

Again, this expression looks similar to the complete markets equilibrium. Consider the additional term $\frac{1}{\Upsilon^k(\theta')}[\cdot]$. Rather than the private cost of insurance into choice j , it is the *social cost* of providing utility to the (k, θ') household. This is the shadow value of outside and differentiated good consumption, evaluated at the shadow value of those resources, Λ_o and Λ_j . This is adjusted by how socially valuable household (k, θ') is.

Proposition 3 summarizes.

Proposition 3 (Pareto Efficient Allocations) *Given utility levels $V^k(\theta')$ for all $k, \theta' \neq i, \theta$, a Pareto efficient allocation is consumption allocations and commodity choice rules $c_j^i(\boldsymbol{\xi}, \theta)$, $q_j^i(\boldsymbol{\xi}, \theta)$, $x_j^i(\boldsymbol{\xi}, \theta)$ for household (i, θ) and all other (k, θ') households, that solve the problem (24) subject to resource constraints (25, 26) and the Pareto constraint in (27).*

The following conditions characterize Pareto efficient allocations:

1. For agent (i, θ) , consumption allocations must satisfy

$$u_c[c^i(j, \theta), q_j^i(\theta); j, \theta] = \Lambda_o \quad \text{and} \quad u_{q_j}[c^i(j, \theta), q_j^i(\theta); j, \theta] = \Lambda_j.$$

2. For agent (k', θ') , consumption allocations must satisfy

$$\Upsilon^k(\theta') u_c[c^k(j, \theta'), q_j^k(\theta'); j, \theta'] = \Lambda_o \quad \text{and} \quad \Upsilon^k(\theta') u_{q_j}[c^k(j, \theta'), q_j^k(\theta'); j, \theta'] = \Lambda_j.$$

3. The commodity choice rule is

$$x_j^k(\boldsymbol{\xi}, \theta') = \begin{cases} 1, & \text{if } u[c^k(j, \theta'), q_j^k(\theta'); j, \theta'] + \xi_j^k - \frac{1}{\Upsilon^k(\theta')} \left[\Lambda_o c^k(j, \theta') + \Lambda_j q_j^k(\theta') \right] \geq \\ & \max_{j'} \left\{ u[c^k(j', \theta'), q_{j'}^k(\theta'); j', \theta'] + \xi_{j'}^k - \frac{1}{\Upsilon^k(\theta')} \left[\Lambda_o c^k(j', \theta') + \Lambda_{j'} q_{j'}^k(\theta') \right] \right\} \\ 0, & \text{otherwise,} \end{cases}$$

which holds for all (k, θ') households. For the (i, θ) household, set the Υ term to one.

4.2. Second Welfare Theorem

We now have a basis for the Second Welfare Theorem.

Proposition 4 (The Second Welfare Theorem.) *Any Pareto efficient allocations can be decentralized as a complete markets equilibrium.*

This is immediate from comparing the content of Proposition 1 with complete markets to Proposition 3. Consumption allocations and choice rules are identical if $\Lambda_o/\Upsilon^k(\theta') = \lambda(\theta')$ and $\Lambda_j/\Upsilon^k(\theta') = \lambda(\theta')p_j$. To align the complete markets allocation to the Pareto efficient allocation, then one simply needs to find the appropriate ex-ante reallocation of resources to attain the correct multipliers on complete markets budget constraints, $\lambda(\theta')$.

4.3. Social Welfare Maximization

We have worked exclusively with the Pareto problem, but it should not be surprising that a planner maximizing a straightforward social welfare function attains the same allocation. Proposition 5 summarizes this result. Appendix D provides all derivations and details.

Proposition 5 (Social Welfare Maximizing Allocations) *Let $\psi^i(\theta)$ be a vector of social welfare weights. Define the **Social Welfare Function** as*

$$\mathcal{W}_\psi = \sum_\theta \int_i \psi^i(\theta) \int_\xi \sum_j x_j^i(\boldsymbol{\xi}, \theta) \left\{ u[c^i(\boldsymbol{\xi}, \theta), q_j^i(\boldsymbol{\xi}, \theta); j, \theta] + \xi_j^i \right\} g(\boldsymbol{\xi}, \theta) d\boldsymbol{\xi} di. \quad (31)$$

*Then a **social welfare maximizing allocation** is consumption allocations and commodity choice rules and $c_j^i(\boldsymbol{\xi}, \theta)$, $q_j^i(\boldsymbol{\xi}, \theta)$, $x_j^i(\boldsymbol{\xi}, \theta)$ for all i, θ to maximize \mathcal{W}_ψ subject to resource constraints. The following conditions characterize the allocation:*

1. For all i, θ , consumption allocations must satisfy:

$$\psi^i(\theta) u_c[c^i(\theta), q_j^i(\theta); j, \theta] = \Lambda_o \quad \text{and} \quad \psi^i(\theta) u_{q_j}[c^i(\theta), q_j^i(\theta); j, \theta] = \Lambda_j. \quad (32)$$

2. For all i, θ the commodity choice rule is

$$x_j^i(\boldsymbol{\xi}, \theta) = \begin{cases} 1, & \text{if } u[c^i(j, \theta), q_j^i(\theta); j, \theta] + \xi_j^i - \frac{1}{\psi^i(\theta)} \left[\Lambda_o c^i(j, \theta) + \Lambda_j q_j^i(\theta) \right] \geq \\ & \max_{j'} \left\{ u[c^i(j', \theta), q_{j'}^i(\theta); j', \theta] + \xi_{j'}^i - \frac{1}{\psi^i(\theta)} \left[\Lambda_o c^i(j', \theta) + \Lambda_{j'} q_{j'}^i(\theta) \right] \right\} \\ 0, & \text{otherwise.} \end{cases} \quad (33)$$

This allocation is a Pareto efficient allocation and coincides with a complete markets allocation under some ex-ante transfers.

The last statement in Proposition 5 follows from a comparison with the results in Proposition 3. There is a clear mapping from $\psi^i(\theta)$ to the multipliers on the utility constraints in the Pareto problem. And a similar comparison between the results in Proposition 1 and the conditions from the results above lead to the conclusion that one needs only to find the correct transfers to align the complete markets allocation and the one that solves the social planning problem above.

This is an important result for applied work, as it shows that the standard objective function researchers use to evaluate welfare across counterfactual policy experiments in discrete choice economies, is maximized in a decentralized competitive equilibrium with complete markets, and that the associated allocations are Pareto efficient. This objective function is usually expressed in applied work as the “*expectation of the max*”.⁹ The implication of this result is that unless the allocations satisfy (1) and (2) in Proposition 5, they are *not maxima* of the objective function for which welfare is being evaluated.

Combined, these results allow researchers to understand social welfare improving policy under incomplete markets in terms of the degree to which it implements a complete markets allocation. In an applied context, when a researcher computes optimal policy in a discrete choice economy under the standard social welfare function, policy that goes toward “completing markets” will be welfare improving. Next section, we show how this can be used to understand how subsidies to production or migration can improve economic efficiency on this margin.

⁹See, for example, Busso et al. (2013, p. 902, paraphrased): “Denote the workers’ social welfare as $V = E_\xi [\max_j \{u_j^i + \xi_j^i\}]$, where the expectation is defined over the ξ_j^i terms.”

5. Optimal policy in environments with externalities

We now show that the above results are important for applied work. Theoretically, we have shown that discrete choice economies feature a generic inefficiency due to market incompleteness. Discrete choice economies are often used in applied work as the building block of supply and demand analysis in environments where (a) there are other inefficiencies—such as spillovers, externalities or imperfect product market competition—, (b) optimal policy is studied. Hence, optimal policy designed to address a particular inefficiency can not be separated from its role in addressing market incompleteness. We show how, using our results above, we can compare optimal policy under complete and incomplete markets economies in order to separate out these effects. We show that doing so changes the view of the magnitude of the optimal policy and associated welfare gains. We achieve this by studying two canonical examples in which discrete choice models and optimal policy have been studied: industrial policy and spatial policy. The key lesson is that policy that one may think of as targeting spillovers primarily delivers welfare gains from resolving market incompleteness, and if designed just to target spillovers optimal policy is substantially different with much smaller welfare gains.

5.1. Goods and Industrial Policy

The first case we study is optimal production subsidies in an environment with production spillovers. As in the following example, this shows that our results so far extend to a production economy.

Environment. The environment is a strict specialization of the general economy that we have studied thus far: (i) two goods, (ii) no outside good, (iii) no differences in ‘characteristics’ of goods, such that utility is only obtained from the quantity consumed of the differentiated good and the idiosyncratic preference: $u(q_j^i) + \xi_j^i$, (iv) Type-I distributed preferences with dispersion η , (v) no permanent type differences.

As opposed to an endowment economy, we assume that each individual inelastically supplies one unit of labor at equilibrium wage W , and production uses only labor. The production technology for Good 1 is $y_1 = z_1 n_1$. The production technology for Good 2 is $y_2 = A_2 z_2 n_2$, where $A_2 = y_2^\phi$ is a spillover that is not internalized by the firm in the competitive equilibrium. We assume that product markets are competitive.¹⁰

We assume $z_2 < z_1$, and hence the externality is associated with a sector that is currently less productive and hence producing higher priced goods but has some knowledge spillovers. A useful example would be electric versus combustion engine cars.

¹⁰One way to think about this is that each individual has preferences for goods produced by one of two sectors $j \in \{1, 2\}$. Within each sector a unit continuum of identical firms are competitive. Let Y_2 be sectoral output, then $y_2 = A(Y_2) z_2 n_2$. Firms in sector 2 choose $\{y_2, n_2\}$, taking Y_2 as given. In equilibrium $y_2 = Y_2$.

Incomplete markets. With incomplete markets, demand $(\rho_1, \rho_2, q_1, q_2)$ is determined by the standard choice probabilities and choice-dependent budget constraints:

$$\rho_1 = \frac{e^{u(q_1)/\eta}}{e^{u(q_1)/\eta} + e^{u(q_2)/\eta}} \quad , \quad \rho_2 = 1 - \rho_1 \quad , \quad q_1 = \frac{W}{P_1} \quad , \quad q_2 = \frac{W}{P_2}$$

Complete markets. We can now apply our results from the previous section, to immediately show that under complete markets demand is determined by (i) the efficient choice probabilities, (ii) an expression that equates the marginal rate of substitution between goods to relative prices, and (iii) the consolidated budget constraint:

$$\rho_1 = \frac{e^{[u(q_1) - u'(q_1)q_1]/\eta}}{e^{[u(q_1) - u'(q_1)q_1]/\eta} + e^{[u(q_2) - u'(q_2)q_2]/\eta}} \quad , \quad \rho_2 = 1 - \rho_1 \quad , \quad \frac{u'(q_1)}{u'(q_2)} = \frac{P_1}{P_2} \quad , \quad \rho_1 P_1 q_1 + \rho_2 P_2 q_2 = W.$$

Policy - Subsidies. In applied cases, a researcher may be interested in optimizing a reduced form tax or subsidy, subject to budget neutrality. Here, we allow a policy maker to choose budget neutral labor taxes and subsidies in each sector. Obviously, the policy maker would like to subsidize labor in sector 2 and tax labor in sector 1. That is, the firms' problems are

$$\Pi_1 = \max_{n_1} P_1 z_1 n_1 - W(1 + \tau_1)n_1 \quad , \quad \Pi_2 = \max_{n_2} P_2 A_2 z_2 n_2 - W(1 - \tau_2)n_2. \quad (34)$$

The requirement of budget neutrality implies that $\tau_1 n_1 = \tau_2 n_2$.

Firms. In a complete markets or incomplete markets economy, firm optimality is the same:

$$P_1 = W(1 + \tau_1) \quad , \quad P_2 = \frac{W(1 - \tau_2)}{A_2 z_2}.$$

A positive subsidy τ_2 directly lowers marginal cost, reducing P_2 . In equilibrium this increases demand for Good 2, and its production, which further reduces P_2 via a higher spillover A_2 .

Comparing welfare. We compute the subsidy τ_2^* that maximizes expected utility under complete and incomplete markets. We do so for a numerical example with $u(c) = c^{1-\sigma}/(1-\sigma)$, $\eta = 1/3$, $z_1 = 1$, $z_2 = 0.80$, $\phi = 0.10$ with baseline value of $\sigma = 3$. Given our previous results, Proposition 5 implies that we can evaluate optimal policy under a social welfare function, knowing that the maximum of the social welfare function is achieved at a Pareto efficient allocation.

Figure 1A plots the change in welfare for different values of the subsidy τ_2 under complete markets (red, dashed) and incomplete markets (blue, solid). The key outcome of our previous results is that we now understand that under complete markets, τ_2 is only addressing the externality. The optimal subsidy is 3.3% and the welfare gain is small, consistent with the gain in welfare from increasing TFP in of both goods by 0.55 percent. Under incomplete markets, however, we find that τ_2^* nearly doubles and the welfare gain increase by an order of magnitude, comparable to a 5.01 percent increase in TFP. Why is this? From what we have learned,

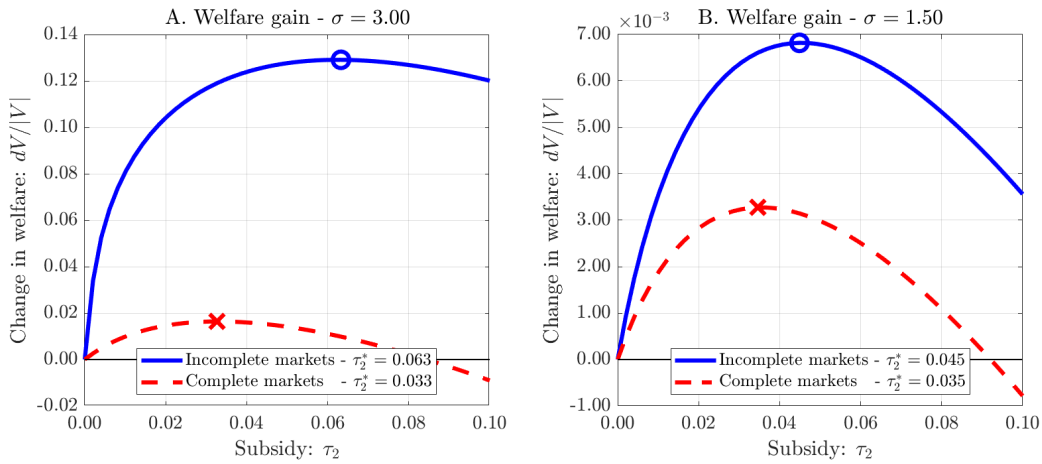


Figure 1: Welfare gains from optimal industrial policy

we know that the optimal τ_2^* under incomplete markets will be performing the additional duty of trying to equate marginal utilities across choices. Raising τ_2 , reduces P_2 , which reduces price dispersion across goods. Hence, across potential choices, individuals' consumption and hence marginal utilities become more similar, which is a shift toward the properties of the efficient allocation.

Implications for applied research. Our results in the previous sections have shown that in addition to the externality that is at the center of the policy question—the production spillover—there is also the inefficiency due to incomplete markets. When markets are incomplete, it is impossible for a researcher to choose τ_2^* to *only* address the externality. In this example, a researcher studying the natural setting—i.e. incomplete markets—may mistakenly conclude that a high subsidy is driven by addressing spillovers, while it is actually driven by redistribution via prices.

Constructively, this example provides two possible recipes for applied work. First, a researcher may study the *positive, empirical* properties of their model under incomplete markets, but when it comes to *normative, welfare* properties of different policies, compute optimal policy under complete markets in order to separate welfare gains. Let's consider the example of industrial policy around Electric (Good 2) vs Combustion (Good 1) cars. The optimal policy under complete markets delivers a subsidy that directly addresses the production spillovers associated with the electric car industry. The optimal policy under incomplete markets mainly distorts prices in order to deliver a subsidy to consumers that happen to like electric cars. We think that understanding how quantitatively different these policy rationales are is important, and that complete markets is a way to do it.

Computing welfare gains under complete markets has the additional benefit of robustness to key parameter values, as opposed to welfare gains under incomplete markets which are highly sensitive to key parameter values. Panel B plots the same exercise under $\sigma = 1.50$. Under

complete markets the welfare gains of the optimal policy are similar for $\sigma = 3.00$ (0.57 percent increase in TFP) and $\sigma = 1.50$ (0.55 percent increase in TFP). Under incomplete markets the welfare gains of the optimal policy are nearly five times larger under $\sigma = 3.00$ (5.01 percent) relative to $\sigma = 1.50$ (1.10 percent). Under incomplete markets, the large amount of redistribution under the optimal policy makes welfare gains sensitive to differences in σ . Under complete markets, the optimal policy focuses on the resolving the spillover which raises consumption for all workers.

Second, as suggested by Panel B, one may make assumptions on preferences that remove the market incompleteness inefficiency. As we show in Section 7, if preferences are logarithmic and there are no additional inefficiencies, then the incomplete markets allocation is efficient. Panel B shows that as we approach log, the optimal subsidy under incomplete and complete markets coincide as the inefficiency due to market incompleteness vanishes, leaving only the inefficiency due to externalities. Thus, if the preference specification is not important for positive implications, then assuming log preferences is another way to isolate the intervention of interest.

5.2. Migration and Spatial Policy

We now consider migration and optimal spatial policies with production spillovers that are location specific. In this case labor is mobile, while goods are not, which substantively differs from the previous case in which goods were traded but in a single location. We also study a policy which is a subsidy to workers instead of firms.

Environment. We generalize the economy we have studied so far to a spatial economy: (i) there are two locations $j \in \{1, 2\}$, (ii) a single tradeable good that sells at price P , (iii) utility is only obtained from the quantity consumed of the traded good and idiosyncratic preferences are over locations: $u(q^i) + \xi_j^i$, (iv) Type-I distributed preferences with dispersion η , (v) no permanent type differences. This is in the spirit of Rosen (1979), Roback (1982).

Each individual inelastically supplies one unit of labor at location wage W_j , and production using only labor. The production technology in location 1 is $y_1 = z_1 n_1$. The production technology in location 2 is $y_2 = A_2 z_2 n_2$, where $A_2 = n_2^\phi$. A_2 is a spillover that is not internalized by the firm in the competitive equilibrium. We assume all markets are competitive.

To mirror the prior example, we assume $z_2 < z_1$, and hence the externality is associated with a location that is higher marginal cost, and hence will be expensive in real terms, but has some knowledge spillovers. A useful example would be an expensive knowledge hub like San Francisco. Here we assume a *place-based transfer* in the form of an income subsidy at rate τ_2 to Location 2 workers, balanced by income taxes at rate τ_1 in Location 1. The core feature of the example is that subsidies to address the spillover are going in the direction of the expensive

location.

Incomplete markets. With incomplete markets, labor supply and product demand $(\rho_1, \rho_2, q_1, q_2)$ are determined by the standard choice probabilities and location-specific budget constraints:

$$\rho_1 = \frac{e^{u(q_1)/\eta}}{e^{u(q_1)/\eta} + e^{u(q_2)/\eta}} \quad , \quad \rho_2 = 1 - \rho_1 \quad , \quad q_1 = \frac{(1 - \tau_1)W_1}{P_1} \quad , \quad q_2 = \frac{(1 + \tau_2)W_2}{P_2}$$

Complete markets. Our previous results allow us to jump to the optimal risk sharing condition. Since the price of the good P is constant across locations, optimal risk-sharing will imply $u'(q_1)/u'(q_2) = 1$, and hence $q_1 = q_2 =: \bar{q}$. Labor supply and product demand $(\rho_1, \rho_2, \bar{q})$ are then determined by (i) the efficient choice probabilities, (ii) the consolidated budget constraint. Here, the argument in the exponent in the efficient choice probability is $[u(q_j) - u'(q_j)(q_j - income_j)]/\eta$. Since q_j is equalized across locations, the $u(\bar{q})/\eta$ terms cancel in the numerator and denominator of the choice probabilities. This extends our novel expressions for efficient choice probabilities (Section 3) to a case when endowments depend on choices:

$$\rho_1 = \frac{e^{u'(\bar{q})(1-\tau_1)W_1/\eta}}{e^{u'(\bar{q})(1-\tau_1)W_1/\eta} + e^{u'(\bar{q})(1+\tau_2)W_2/\eta}} \quad , \quad \rho_2 = 1 - \rho_1 \quad , \quad \bar{q} = \rho_1(1 - \tau_1)W_1 + \rho_2(1 + \tau_2)W_2.$$

Differently to the industrial policy example, even under log utility the incomplete and complete markets choice probabilities do not coincide.¹¹

Firms. In complete and incomplete markets economies, firms optimally equate price P to marginal cost, ignoring the spillover. This implies the following wages in each location:

$$W_1 = P \times z_1 \quad , \quad W_2 = P \times A_2 z_2.$$

Since $z_2 < z_1$, wages are lower in Location 2, and so with incomplete markets consumption will also be lower in Location 2.

Comparing welfare. As above, we compare the optimal subsidy τ_2^* under complete and incomplete markets. We do so for a numerical example with $u(c) = c^{1-\sigma}/(1 - \sigma)$, $\eta = 1/5$, $z_1 = 1$, $z_2 = 0.80$, $\phi = 0.08$ with baseline value of $\sigma = 3$. Again using Proposition 5, we compute optimal policy by maximizing social welfare.

Figure 2 shows a similar set of results to the industrial policy case. First, under incomplete markets, the subsidy props up consumption of individuals that choose to live in Location 2. It does so via two forms of redistribution: directly via taxes (holding wages fixed), and indirectly via increasing wages in Location 2. In the industrial policy case, the subsidy was achieving redistribution via prices. In the spatial policy case, the subsidy is redistributing more directly

¹¹A straight-forward way to observe this is as follows. As long as u is concave, efficiency will imply constant consumption regardless of location choices. Regardless of u , this is not satisfied under incomplete markets as long as there are some wage differences across locations.

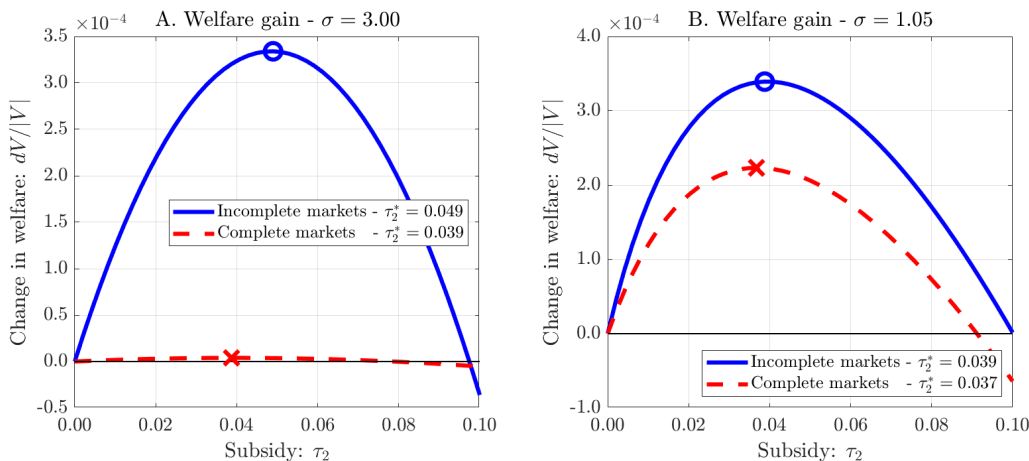


Figure 2: Welfare gains from optimal spatial policy

and also via prices. These effects compress consumption across space, which improves welfare and rationalizes a higher optimal subsidy than under complete markets, where the focus is purely on the inefficiency caused by the spillover. Second, the optimal policy under complete markets delivers much smaller welfare gains, as gains are derived only from addressing the spillover.¹²

Implications for applied research. Again, when markets are incomplete—which we view as the empirically relevant case—it is impossible for a researcher to choose place-based transfers to *only* address the externality. In this example, one may mistakenly conclude that a high subsidy is driven by addressing spillovers, while it is actually driven by redistribution. For applied research this again suggests that computing the optimal policy response to the spillover directly could be obtained by conducting the optimal subsidy analysis in the presence of complete markets.

5.3. Insurance vs. Redistribution.

One might view tastes ξ as permanent characteristics of individuals rather than shocks. It should be clear that maximizing social welfare via $\mathbb{E} [u(q_j^i) + \xi_j^i]$ —which is the practice in the majority of applied work—delivers the same optimal subsidies in both of our examples. What the policy is doing, however, would be described as *redistribution* rather than *insurance*. In the case of permanent characteristics, the policymaker uses τ_2 to *redistribute via prices* in both of our examples. It does so because, under Utilitarian weights, it views it as unfair that one individual might like the expensive Good 2 (Location 2), while another happens to like the cheap Good 1 (Location 1).

¹²The welfare gain under complete markets optimal policy is equivalent to a 0.0045% increase in TFP. The welfare gain under incomplete markets optimal policy is equivalent to a 0.38% increase in TFP.

6. Measuring the welfare effects of price changes

We have showed that understanding the incomplete markets nature of discrete choice economies is important for welfare analysis in these economies. Here, we show that, in a discrete choice, incomplete markets setting, the empirical approach of measuring welfare effects of price changes from data on price changes and budget shares fails. As above, we conclude with concrete proposals for applied work based on the theory.

We start from standard results in consumer theory regarding the measurement of the welfare effects of change in prices. Deaton (1989) is the seminal reference, with the insight that initial budget shares are the sufficient statistics needed to compute the welfare effects of price changes to a first order. This approach has gained increasing popularity due to recent availability of itemized household consumption data and the seemingly assumption-free nature of the exercise. Del Canto et al. (2023), Fagereng et al. (2022), and Borusyak and Jaravel (2021) are recent, important applications of this approach. We extend this approach to a discrete choice economy and show that the applicability of Deaton (1989)-like results depend crucially on market structure. In the empirically irrelevant case of complete markets, these formulas hold. In the empirically relevant case of incomplete markets, these formulas fail.

Equivalent variation. Following the literature, we focus on the equivalent variation welfare metric (EV). To simplify the presentation, we strip back our notation and focus on a simplified economy: there is a single household type, only J goods with prices $\mathbf{p} = (p_1, \dots, p_J)$, and no outside good. Income y is denominated in some numeraire. Let $V(\mathbf{p}, y)$ be the indirect utility function from the household's choice problem. Consider some alternative vector of prices $\mathbf{p}' = \mathbf{p} + d\mathbf{p}$. EV is defined as the percent change in income ϕ at initial prices that satisfies

$$V(\mathbf{p} + d\mathbf{p}, y) = V(\mathbf{p}, y + dy), \quad \phi = \frac{dy}{y}. \quad (35)$$

Define $\tilde{\phi}$ as the measure of EV to a first order, around $d\mathbf{p} = \mathbf{0}$ and $dy = 0$. Then

$$\tilde{\phi} = \sum_{j=1}^J \frac{V_{p_j}(\mathbf{p}, y)}{V_y(\mathbf{p}, y)} \frac{p_j}{y} \Delta \log p_j. \quad (36)$$

This is general. We now compute $V_{p_j}(\mathbf{p}, y)$ and $V_y(\mathbf{p}, y)$ under different assumptions about the choice problem (discrete or continuous choice) and market structure (complete or incomplete markets). Thus, our general approach illustrates how these different problems and settings relate to each other.

6.1. Continuous choice model.

First, we consider the continuous choice model, which provides an overview of the standard results like those in Deaton (1989). Let $u(q_1, \dots, q_J)$ be the utility function from consuming *all* the J goods in continuous quantities that satisfies standard assumptions. The budget constraint is $\sum_j p_j q_j = y$, and the envelope theorem delivers $V_{p_j}(\mathbf{p}, y)$ and $V_y(\mathbf{p}, y)$:

$$V_{p_j}(\mathbf{p}, y) = -\lambda q_j \quad , \quad V_y(\mathbf{p}, y) = -\lambda, \quad (37)$$

where λ is the multiplier on the household's budget constraint. Using (36), EV is

$$\tilde{\phi} = -\sum_{j=1}^J \left(\frac{p_j q_j}{y} \right) \Delta \log p_j. \quad (38)$$

Equivalent variation is how much prices change weighted by initial budget shares. This is the standard result implemented in applied empirical studies like those cited above.

6.2. Discrete Choice - Complete Markets.

Now, consider the discrete choice model under complete markets as described in Section 3. The consolidated budget constraint is $\int \sum_j \varphi(\boldsymbol{\xi}) x_j(\boldsymbol{\xi}) p_j q_j(\boldsymbol{\xi}) d\boldsymbol{\xi} = y$. As we did above, we can compute $V_{p_j}(\mathbf{p}, y)$ and $V_y(\mathbf{p}, y)$:

$$V_{p_j}(\mathbf{p}, y) = -\lambda \int \sum_j \varphi(\boldsymbol{\xi}) x_j(\boldsymbol{\xi}) q_j(\boldsymbol{\xi}) d\boldsymbol{\xi} = -\lambda \rho_j q_j, \quad V_y(\mathbf{p}, y) = \lambda \int \varphi(\boldsymbol{\xi}) d\boldsymbol{\xi} = \lambda. \quad (39)$$

Here, λ is the multiplier on the unified budget constraint across all possible realizations of the taste shock. The final equalities result from arguments similar to those Section 3: equilibrium state prices $\varphi(\boldsymbol{\xi})$ are actuarially fair, and $q_j(\boldsymbol{\xi})$ is independent of preferences $\boldsymbol{\xi}$. Inserting these into (36), the first order equivalent variation measure is

$$\tilde{\phi} = -\sum_{j=1}^J \left(\frac{\rho_j p_j q_j}{y} \right) \Delta \log p_j. \quad (40)$$

Our key result is that complete markets are necessary to deliver a formula similar to that used in applied work (38). Complete markets allow households to face a unified budget constraint and effectively behave and experience welfare gains *as if* consuming all goods at once.

The key difference between (38) and (40) is that the choice probability now enters into the share calculation. One interpretation of (40) is from the perspective of an individual, the expected budget share $\frac{\rho_j p_j q_j}{y}$ is the welfare-relevant share. A second interpretation of (40) is that it is the aggregated budget share across all households in the economy. Some individuals purchase bananas, some purchase apples, but when aggregating budget shares across these discrete choices, they are $\rho_j \times \frac{p_j q_j}{y}$.

This last observation makes how to interpret the applied practice of aggregating budget shares across households when underlying microdata are discrete. Suppose a researcher had scanner or survey data on a *group* of individuals of this type. For example, Del Canto et al. (2023) group together high-school-educated households aged 35-40 and use the CEX to compute the budget shares of *that group* on a collection of goods. One interpretation behind this measurement is that each individual consumes a little of every single good. An alternative interpretation is that individuals are making discrete choices across the different goods. What equations (38) and (40) say is that these two approaches measure the same object, $\tilde{\phi}$, only under complete markets. From this perspective, we have generalized the interpretation and foundation behind results like the ones in Deaton (1989). The key restriction, however, is that one must believe that complete markets are available.

6.3. Discrete Choice - Incomplete Markets.

The empirically relevant case of the discrete choice model with incomplete markets delivers a substantively different formula relative to (40) and hence also incompatible with empirical approaches. Computing $V_{p_j}(\mathbf{p}, y)$ and $V_y(\mathbf{p}, y)$ with incomplete markets gives

$$V_{p_j}(\mathbf{p}, y) = - \int x_j(\boldsymbol{\xi}) \lambda_j(\boldsymbol{\xi}) q_j(\boldsymbol{\xi}) d\boldsymbol{\xi} = -\rho_j \frac{u'(q_j) q_j}{p_j}, \quad V_y(\mathbf{p}, y) = \int \sum_j x_j(\boldsymbol{\xi}) \lambda_j(\boldsymbol{\xi}) d\boldsymbol{\xi} = \sum_j \rho_j \frac{u'(q_j)}{p_j}.$$

The consequence of incomplete markets is twofold. First, the cost of a price change is siloed only into states where j is consumed, rather than equalized as in the complete markets case. Second, the value of income y varies across purchases. Combining these terms, we obtain our third expression for equivalent variation:

$$\tilde{\phi} = - \sum_{j=1}^J \left(\frac{\rho_j u'(q_j)/p_j}{\sum_k \rho_k u'(q_k)/p_k} \right) \left(\frac{p_j q_j}{y} \right) \Delta \log p_j. \quad (41)$$

With incomplete markets an additional factor twists the weighting toward states in which the individual has a high marginal value of income.¹³ Intuitively, if the individual cannot move resources into states where they like expensive choices (a high $\lambda_j(\boldsymbol{\xi})$), then further increases in prices in these states will be more costly.

Relationship with Small and Rosen (1981). Note that the ‘standard formula’ (40) obtains once more if one takes the approach of Small and Rosen (1981). Small and Rosen (1981) operate in an environment where the marginal value of wealth would vary across choices, which is the environment that we have also studied here. However, to compute welfare formulas they

¹³Inspecting (41) connects to our Anderson et al. (1987) example, where the incomplete markets allocation aligned with the complete markets allocation. Specifically, with log preferences $u'(q_j)/p_j = 1/y$ for all j goods and (41) collapses back to the complete markets expression. Again, this is irrespective of the distribution of idiosyncratic preferences.

impose an additional assumption that the marginal value of wealth is constant across choices.¹⁴ In our context, this is the multiplier on the incomplete markets budget constraint λ_j , which equals $u'(q_j)/p_j$. If this is constant across choices in (41), then the standard formula (40) is obtained.

6.4. Implications for applied work.

Regardless of the distribution of preferences, these results show that ‘off-the-shelf’ formulas for the welfare effects of price changes derived in Deaton-like environments only hold in discrete choice economies under assumptions of complete markets. If we believe that markets are instead incomplete—as is no doubt the empirically relevant case—such formulas do not apply. This presents a significant issue for applied work, as evaluating (41) requires additional assumptions on preferences that are otherwise avoided in the Deaton-like formula. We now try to understand what else a researcher must compute to evaluate (41).

In the empirically relevant case of incomplete markets the following conventional wisdom fails: *elasticities of demand do not need to be computed to evaluate the first order welfare effects of price changes*. In the natural setting of incomplete markets, elasticities of demand are necessary inputs.

To see this, define ε_j as what a researcher would measure as a the price elasticity of demand:

$$\varepsilon_j = -\frac{\partial \log \text{Quantity Sold}_j}{\partial \log p_j} = -\frac{\partial \log [\rho_j q_j]}{\partial \log p_j} = \underbrace{\left[-\frac{\partial \log \rho_j}{\partial \log p_j} \right]}_{\text{Extensive margin}} + \underbrace{\left[-\frac{\partial \log q_j}{\partial \log p_j} \right]}_{\text{Intensive margin}}. \quad (42)$$

If we additionally assume that taste shocks are type 1 extreme value distributed, the price elasticity of demand is: $\varepsilon_j = \eta (1 - \rho_j) u'(q_j) q_j - 1$. Using this in the incomplete markets expression for equivalent variation (41), we arrive at

$$\tilde{\phi} = -\sum_j \left(\frac{\left(\frac{\rho_j}{1-\rho_j} \right) (1 + \varepsilon_j)}{\sum_k \left(\frac{\rho_k}{1-\rho_k} \right) (1 + \varepsilon_k)} \right) \Delta \log p_j. \quad (43)$$

To a first order, elasticities of demand appear in the computation of equivalent variation *under incomplete markets*. Marginal utility is not equated across goods choices, and thus changes in prices and the substitution that they induce have first-order welfare impacts. Under complete markets, these margins are equated and equivalent variation aligns with the conventional wisdom that budget shares are the sufficient to characterize the welfare effects of price changes. We conclude that if seeking ‘model-free’ estimates of *first order* welfare effects of price changes, applied researchers require good estimates of elasticities of demand.

Our result is worth contrasting with Auer et al.’s (2022) presentation of the welfare effects of

¹⁴On page 124 of Small and Rosen (1981) they “make three assumptions” including that the “marginal utility of income, is approximately independent of the price of the good”.

price changes. Starting from a standard continuous choice problem like that described above, their first order components are as in (36). They emphasize that elasticities of demand show up in *second order* components of total equivalent variation, ϕ . Their contribution is to show, empirically, that these second order components are large and systematically vary across rich and poor households. Our contribution is to show that in a discrete choice setting with incomplete markets, elasticities of demand show up to the *first order*, $\tilde{\phi}$. Hence, our results are complementary, in that both foreground the importance of elasticity of demand estimates in computing the welfare effects of price changes.

7. Aggregation and efficiency in discrete choice models

As a final application, we consider issues of aggregation in discrete choice economies, revisiting a famous result in the theoretical literature. Issues of efficiency and the existence of a representative household are sometimes linked. Our theoretical results allow us to assess efficiency in discrete choice economies and hence understand when aggregation and efficiency go hand-in-hand and when they do not.

We study (i) the Anderson et al. (1987, henceforth, ADT) economy which yields demand equivalent to a representative household with CES preferences, (ii) extend this example to the case of multiplicative shocks with Fréchet distributed preferences. In the first case, we contribute new results that show how the functional form assumption on utility in ADT—rather than features of the distribution of preferences—imply a knife-edge case where the incomplete markets allocation is efficient. From this well known example, one might think that aggregation and efficiency go hand-in-hand. In the second case, we show that this is not true. In the case of multiplicative Fréchet distributed preferences, aggregation to a CES representative household occurs under both incomplete and complete markets, but with different coefficients.

Anderson et al. (1987) economy. A key example of interest in the study of discrete choice models is the ADT economy. In their economy aggregated demands for each commodity are as if they came from a representative consumer with a CES utility function over the commodities. We show that the ADT economy has the unique feature that the incomplete markets allocation aligns with the complete markets allocation, and hence is efficient. We show that this result is independent of the taste shock distribution, but instead depends on specific functional form for preferences over the differentiated good.

As a special case of our general discrete choice economy, the ADT economy has the following features: (i) no outside good, (ii) continuous choice in q_j , (iii) no quality differences, (iv) common preferences across θ types, and (iv) different endowments across θ types. Finally, and most importantly, preferences in the ADT economy are log over the differentiated good, $u(q_j) = \log q_j$.

Given these assumptions, we have the following objective function and budget constraints:

$$\int \sum_j x_j^i(\boldsymbol{\xi}, \theta) \left[u(q_j(\boldsymbol{\xi}, \theta)) + \xi_j^i \right] dG(\boldsymbol{\xi}) , \quad \underbrace{p_j q_j(\theta) = W(\theta)}_{\text{Incomplete markets, for each } j} \quad \underbrace{\int_{\boldsymbol{\xi}} \varphi(\boldsymbol{\xi}) x_j^i(\boldsymbol{\xi}, \theta) p_j q_j(\theta) d\boldsymbol{\xi} = W(\theta)}_{\text{Complete markets}} .$$

Consumption policy. In this economy, the allocations under incomplete markets and complete markets satisfy the same conditions. Consider the first order condition for $q_j(\boldsymbol{\xi}, \theta)$ with incomplete markets:

$$\lambda_j(\theta) = \frac{u'(q_j(\theta))}{p_j} = \frac{1}{p_j q_j(\theta)} = \frac{1}{W(\theta)} = \frac{1}{p_{j'} q_{j'}(\theta)} = \frac{u'(q_j(\theta))}{p_j} = \lambda_{j'}(\theta). \quad (44)$$

This delivers the main intuition for the result. Simply consuming $q_j(\theta) = W(\theta)/p_j$, as dictated by the budget constraint, also equalizes the marginal value of wealth across goods (as in (19)) when utility is log. When choosing a good that is 1 percent more expensive, the budget constraint dictates that consumption q_j is 1 percent lower, and log preferences dictate that marginal utility increases by exactly 1 percent. Hence, across goods of different prices, the ratio of marginal utilities is equated to the ratio of relative prices. Note that this is independent of assumptions on distributions of idiosyncratic preferences. With a constant multiplier, regardless of the good that is chosen, the key condition of efficiency is obtained.

Discrete choice rule. Starting from the complete markets choice rule (20), we substitute in the first order condition for $q_j(\theta)$: $\lambda(\theta)p_j = u'(q_j(\theta))$. This gives the first expression here:

$$x_j(\boldsymbol{\xi}, \theta) = \begin{cases} 1, & \text{if } u(q_j(\theta)) + \xi_j^i - u'(q_j(\theta))q_j(\theta) \geq \\ & \max_k [u(q_k(\theta)) + \xi_k - u'(q_k(\theta))q_k(\theta)] \\ 0, & \text{otherwise.} \end{cases} = \begin{cases} 1, & \text{if } u(q_j(\theta)) + \xi_j^i \geq \\ & \max_k [u(q_k(\theta)) + \xi_k] \\ 0, & \text{otherwise.} \end{cases} \quad (45)$$

Choosing good j removes $p_j q_j$ from the complete markets budget constraint, the marginal cost of which is $\lambda_j p_j q_j$. Optimality of q_j implies that $\lambda_j = u'(q_j)/p_j$, which with log preferences is $1/p_j q_j$. Hence the marginal cost is constant across choices and the additional term in the complete markets choice rule drops out across choices, leaving the incomplete markets choice rule (i.e. $u'(q_j)q_j = 1$). Thus, applying our earlier results in this context shows that the ADT economy has the unique feature that the incomplete markets allocation aligns with the complete markets allocation and hence is efficient. Furthermore, this feature depends only on the form of the utility function and not distributional assumptions on preferences.

This delivers a striking lesson regarding the aggregation result of Anderson et al. (1992). We did not invoke a distributional assumption on the taste shocks. Hence, the alignment between complete and incomplete markets in this setting is only about the shape of the utility function.

We therefore have shown that log utility *and* Type I extreme value preferences yield aggregation. But the efficiency properties are independent of the shocks.

Implications for applied work - Aggregation to CES \neq Efficiency. A common misconception is that the ability to aggregate to a representative consumer—as in the ADT case—and efficiency somehow go hand-in-hand. We provide two examples that show that this is not the case.

First, in the ADT case, we showed that the economy is efficient regardless of the distribution of taste shocks, but only obtains a CES aggregation under a Type I extreme value distribution. Hence, we have provided examples of an economy which is efficient, but may or may not to CES demand. Second, in Appendix E we study the multiplicative case, where utility is $u(q_j(\theta)\xi_j^i)$. We provide general results and then specialize to the case where ξ_j^i are drawn from independent Fréchet distributions and u is isoelastic. In this special case we show that *both* the complete markets economy and the incomplete markets economy deliver an aggregate demand system as if it was obtained from a CES representative consumer. However (a) the incomplete markets remains inefficient, and (b) the representative consumer CES utility functions are different across market structures. Hence, we have also provided an example of efficient and inefficient economies that both aggregate to CES demand.

Summary. In summary, our theoretical results allow us to distinguish between efficient and inefficient discrete choice economies. We have applied these results to show that efficient and inefficient economies can both aggregate to CES demand. We have applied these results to show that efficient and inefficient economies can both fail to aggregate to CES demand. Hence, inferences about efficiency cannot be made by reference to the existence of a representative CES household.

8. Conclusion

We started from two observations. First, discrete choice models are a powerful and simple framework for thinking about choice problems. Second, seminars or papers about the welfare impacts from various shocks in this class of models would raise a persistent question: what are properties of the allocations from which welfare is evaluated? Given the additivity of the shock, the answers seemed not interesting. This was not the case. These economies are generically inefficient. We then showed how the complete markets and Pareto efficient allocations take simple intuitive forms and how they can be put to use in a broad array of settings to understand departures from efficiency. Importantly for applied work, we showed that complete markets can be used as a tool to separate out market incompleteness from other inefficiencies—e.g. production externalities, amenity spillovers, lack of perfect competition—that researchers are interested in.

The purpose of this paper is not to take a stand, per se, on the extent of market incompleteness.

ness—on the contrary. Discrete choice models are often used to (i) account for heterogeneity in behavior in a parsimonious and flexible way and (ii) evaluate a specific policy— for example, a place-based policy or merger. But how does one evaluate the welfare gains from the policy versus the welfare effects that arise because of how heterogeneity is introduced? This is what our complete markets formulation achieved. And we hope that by providing an appropriate benchmark to isolate and evaluate the effects of the specific policy, it is useful.

There are several areas for future research. One is the extension to dynamic frameworks. In the background, and as alluded to at various points in the text, we very much thought of the static problem as being a representation of dynamic environments. That’s not quite the case, though, because the choice today often influences values in the future, and thus efficient and complete markets allocations may be different than what we characterized. The characterization of socially optimal rural-urban migration flows in Lagakos et al. (2023) is one example; Donald et al. (2023a) is another example, studying optimal policy in dynamic spatial models.

The second question regards the role of partial / self-insurance and discrete choices. The standard incomplete markets model, where households can imperfectly insure against labor income risk, has proven to be a very useful laboratory for thinking about household heterogeneity and consumption dynamics. Given our results, we think that there are many interesting questions for research regarding how partial insurance influences discrete choices such as where one lives, which good you buy or employer you work for.

References

- ALES, L. AND C. SLEET (2022): "Optimal Taxation of Income-Generating Choice," *Econometrica*, 90, 2397–2436.
- ALLEN, R. AND J. REHBECK (2019): "Identification With Additively Separable Heterogeneity," *Econometrica*, 87, 1021–1054.
- ANDERSON, S. P., A. DE PALMA, AND J.-F. THISSE (1987): "The CES is a discrete choice model?," *Economics Letters*, 24, 139–140.
- (1992): *Discrete choice theory of product differentiation*, MIT press.
- ARTUÇ, E., S. CHAUDHURI, AND J. MCLAREN (2010): "Trade shocks and labor adjustment: A structural empirical approach," *The American Economic Review*, 100, 1008–1045.
- AUER, R., A. BURSTEIN, S. M. LEIN, AND J. VOGEL (2022): "Unequal expenditure switching: Evidence from Switzerland," Tech. rep., National Bureau of Economic Research.
- BACKUS, D. K. AND G. W. SMITH (1993): "Consumption and real exchange rates in dynamic economies with non-traded goods," *Journal of International Economics*, 35, 297 – 316.
- BERGER, D., K. HERKENHOFF, AND S. MONGEY (2022): "Labor market power," *American Economic Review*, 112, 1147–93.
- BERRY, S., J. LEVINSOHN, AND A. PAKES (1995): "Automobile prices in market equilibrium," *Econometrica*, 63, 841–890.
- BHATTACHARYA, D. (2021): "The Empirical Content of Binary Choice Models," *Econometrica*, 89, 457–474.
- BORUSYAK, K. AND X. JARAVEL (2021): "The distributional effects of trade: Theory and evidence from the united states," Tech. rep., National Bureau of Economic Research.
- BUSSO, M., J. GREGORY, AND P. KLINE (2013): "Assessing the Incidence and Efficiency of a Prominent Place Based Policy," *American Economic Review*, 103, 897–947.
- CALIENDO, L., M. DVORKIN, AND F. PARRO (2019): "Trade and labor market dynamics: General equilibrium analysis of the China trade shock," *Econometrica*, 87, 741–835.
- CARD, D., A. R. CARDOSO, J. HEINING, AND P. KLINE (2018): "Firms and labor market inequality: Evidence and some theory," *Journal of Labor Economics*, 36, S13–S70.
- DAVIS, M. AND J. M. GREGORY (2023): "Place-Based Redistribution in Location-Choice Models," Tech. rep.

- DEATON, A. (1989): "Rice prices and income distribution in Thailand: A non-parametric analysis," *Economic Journal*, 99, 1–37.
- DEL CANTO, F. N., J. R. GRIGSBY, E. QIAN, AND C. WALSH (2023): "Are inflationary shocks regressive? A feasible set approach," Tech. rep., National Bureau of Economic Research.
- DIAMOND, R. (2016): "The determinants and welfare implications of US workers' diverging location choices by skill: 1980–2000," *American Economic Review*, 106, 479–524.
- DONALD, E., M. FUKU, AND Y. MIYAUCHI (2023a): "Optimal Dynamic Spatial Policy," Tech. rep.
- (2023b): "Unpacking Aggregate Welfare in a Spatial Economy," Tech. rep.
- FAGERENG, A., M. GOMEZ, E. GOUIN-BONENFANT, M. HOLM, B. MOLL, AND G. NATVIK (2022): "Asset-price redistribution," Tech. rep., Working Paper.
- FAJGELBAUM, P. D. AND C. GAUBERT (2020): "Optimal spatial policies, geography, and sorting," *The Quarterly Journal of Economics*, 135, 959–1036.
- KENNAN, J. AND J. R. WALKER (2011): "The effect of expected income on individual migration decisions," *Econometrica*, 79, 211–251.
- KLINE, P. AND E. MORETTI (2014): "People, places, and public policy: Some simple welfare economics of local economic development programs," *Annu. Rev. Econ.*, 6, 629–662.
- LAGAKOS, D., M. MOBARAK, AND M. E. WAUGH (2023): "The welfare effects of encouraging rural-urban migration," *Econometrica*, 90, 803–837.
- MCFADDEN, D. (1974): "Conditional logit analysis of qualitative choice behavior," in *Frontiers in Econometrics*, edited by P. Zarembka, 105–142, Academic Press.
- NAKAMURA, E. AND D. ZEROM (2010): "Accounting for incomplete pass-through," *Review of Economic Studies*, 77, 1192–1230.
- NEVO, A. (2000): "A practitioner's guide to estimation of random-coefficients logit models of demand," *Journal of Economics & Management Strategy*, 9, 513–548.
- REDDING, S. J. AND E. ROSSI-HANSBERG (2017): "Quantitative spatial economics," *Annual Review of Economics*, 9, 21–58.
- ROBACK, J. (1982): "Wages, rents, and the quality of life," *Journal of political Economy*, 90, 1257–1278.

- ROGERSON, R. (1988): "Indivisible labor, lotteries and equilibrium," *Journal of monetary Economics*, 21, 3–16.
- ROSEN, S. (1979): "Wage-based indexes of urban quality of life," *Current issues in urban economics*, 74–104.
- SMALL, K. A. AND H. S. ROSEN (1981): "Applied Welfare Economics with Discrete Choice Models," *Econometrica*, 49, 105–130.
- TOWNSEND, R. M. (1994): "Risk and insurance in village India," *Econometrica*, 539–591.
- TRAIN, K. E. (2009): *Discrete choice methods with simulation*, Cambridge university press.

Appendix

A. Appendix: Incomplete Markets

Here walk through the arguments regarding how we solve the agents problem in incomplete markets. Let us state the problem:

$$\max_{c_j^i(\boldsymbol{\xi}, \theta), q_j^i(\boldsymbol{\xi}, \theta), x_j^i(\boldsymbol{\xi}, \theta)} \int_{\boldsymbol{\xi}} \sum_j x_j^i(\boldsymbol{\xi}, \theta) \left\{ u[c^i(\boldsymbol{\xi}, \theta), q_j^i(\boldsymbol{\xi}, \theta); j, \theta] + \xi_j^i \right\} g(\boldsymbol{\xi}; \theta) d\boldsymbol{\xi}, \quad (\text{A1})$$

$$\text{subject to: } [\lambda^i(\boldsymbol{\xi}, \theta)] : \sum_j x_j^i(\boldsymbol{\xi}, \theta) [c^i(\boldsymbol{\xi}, \theta) + p_j q_j^i(\boldsymbol{\xi}, \theta)] \leq W(\theta), \quad \text{for all } \boldsymbol{\xi}. \quad (\text{A2})$$

A household chooses consumption quantities, and commodity choices for every possible state subject to the constraint is that consumption equals the households endowment — state by state. The Lagrangian associated with this problem is

$$\mathcal{L} = \max_{c_j(\boldsymbol{\xi}), q_j(\boldsymbol{\xi}), x_j(\boldsymbol{\xi})} \int_{\boldsymbol{\xi}} \sum_j x_j^i(\boldsymbol{\xi}, \theta) \left\{ u[c^i(\boldsymbol{\xi}, \theta), q_j^i(\boldsymbol{\xi}, \theta); j, \theta] + \xi_j^i \right\} g(\boldsymbol{\xi}; \theta) d\boldsymbol{\xi}, \quad (\text{A3})$$

$$+ \int_{\boldsymbol{\xi}} \lambda^i(\boldsymbol{\xi}, \theta) \left\{ W(\theta) - \sum_j x_j^i(\boldsymbol{\xi}, \theta) [c^i(\boldsymbol{\xi}, \theta) + p_j q_j^i(\boldsymbol{\xi}, \theta)] \right\} g(\boldsymbol{\xi}, \theta) d\boldsymbol{\xi}. \quad (\text{A4})$$

The strategy is to characterize necessary conditions that determine consumption and then characterize the rule which determines which good to chose.

The first order condition for consumption is

$$x_j^i(\boldsymbol{\xi}, \theta) u_{q_j} [c^i(\boldsymbol{\xi}, \theta), q_j^i(\boldsymbol{\xi}, \theta); j, \theta] g(\boldsymbol{\xi}, \theta) = \lambda^i(\boldsymbol{\xi}, \theta) x_j^i(\boldsymbol{\xi}, \theta) p_j g(\boldsymbol{\xi}, \theta), \quad (\text{A5})$$

Notice that we do not have to take a stand on the value of $x_j^i(\boldsymbol{\xi}, \theta)$. To see this notice that there are two cases: (i) the first order condition is trivially satisfied as $x_j^i(\boldsymbol{\xi}, \theta) = 0$ or (ii) $x_j^i(\boldsymbol{\xi}, \theta) = 1$. Furthermore, notice how the pdf for the taste shock shows up on both the left and the right hand side, hence, we can canceling terms giving

$$u_{q_j} [c^i(\boldsymbol{\xi}, \theta), q_j^i(\boldsymbol{\xi}, \theta); j, \theta] = \lambda^i(\boldsymbol{\xi}, \theta) p_j. \quad (\text{A6})$$

Then similarly for the non-differentiated good

$$x_j^i(\boldsymbol{\xi}, \theta) u_c [c^i(\boldsymbol{\xi}, \theta), q_j^i(\boldsymbol{\xi}, \theta); j, \theta] g(\boldsymbol{\xi}, \theta) = \lambda^i(\boldsymbol{\xi}, \theta) x_j^i(\boldsymbol{\xi}, \theta) g(\boldsymbol{\xi}, \theta), \quad (\text{A7})$$

and then canceling terms we have

$$u_c [c^i(\boldsymbol{\xi}, \theta), q_j^i(\boldsymbol{\xi}, \theta); j, \theta] = \lambda^i(\boldsymbol{\xi}, \theta). \quad (\text{A8})$$

Then inspecting these conditions we make several observations. First, all that matters for the multiplier is the choice, not the realization $\boldsymbol{\xi}$, since any $\boldsymbol{\xi}$ and $\boldsymbol{\xi}'$ that lead to the same choice have the same budget constraint. Second, and because there is no dependence on the left-hand side or right-hand side on the shock, we can drop the dependence of consumption on the realization $\boldsymbol{\xi}$. However, it does depend upon the choice j . Third, we can drop the indexing of the multiplier by i . The argument is that the only thing that might make i households different are the realizations $\boldsymbol{\xi}$, but now the consumption allocations for all i don't depend upon the realization, endowments are the same across i , so the multiplier does not depend upon i . These arguments imply that the first order conditions characterizing consumption allocations are

$$u_c [c(j, \theta), q_j(\theta); j, \theta] = \lambda_j(\theta) \quad \text{and} \quad u_{q_j} [c(j, \theta), q_j(\theta); j, \theta] = \lambda_j(\theta)p_j. \quad (\text{A9})$$

The next step is to characterize the $x_j(\boldsymbol{\xi}, \theta)$ that is the rule mapping the realization of the taste shock into which good should be chosen. We do this incrementally on the Lagrangian by thinking through which j gives the most utility for a given $\boldsymbol{\xi}$). So fix a realization, $\boldsymbol{\xi}$, then compare utility across those events in the Lagrangian...

$$[u [c^i(\boldsymbol{\xi}, \theta), q_1^i(\boldsymbol{\xi}, \theta); 1, \theta] + \xi_1^i]g(\boldsymbol{\xi}, \theta) + \lambda^i(\boldsymbol{\xi}, \theta)g(\boldsymbol{\xi}, \theta) [W(\theta) - c^i(\boldsymbol{\xi}, \theta) - p_1q_1^i(\boldsymbol{\xi}, \theta)] \quad \text{vs.} \quad (\text{A10})$$

$$[u [c^i(\boldsymbol{\xi}, \theta), q_2^i(\boldsymbol{\xi}, \theta); 2, \theta] + \xi_2^i]g(\boldsymbol{\xi}, \theta) + \lambda^i(\boldsymbol{\xi}, \theta)g(\boldsymbol{\xi}, \theta) [W(\theta) - c^i(\boldsymbol{\xi}, \theta) - p_2q_2^i(\boldsymbol{\xi}, \theta)] \dots \quad (\text{A11})$$

Now the budget constraint always binds for every realization $\boldsymbol{\xi}$ independent of the good chosen. This implies that the second terms in the comparison are always zero. Then inserting our arguments about how consumption, the rule describing which good to consume is

$$x_j^i(\boldsymbol{\xi}, \theta) = \begin{cases} 1, & \text{if } u [c(j, \theta), q_j(\theta); j, \theta] + \xi_j^i \geq \\ & \max_{j'} \left\{ u [c(j', \theta), q_{j'}(\theta); j', \theta] + \xi_{j'}^i \right\} \\ 0. & \text{otherwise} \end{cases} \quad (\text{A12})$$

As emphasized in the text, this is the same choice rule that is the starting point from the literature: see what the realized shocks are and then the utility associated with consumption amongst the different choices and then chose the one that delivers highest utility. however, we

did not employ an ex-post approach, but an ex-ante approach which, with incomplete markets, delivers the same choice rule.

A. Appendix: Complete Markets

Here walk through several arguments. First, we state the households problem and then provide an argument that the complete markets equilibrium is Pareto efficient. Then the second is a characterization of the necessary conditions behind the households optimization problem.

A.1. The Households Problem

Here we state the households problem. It is

$$\max_{a_j(\boldsymbol{\xi}), c(\boldsymbol{\xi}), q_j(\boldsymbol{\xi}), x_j(\boldsymbol{\xi})} \int_{\boldsymbol{\xi}} \sum_j x_j^i(\boldsymbol{\xi}, \theta) \left\{ u[c^i(\boldsymbol{\xi}, \theta), q_j^i(\boldsymbol{\xi}, \theta); j, \theta] + \xi_j^i \right\} g(\boldsymbol{\xi}; \theta) d\boldsymbol{\xi}, \quad \text{subject to:} \quad (\text{A1})$$

$$\sum_j x_j^i(\boldsymbol{\xi}, \theta) [p_o c^i(\boldsymbol{\xi}, \theta) + p_j q_j^i(\boldsymbol{\xi}, \theta)] \leq W(\theta) + a^i(\boldsymbol{\xi}, \theta) \quad \forall \boldsymbol{\xi}, \quad (\text{A2})$$

$$\int_{\boldsymbol{\xi}} \varphi(\boldsymbol{\xi}, \theta) a^i(\boldsymbol{\xi}, \theta) d\boldsymbol{\xi} = 0, \quad (\text{A3})$$

where $\varphi(\boldsymbol{\xi}, \theta)$ are the actuarially fair state prices. A household chooses contingent claims, consumption quantities, and commodity choices for every possible state subject to two constraints. The first constraint is that consumption equals endowments and amount of insurance purchased — state by state. The second constraint says that the household's net asset position must be zero.

Then rewrite the problem in (A1) by substituting in the budget constraints from (A2) into the constraint for assets (A3). The resulting problem is

$$\max_{a_j(\boldsymbol{\xi}), c_j(\boldsymbol{\xi}), q_j(\boldsymbol{\xi}), x_j(\boldsymbol{\xi})} \int_{\boldsymbol{\xi}} \sum_j x_j^i(\boldsymbol{\xi}, \theta) \left\{ u[c^i(\boldsymbol{\xi}, \theta), q_j^i(\boldsymbol{\xi}, \theta); j, \theta] + \xi_j^i \right\} g(\boldsymbol{\xi}; \theta) d\boldsymbol{\xi}, \quad \text{subject to:} \quad (\text{A4})$$

$$[\lambda^i(\theta)] \int_{\boldsymbol{\xi}} \varphi(\boldsymbol{\xi}, \theta) \left\{ W(\theta) - \sum_j x_j^i(\boldsymbol{\xi}, \theta) [c^i(\boldsymbol{\xi}, \theta) + p_j q_j^i(\boldsymbol{\xi}, \theta)] \right\} d\boldsymbol{\xi} = 0. \quad (\text{A5})$$

where the new constraint has the interpretation of the “lifetime” budget constraint and now there is just one multiplier $\lambda^i(\theta)$ on this one constraint unlike the incomplete markets problem where there are state by state multipliers. Then from here we can define the Complete Markets

Equilibrium:

Definition 3 (The Complete Markets Equilibrium.) A complete markets equilibrium are allocations $c^i(\boldsymbol{\xi}, \theta)$, $q_j^i(\boldsymbol{\xi}, \theta)$, $a^i(\boldsymbol{\xi}, \theta)$, $x_j^i(\boldsymbol{\xi}, \theta)$ for each i and type θ ; goods prices p_j and state prices $\varphi(\boldsymbol{\xi}, \theta)$ such that

i The consumption allocations $(c^i(\boldsymbol{\xi}, \theta)$, $q_j^i(\boldsymbol{\xi}, \theta)$, and $x_j^i(\boldsymbol{\xi}, \theta)$) and asset positions $a^i(\boldsymbol{\xi}, \theta)$ satisfy the household's problem in (A4);

ii State prices satisfy (12);

iii Goods and asset markets clear.

A.2. Proof of First Welfare Theorem

Given our definition of a Complete Markets equilibrium, we argue that it is Pareto efficient. To establish our argument, first note that any alternative allocation $\{\tilde{c}^i(\boldsymbol{\xi}, \theta), \tilde{q}_j^i(\boldsymbol{\xi}, \theta), \tilde{a}^i(\boldsymbol{\xi}, \theta), \tilde{x}_j^i(\boldsymbol{\xi}, \theta)\}$ that is preferred by agent i (in the sense that it provides higher utility than (16)) must imply that

$$\int_{\boldsymbol{\xi}} \varphi(\boldsymbol{\xi}, \theta) \left\{ W(\theta) - \sum_j \tilde{x}_j^i(\boldsymbol{\xi}, \theta) [c^i(\boldsymbol{\xi}, \theta) + p_j \tilde{q}_j^i(\boldsymbol{\xi}, \theta)] \right\} d\boldsymbol{\xi} < 0. \quad (\text{A6})$$

In words, this alternative and preferred allocation is not budget feasible. Now consider an alternative allocation $\tilde{c}^i(\boldsymbol{\xi}, \theta), \tilde{q}_j^i(\boldsymbol{\xi}, \theta), \tilde{a}^i(\boldsymbol{\xi}, \theta), \tilde{x}_j^i(\boldsymbol{\xi}, \theta)$ for all i and θ types that is (i) feasible and (ii) is preferred by all i and θ types. Feasibility implies that

$$\sum_{\theta} \int_0^{\mu(\theta)} y_o^i(\theta) di = \sum_{\theta} \int_{\boldsymbol{\xi}} \int_0^{\mu(\theta)} c^i(\boldsymbol{\xi}, \theta) g(\boldsymbol{\xi}, \theta) di d\boldsymbol{\xi} \quad (\text{A7})$$

$$\sum_{\theta} \int_0^{\mu(\theta)} y_j^i(\theta) di = \sum_{\theta} \int_{\boldsymbol{\xi}} \int_0^{\mu(\theta)} \tilde{x}_j^i(\boldsymbol{\xi}, \theta) q_j^i(\boldsymbol{\xi}, \theta) g(\boldsymbol{\xi}, \theta) di d\boldsymbol{\xi}, \forall j \quad (\text{A8})$$

then combine the outside good and differentiated good feasibility constraints evaluated at prices

$$\sum_{\theta} \int_0^{\mu(\theta)} y_o^i(\theta) di + \sum_{\theta} \int_0^{\mu(\theta)} p_j y_j^i(\theta) di = \quad (\text{A9})$$

$$\sum_{\theta} \int_{\xi} \int_0^{\mu(\theta)} c^i(\xi, \theta) g(\xi, \theta) di d\xi + \sum_j \sum_{\theta} \int_{\xi} \int_0^{\mu(\theta)} \tilde{x}_j^i(\xi, \theta) p_j q_j^i(\xi, \theta) g(\xi; \theta) di d\xi \quad (\text{A10})$$

and then rearrange everything so one can see it i by i so we have

$$\sum_{\theta} \int_0^{\mu(\theta)} W(\theta) di = \sum_{\theta} \int_{\xi} \int_0^{\mu(\theta)} \left[\sum_j x_j^i(\xi, \theta) [\tilde{c}^i(\xi, \theta) + p_j q_j^i(\xi, \theta)] \right] g(\xi; \theta) di d\xi \quad (\text{A11})$$

then make the observation that asset trades occur at actuarially fair prices and then change the order of integration giving

$$\sum_{\theta} \int_0^{\mu(\theta)} W(\theta) di = \sum_{\theta} \int_0^{\mu(\theta)} \int_{\xi} \varphi(\xi, \theta) \left[\sum_j x_j^i(\xi, \theta) [\tilde{c}^i(\xi, \theta) + p_j q_j^i(\xi, \theta)] \right] d\xi di. \quad (\text{A12})$$

so feasibility then implies that each households budget constraint is satisfied. But if the allocation is preferred than

$$\sum_{\theta} \int_0^{\mu(\theta)} W(\theta) di = \sum_{\theta} \int_0^{\mu(\theta)} \int_{\xi} \varphi(\xi, \theta) \left[\sum_j x_j^i(\xi, \theta) [\tilde{c}^i(\xi, \theta) + p_j q_j^i(\xi, \theta)] \right] d\xi di > \sum_{\theta} \int_0^{\mu(\theta)} W(\theta) di \quad (\text{A13})$$

where the right-hand side inequality follows from (A6) and we have our contradiction. Thus, an allocation that satisfies the definition of a complete markets equilibrium is a Pareto efficient allocation.

A.3. Solving the Households Problem

Here walk through the arguments regarding how we characterize the agents problem in complete markets. The households Lagrangian associated with the complete markets problem is

$$\mathcal{L} = \max_{c_j(\boldsymbol{\xi}), q_j(\boldsymbol{\xi}), x_j(\boldsymbol{\xi})} \int_{\boldsymbol{\xi}} \sum_j x_j^i(\boldsymbol{\xi}, \theta) \left\{ u[c^i(\boldsymbol{\xi}, \theta), q_j^i(\boldsymbol{\xi}, \theta); j, \theta] + \xi_j^i \right\} g(\boldsymbol{\xi}; \theta) d\boldsymbol{\xi}, \quad (\text{A14})$$

$$+ \lambda^i(\theta) \int_{\boldsymbol{\xi}} \varphi(\boldsymbol{\xi}, \theta) \left\{ W(\theta) - \sum_j x_j^i(\boldsymbol{\xi}, \theta) [c^i(\boldsymbol{\xi}, \theta) + p_j q_j^i(\boldsymbol{\xi}, \theta)] \right\} d\boldsymbol{\xi}. \quad (\text{A15})$$

The first order condition for consumption is

$$x_j^i(\boldsymbol{\xi}, \theta) u_{q_j} [c^i(\boldsymbol{\xi}, \theta), q_j^i(\boldsymbol{\xi}, \theta); j, \theta] g(\boldsymbol{\xi}, \theta) = \lambda^i(\theta) \varphi(\boldsymbol{\xi}, \theta) x_j^i(\boldsymbol{\xi}, \theta) p_j. \quad (\text{A16})$$

Now insert the fact that the state prices are actuarially fair and then canceling terms we have

$$u_{q_j} [c^i(\boldsymbol{\xi}, \theta), q_j^i(\boldsymbol{\xi}, \theta); j, \theta] = \lambda^i(\theta) p_j. \quad (\text{A17})$$

Then similarly for the non-differentiated good

$$x_j^i(\boldsymbol{\xi}, \theta) u_c [c^i(\boldsymbol{\xi}, \theta), q_j^i(\boldsymbol{\xi}, \theta); j, \theta] g(\boldsymbol{\xi}, \theta) = \lambda^i(\theta) \varphi(\boldsymbol{\xi}, \theta) x_j^i(\boldsymbol{\xi}, \theta), \quad (\text{A18})$$

and then with actuarially fair state prices and canceling terms we have

$$u_c [c^i(\boldsymbol{\xi}, \theta), q_j^i(\boldsymbol{\xi}, \theta); j, \theta] = \lambda^i(\theta). \quad (\text{A19})$$

Then the same arguments applied in the incomplete markets problem regarding how certain indexing can be dropped apply here as well. These arguments imply that the first order conditions characterizing consumption allocations are

$$u_c [c(j, \theta), q_j(\theta); j, \theta] = \lambda(\theta) \quad \text{and} \quad u_{q_j} [c(j, \theta), q_j(\theta); j, \theta] = \lambda(\theta) p_j. \quad (\text{A20})$$

Then this condition leads to a risk-sharing type result where marginal utilities are equated across *all* events — both explicitly as there is no dependence upon $\boldsymbol{\xi}$ and then implicitly as there is no dependence upon the choice.

We can follow the same arguments in the incomplete markets problem to characterize $x_j(\boldsymbol{\xi}, \theta)$.

So fix an event, then compare utility across those events in the Lagrangian...

$$[u[c^i(\boldsymbol{\xi}, \theta), q_1^i(\boldsymbol{\xi}, \theta); 1, \theta] + \xi_1^i]g(\boldsymbol{\xi}, \theta) + \lambda^i(\theta)\varphi(\boldsymbol{\xi}, \theta)[W(\theta) - c^i(\boldsymbol{\xi}, \theta) - p_1q_1^i(\boldsymbol{\xi}, \theta)] \quad \text{vs.} \quad (\text{A21})$$

$$[u[c^i(\boldsymbol{\xi}, \theta), q_2^i(\boldsymbol{\xi}, \theta); 2, \theta] + \xi_2^i]g(\boldsymbol{\xi}, \theta) + \lambda^i(\theta)\varphi(\boldsymbol{\xi}, \theta)[W(\theta) - c^i(\boldsymbol{\xi}, \theta) - p_2q_2^i(\boldsymbol{\xi}, \theta)], \dots \quad (\text{A22})$$

Then using the observation that state prices are actuarially fair means the comparison reduces to:

$$[u[c^i(\boldsymbol{\xi}, \theta), q_1^i(\boldsymbol{\xi}, \theta); 1, \theta] + \xi_1^i] - \lambda^i(\theta)[c^i(\boldsymbol{\xi}, \theta) + p_1q_1^i(\boldsymbol{\xi}, \theta)] \quad \text{vs.} \quad (\text{A23})$$

$$[u[c^i(\boldsymbol{\xi}, \theta), q_2^i(\boldsymbol{\xi}, \theta); 2, \theta] + \xi_2^i] - \lambda^i(\theta)[c^i(\boldsymbol{\xi}, \theta) + p_1q_2^i(\boldsymbol{\xi}, \theta)] \dots \quad (\text{A24})$$

At this point, there are still lots of difficulties, specifically how are the quantities varying with the realization of the shock. This problem is solved by inserting the observations about how consumption does not depend upon $\boldsymbol{\xi}$ which gives the following choice rule

$$x_j^i(\boldsymbol{\xi}, \theta) = \begin{cases} 1, & \text{if } u[c(j, \theta), q_j(\theta); j, \theta] + \xi_j^i - \lambda(\theta)[c(j, \theta) + p_jq_j(\theta)] \geq \\ & \max_{j'} \left\{ u[c(j', \theta), q_{j'}(\theta); j', \theta] + \xi_{j'}^i - \lambda(\theta)[c(j', \theta) + p_{j'}q_{j'}(\theta)] \right\} \\ 0. & \text{otherwise} \end{cases} \quad (\text{A25})$$

The key novelty are the $\lambda(\theta)$ term which is not in the rule for the standard allocation (equation 9). This additional term reflects the cost of choosing that commodity on the consolidated budget constraint.

The final property we consider are the asset positions. Interestingly, a households positions are only contingent the choice, not the particular shock realization $\boldsymbol{\xi}$. This follows from and the properties of consumption discussed above and the budget constraint, so

$$a(j, \theta) = W(\theta) - [c(j, \theta) + p_jq_j(\theta)]. \quad (\text{A26})$$

Finally, notice that one can insert (A26) into the choice rule giving a very simple representation

$$x_j^i(\boldsymbol{\xi}, \theta) = \begin{cases} 1, & \text{if } u[c(j, \theta), q_j(\theta); j, \theta] + \xi_j^i - \lambda(\theta)a(j, \theta) \geq \\ & \max_{j'} \left\{ u[c(j', \theta), q_{j'}(\theta); j', \theta] + \xi_{j'}^i - \lambda(\theta)a(j', \theta) \right\} \\ 0, & \text{otherwise} \end{cases} \quad (\text{A27})$$

which gives interpretation on the last terms as the cost of purchasing insurance associated with the choices j . If we impose the type 1 extreme value distributional assumption and integrate up we obtain the efficient choice probabilities:

$$\exp\left(\frac{u[c(j, \theta), q_j(\theta), j, \theta] - \lambda(\theta)a(j, \theta)}{\eta_\theta}\right) / \sum_{j'} \exp\left(\frac{u[c(j', \theta), q_{j'}(\theta), j', \theta] - \lambda(\theta)a(j', \theta)}{\eta_\theta}\right). \quad (\text{A28})$$

Below we summarize these results.

Proposition 6 (Complete Markets Allocations) *The following conditions characterize the complete markets allocations:*

1. *Consumption allocations must satisfy*

$$u_c[c(\theta, j), q_j(\theta); j, \theta] = \lambda(\theta) \quad \text{and} \quad u_{q_j}[c(\theta, j), q_j(\theta); j, \theta] = \lambda(\theta)p_j,$$

2. *The commodity choice rule is*

$$x_j^i(\boldsymbol{\xi}, \theta) = \begin{cases} 1, & \text{if } u[c(j, \theta), q_j(\theta); j, \theta] + \xi_j^i - \lambda(\theta)a(j, \theta) \geq \\ & \max_{j'} \left\{ u[c(j', \theta), q_{j'}(\theta); j', \theta] + \xi_{j'}^i - \lambda(\theta)a(j', \theta) \right\} \\ 0, & \text{otherwise} \end{cases}$$

3. *“Arrow Vouchers.” Asset positions are given by*

$$a(j, \theta) = W(\theta) - [c(j, \theta) + p_j q_j(\theta)],$$

which are contingent only on the choice j , not the taste shock $\boldsymbol{\xi}$. The state prices for the

Arrow Vouchers are the choice probabilities which with the Type 1 EV assumption are

$$\exp\left(\frac{u[c(j, \theta), q_j(\theta), j, \theta] - \lambda(\theta)a(j, \theta)}{\eta_\theta}\right) / \sum_{j'} \exp\left(\frac{u[c(j', \theta), q_{j'}(\theta), j', \theta] - \lambda(\theta)a(j', \theta)}{\eta_\theta}\right).$$

C. Appendix: The Pareto Problem

The Pareto problem is to find allocations that maximize somebody's utility subject to a resource constraint and that the allocation must respect given utility levels for all other households in the economy. The resulting allocation is then by definition of the problem a Pareto efficient allocation.

Utility for household i of type θ is

$$V^i(\theta) = \int_{\xi} \sum_j x_j^i(\xi, \theta) \left\{ u[c^i(\xi, \theta), q_j^i(\xi, \theta); j, \theta] + \xi_j^i \right\} g(\xi; \theta) d\xi, \quad (C1)$$

and then we will index all other households by the label k, θ' with given utility level $V^k(\theta')$. Then the Pareto problem is

$$\max_{c_j(\xi, \theta), q_j(\xi, \theta), x_j(\xi, \theta) \forall i, k, \theta'} \int_{\xi} \sum_j x_j^i(\xi, \theta) \left\{ u[c^i(\xi, \theta), q_j^i(\xi, \theta); j, \theta] + \xi_j^i \right\} g(\xi; \theta) d\xi \quad (C2)$$

subject to the resource constraints:

$$[\Lambda_o] \quad \sum_{\theta'} \int_0^{\mu(\theta')} y_o^k(\theta') dk \geq \sum_{\theta'} \int_{\xi} \int_k x_j^k(\xi, \theta') c_j^k(\xi, \theta') dk g(\xi, \theta') d\xi, \quad (C3)$$

$$[\Lambda_j] \quad \sum_{\theta'} \int_0^{\mu(\theta')} y_j^k(\theta') dk \geq \sum_{\theta'} \int_{\xi} \int_k x_j^k(\xi, \theta') q_j^k(\xi, \theta') dk g(\xi, \theta') d\xi \quad \forall j. \quad (C4)$$

which says that goods supply must be greater than or equal to goods demand. Associated with these constraints are the Lagrange multiplier Λ_o and Λ_j for each good j . The next constraint is the Pareto constraint:

$$[\Upsilon^k(\theta')] \quad V^k(\theta') \leq \int_{\xi} \sum_j x_j^k(\xi, \theta') \left\{ u[c^k(\xi, \theta'), q_j^k(\xi, \theta'); j, \theta'] + \xi_j^k \right\} g(\xi, \theta') d\xi \quad \forall k, \theta' \neq i, \theta. \quad (C5)$$

This says that at any allocation, it has to deliver utility level $V^k(\theta')$ (or be better) for every k, θ' household. Associated with each of these constraints is the Lagrange multiplier $\Upsilon^k(\theta')$. Putting

the problem all together we have

$$\mathcal{L} = \max_{c_j(\boldsymbol{\xi}, \theta), q_j(\boldsymbol{\xi}, \theta), x_j(\boldsymbol{\xi}, \theta) \forall i, k, \theta'} \int_{\boldsymbol{\xi}} \sum_j x_j^i(\boldsymbol{\xi}, \theta) \left\{ u \left[c^i(\boldsymbol{\xi}, \theta), q_j^i(\boldsymbol{\xi}, \theta); j, \theta \right] + \xi_j^i \right\} g(\boldsymbol{\xi}; \theta) d\boldsymbol{\xi} \quad (\text{C6})$$

$$+ \Lambda_o \left[\sum_{\theta'} Y_o(\theta') - \sum_{\theta'} \int_{\boldsymbol{\xi}} \int_k x_j^k(\boldsymbol{\xi}, \theta') c_j^k(\boldsymbol{\xi}, \theta') dk g(\boldsymbol{\xi}, \theta') d\boldsymbol{\xi} \right] \quad (\text{C7})$$

$$+ \sum_j \Lambda_j \left[\sum_{\theta'} Y_j(\theta') - \sum_{\theta'} \int_{\boldsymbol{\xi}} \int_k x_j^k(\boldsymbol{\xi}, \theta') q_j^k(\boldsymbol{\xi}, \theta') dk g(\boldsymbol{\xi}, \theta') d\boldsymbol{\xi} \right] \quad (\text{C8})$$

$$+ \sum_{\theta'} \int_k \Upsilon^k(\theta') \left[\int_{\boldsymbol{\xi}} \sum_j x_j^k(\boldsymbol{\xi}, \theta') \left\{ u \left[c^k(\boldsymbol{\xi}, \theta'), q_j^k(\boldsymbol{\xi}, \theta'); j, \theta' \right] + \xi_j^k \right\} g(\boldsymbol{\xi}, \theta') d\boldsymbol{\xi} - V^k(\theta') \right] dk g(\boldsymbol{\xi}) d\boldsymbol{\xi}. \quad (\text{C9})$$

As a recap: This planner chooses allocations for everybody: i, θ and all k, θ' to maximize welfare for i, θ given the resource constraint and then the idea that all k, θ' households must be delivered at least $V^k(\theta')$ level of utility.

The next steps derive first order conditions and characterize the rule describing which goods are chosen and under which circumstances. We do this below in steps.

Conditional on a choice, the first order condition for i 's consumption of variety q_j is

$$x_j^i(\boldsymbol{\xi}, \theta) u_{q_j} \left[c^i(\boldsymbol{\xi}, \theta), q_j^i(\boldsymbol{\xi}, \theta); j, \theta \right] g(\boldsymbol{\xi}, \theta) = \Lambda_j x_j^i(\boldsymbol{\xi}, \theta) g(\boldsymbol{\xi}, \theta) \quad (\text{C10})$$

and then canceling terms we have

$$u_{q_j} \left[c^i(\boldsymbol{\xi}, \theta), q_j^i(\boldsymbol{\xi}, \theta); j, \theta \right] = \Lambda_j. \quad (\text{C11})$$

And then a similar condition holds for the non-differentiated commodity

$$x_j^i(\boldsymbol{\xi}, \theta) u_c \left[c^i(\boldsymbol{\xi}, \theta), q_j^i(\boldsymbol{\xi}, \theta); j, \theta \right] g(\boldsymbol{\xi}, \theta) = \Lambda_o x_j^i(\boldsymbol{\xi}, \theta) g(\boldsymbol{\xi}, \theta), \quad (\text{C12})$$

and then canceling terms we have

$$u_c \left[c^i(\boldsymbol{\xi}, \theta), q_j^i(\boldsymbol{\xi}, \theta); j, \theta \right] = \Lambda_o. \quad (\text{C13})$$

Then we make the observation that the quantities don't depend upon the particular shock real-

ization ξ since the multipliers don't depend on the shock giving

$$u_c [c^i(j, \theta), q_j^i(\theta); j, \theta] = \Lambda_o. \quad (\text{C14})$$

The next question is then what about k, θ' 's consumption. For variety j we have

$$\Upsilon^k(\theta') x_j^k(\xi, \theta') u_{q_j} [c^k(\xi, \theta'), q_j^k(\xi, \theta'); j, \theta'] g(\xi, \theta') = \Lambda_j x_j^k(\xi, \theta') g(\xi, \theta'), \quad (\text{C15})$$

and then canceling terms gives

$$u_{q_j} [c^k(\xi, \theta), q_j^k(\xi, \theta); j, \theta'] = \frac{\Lambda_j}{\Upsilon^k(\theta')}. \quad (\text{C16})$$

then re-indexing with the dropping of the choice we have

$$u_{q_j} [c^k(j, \theta), q_j^k(\theta); j, \theta'] = \frac{\Lambda_j}{\Upsilon^k(\theta')}. \quad (\text{C17})$$

And then we have the similar condition for the non-differentiated commodity

$$u_c [c^k(j, \theta'), q_j^k(\theta'); j, \theta'] = \frac{\Lambda_o}{\Upsilon^k(\theta')}. \quad (\text{C18})$$

Similar to above, quantities for the k, θ' guys are set so that marginal utility equals the multiplier on the resource constraint adjusted by the multiplier on the Pareto constraint. The adjustment for the multiplier on the Pareto constraint then has the interpretation as the weight that the planner places on agent k, θ' .

From here we can take ratios of these conditions and arrive at results that any Pareto efficient allocation must satisfy. Specifically

$$\frac{u_{q_j} [c^k(j, \theta), q_j^k(\theta); j, \theta']}{u_{q_{j'}} [c^k(j, \theta), q_{j'}^k(\theta); j', \theta']} = \frac{\Lambda_j}{\Lambda_{j'}}. \quad (\text{C19})$$

and this condition holds for any k, θ' agent (including i, θ), any shock realization, and any ratio of commodity j, j' or o . This says that the ratio of marginal rate of substitution between these commodities, for any person, must equal the shadow cost of those commodities.

The next step is to characterize the commodity choice rule or the $x_j^i(\xi, \theta)$ s. Like in the cases above, the variational approach can be applied as well. So we just inspect the Lagrangian,

choice by choice for the i, θ guy under a particular realization of the shock

$$[u[c^i(\boldsymbol{\xi}, \theta), q_1^i(\boldsymbol{\xi}, \theta); 1, \theta] + \xi_1^i]g(\boldsymbol{\xi}, \theta) - \left[\Lambda_o c^i(\boldsymbol{\xi}, \theta) + \Lambda_1 q_1^i(\boldsymbol{\xi}, \theta) \right]g(\boldsymbol{\xi}, \theta) \quad \text{vs.} \quad (\text{C20})$$

$$[u[c^i(\boldsymbol{\xi}, \theta), q_2^i(\boldsymbol{\xi}, \theta); 2, \theta] + \xi_2^i]g(\boldsymbol{\xi}, \theta) - \left[\Lambda_o c^i(\boldsymbol{\xi}, \theta) + \Lambda_2 q_2^i(\boldsymbol{\xi}, \theta) \right]g(\boldsymbol{\xi}, \theta), \dots \quad (\text{C21})$$

Then canceling the densities and inserting the result that the consumption does not depend upon $\boldsymbol{\xi}$

$$x_j^i(\boldsymbol{\xi}, \theta) = \begin{cases} 1, & \text{if } u[c^i(j, \theta), q_j^i(\theta); j, \theta] + \xi_j^i - \left[\Lambda_o c^i(j, \theta) + \Lambda_j q_j^i(\theta) \right] \geq \\ & \max_{j'} \left\{ u[c^i(j', \theta), q_{j'}^i(\theta); j', \theta] + \xi_{j'}^i - \left[\Lambda_o c^i(j', \theta) + \Lambda_{j'} q_{j'}^i(\theta) \right] \right\} \\ 0, & \text{otherwise} \end{cases} \quad (\text{C22})$$

Now the interesting case is the k, θ' guy. Again, use the same variational argument with

$$\Upsilon^k(\theta') [u[c^k(\boldsymbol{\xi}, \theta'), q_1^k(\boldsymbol{\xi}, \theta'); 1, \theta'] + \xi_1^k]g(\boldsymbol{\xi}, \theta') - \left[\Lambda_o c^k(\boldsymbol{\xi}, \theta') + \Lambda_1 q_1^k(\boldsymbol{\xi}, \theta') \right]g(\boldsymbol{\xi}, \theta') \quad \text{vs.} \quad (\text{C23})$$

$$\Upsilon^k(\theta') [u[c^k(\boldsymbol{\xi}, \theta'), q_2^k(\boldsymbol{\xi}, \theta'); 2, \theta'] + \xi_2^k]g(\boldsymbol{\xi}, \theta') - \left[\Lambda_o c^k(\boldsymbol{\xi}, \theta') + \Lambda_2 q_2^k(\boldsymbol{\xi}, \theta') \right]g(\boldsymbol{\xi}, \theta') \dots \quad (\text{C24})$$

Then we follow the same steps as with the i case. Terms cancel, insert the result that the consumption allocation does not depend upon the shock, then divide through by the multiplier $\Upsilon^k(\theta')$. This then gives rise to the following commodity choice rule:

$$x_j^k(\boldsymbol{\xi}, \theta') = \begin{cases} 1, & \text{if } u[c^k(j, \theta'), q_j^k(\theta'); j, \theta'] + \xi_j^k - \frac{1}{\Upsilon^k(\theta')} \left[\Lambda_o c^k(j, \theta') + \Lambda_j q_j^k(\theta') \right] \geq \\ & \max_{j'} \left\{ u[c^k(j', \theta'), q_{j'}^k(\theta'); j', \theta'] + \xi_{j'}^k - \frac{1}{\Upsilon^k(\theta')} \left[\Lambda_o c^k(j', \theta') + \Lambda_{j'} q_{j'}^k(\theta') \right] \right\} \\ 0, & \text{otherwise} \end{cases} \quad (\text{C25})$$

which holds for all k, θ' households. And note that this rule is equivalent to the i, θ household by setting the Υ term to one. An interesting step is to notice how the ratio of the multipliers then connects with the marginal utility of consumption giving the same exact choice rule as

with the i guy

$$x_j^k(\boldsymbol{\xi}, \theta') = \begin{cases} 1, & \text{if } u[c^k(j, \theta'), q_j^k(\theta'); j, \theta'] + \xi_j^k - \left[u_c^k(j, \theta') c^k(j, \theta') + u_{q_j}^k(\theta') q_j^k(\theta') \right] \geq \\ & \max_{j'} \left\{ u[c^k(j', \theta'), q_{j'}^k(\theta'); j', \theta'] + \xi_{j'}^k - \left[u_c^k(j, \theta') c^k(j', \theta') + u_{q_{j'}}^k(\theta') q_{j'}^k(\theta') \right] \right\} \\ 0, & \text{otherwise} \end{cases} \quad (\text{C26})$$

And this rule then looks exactly the same as for the i, θ agent. One interesting feature of this characterization is that the welfare weights (or the $\Upsilon^k(\theta')$) don't directly show up. What is going on is that welfare weights determine how much or little consumption k, θ' person receives, and thus their marginal utility for the commodity. Then these marginal utilities appropriately determines the social cost of choosing a particular variety.

Proposition 7 summarizes the result below.

Proposition 7 (Pareto Efficient Allocations) *Given utility levels $V^k(\theta')$ for all $k, \theta' \neq i, \theta$, a Pareto efficient allocation is consumption allocations and commodity choice rules $c_j^i(\boldsymbol{\xi}, \theta)$, $q_j^i(\boldsymbol{\xi}, \theta)$, $x_j^i(\boldsymbol{\xi}, \theta)$ and i, θ and for all other all k, θ' , that solve the problem (24) subject to resource constraints (25, 26) and the Pareto constraint in (27).*

The following conditions characterize Pareto efficient allocations:

1. For agent i, θ , consumption allocations must satisfy

$$u_c[c^i(j, \theta), q_j^i(\theta); j, \theta] = \Lambda_o \quad \text{and} \quad u_{q_j}[c^i(j, \theta), q_j^i(\theta); j, \theta] = \Lambda_j.$$

2. For agent k', θ' , consumption allocations must satisfy:

$$\Upsilon^k(\theta') u_c[c^k(j, \theta), q_j^k(\theta); j, \theta] = \Lambda_o \quad \text{and} \quad \Upsilon^k(\theta') u_{q_j}[c^k(j, \theta), q_j^k(\theta); j, \theta] = \Lambda_j.$$

3. The commodity choice rule is

$$x_j^k(\boldsymbol{\xi}, \theta') = \begin{cases} 1, & \text{if } u[c^k(j, \theta'), q_j^k(\theta'); j, \theta'] + \xi_j^k - \frac{1}{\Upsilon^k(\theta')} \left[\Lambda_o c^k(j, \theta') + \Lambda_j q_j^k(\theta') \right] \geq \\ & \max_{j'} \left\{ u[c^k(j', \theta'), q_{j'}^k(\theta'); j', \theta'] + \xi_{j'}^k - \frac{1}{\Upsilon^k(\theta')} \left[\Lambda_o c^k(j', \theta') + \Lambda_j q_{j'}^k(\theta') \right] \right\} \\ 0, & \text{otherwise} \end{cases}$$

which holds for all k, θ' households, and for the i, θ household set the Υ term to one.

We now have a basis for the *Second Welfare Theorem*. This is immediate from comparing the contents of Proposition 1 and from Complete Markets to Proposition 3. Consumption allocations and choice rules are identical if $\Lambda_0/\Upsilon^k(\theta) = \lambda(\theta)$, and $\Lambda_j/\Upsilon^k(\theta) = \lambda(\theta)p_j$. To align the complete markets allocation to the Pareto efficient allocation, then one simply needs to find the appropriate ex-ante reallocation of resources to attain the correct multipliers on complete markets budget constraints, $\lambda(\theta)$.

D. Appendix: The Social Planning Problem

In this section we solve the problem of a planner who maximizes a linear social welfare function with welfare weights $\psi^k(\theta')$ for all k and θ' . Inspection of the conditions of this problem then provides an equivalence between Pareto efficient allocations and maximizing linear social welfare functions under some welfare weights.

The planning problem is the following

$$\max_{c^k(\xi, \theta'), q_j^k(\xi, \theta'), x_j^k(\xi, \theta')} \sum_{\theta'} \int_k \psi^k(\theta') \int_{\xi} \sum_j x_j^k(\xi, \theta') \left\{ u[c^k(\xi, \theta'), q_j^k(\xi, \theta'); j, \theta'] + \xi_j^k \right\} g(\xi, \theta') d\xi dk, \quad (\text{D1})$$

$$\text{subject to: } \sum_{\theta'} Y_j(\theta') \geq \sum_{\theta'} \int_{\xi} \int_k x_j^k(\xi, \theta') q_j^k(\xi, \theta') dk g(\xi, \theta') d\xi \quad \forall j \quad (\text{D2})$$

$$\sum_{\theta'} Y_o(\theta') \geq \sum_{\theta'} \int_{\xi} \int_k \sum_j x_j^k(\xi, \theta') c^k(\xi, \theta') dk g(\xi, \theta') d\xi \quad (\text{D3})$$

Where the social planner chooses consumption and commodities for every event ξ and every k, θ' agent in the economy. These allocations are chosen to maximize social welfare in (D1) which is a weighted average of each individual agents expected utility, with the weights being the given social welfare weights $\psi^k(\theta')$. This objective function is maximized subject to the resource constraints.

The Lagrangian associated with this problem is

$$\mathcal{L} = \max_{c^k(\boldsymbol{\xi}, \theta'), q_j^k(\boldsymbol{\xi}, \theta'), x_j^k(\boldsymbol{\xi}, \theta')} \sum_{\theta'} \int_k \psi^k(\theta') \int_{\boldsymbol{\xi}} \sum_j x_j^k(\boldsymbol{\xi}, \theta') \left\{ u[c^k(\boldsymbol{\xi}, \theta'), q_j^k(\boldsymbol{\xi}, \theta'); j, \theta'] + \xi_j^k \right\} dk g(\boldsymbol{\xi}, \theta') d\boldsymbol{\xi}, \quad (\text{D4})$$

$$+ \sum_j \Lambda_j \left[\sum_{\theta'} Y_j(\theta') - \sum_{\theta'} \int_{\boldsymbol{\xi}} \int_k x_j^k(\boldsymbol{\xi}, \theta') q_j^k(\boldsymbol{\xi}, \theta') dk g(\boldsymbol{\xi}, \theta') d\boldsymbol{\xi} \right] \quad (\text{D5})$$

$$+ \Lambda_o \left[\sum_{\theta'} Y_o(\theta') - \sum_{\theta'} \int_{\boldsymbol{\xi}} \int_k \sum_j x_j^k(\boldsymbol{\xi}, \theta') c^k(\boldsymbol{\xi}, \theta') dk g(\boldsymbol{\xi}, \theta') d\boldsymbol{\xi} \right]. \quad (\text{D6})$$

The first order condition for consumption of commodity j becomes (after canceling terms using the same arguments above)

$$\psi^k(\theta') u_{q_j} [c^k(j, \theta'), q_j^k(\theta'); j, \theta] = \Lambda_j, \quad (\text{D7})$$

and then for the non-differentiated commodity we have

$$\psi^k(\theta') u_c [c^k(j, \theta'), q_j^k(\theta'); j, \theta] = \Lambda_o. \quad (\text{D8})$$

The planner sets social-welfare-weighted marginal utility equal to the shadow cost of consuming that commodity. Inspecting these conditions (D7, D8) and (C11, C13) shows that for a given Pareto efficient allocation, the multipliers on the Pareto constraint maps directly into the social welfare weights associated with Planning problem.

Now social welfare maximizing choice rule $x_j^k(\boldsymbol{\xi}, \theta')$ is

$$\psi^k(\theta') [u[c^k(\boldsymbol{\xi}, \theta'), q_1^k(\boldsymbol{\xi}, \theta'); 1, \theta'] + \xi_1^k] g(\boldsymbol{\xi}, \theta') - \left[\Lambda_o c^k(\boldsymbol{\xi}, \theta') + \Lambda_1 q_1^k(\boldsymbol{\xi}, \theta') \right] g(\boldsymbol{\xi}, \theta') \quad \text{vs.} \quad (\text{D9})$$

$$\psi^k(\theta') [u[c^k(\boldsymbol{\xi}, \theta'), q_2^k(\boldsymbol{\xi}, \theta'); 2, \theta'] + \xi_2^k] g(\boldsymbol{\xi}, \theta') - \left[\Lambda_o c^k(\boldsymbol{\xi}, \theta') + \Lambda_2 q_2^k(\boldsymbol{\xi}, \theta') \right] g(\boldsymbol{\xi}, \theta') \dots \quad (\text{D10})$$

Then terms cancel, insert the result that the consumption allocation does not depend upon the

shock, then divide through by the social welfare weight $\psi^k(\theta')$. The choice rule

$$x_j^k(\boldsymbol{\xi}, \theta') = \begin{cases} 1, & \text{if } u[c^k(j, \theta'), q_j^k(\theta'); j, \theta'] + \xi_j^k - \frac{1}{\psi^k(\theta')} \left[\Lambda_o c^k(j, \theta') + \Lambda_j q_j^k(\theta') \right] \geq \\ & \max_{j'} \left\{ u[c^k(j', \theta'), q_{j'}^k(\theta'); j', \theta'] + \xi_{j'}^k - \frac{1}{\psi^k(\theta')} \left[\Lambda_o c^k(j', \theta') + \Lambda_{j'} q_{j'}^k(\theta') \right] \right\} \\ 0, & \text{otherwise} \end{cases} \quad (\text{D11})$$

And the key observation here is that this choice rule is exactly the same as that in the Pareto problem, and thus the complete markets allocation. Proposition 8 summarizes the result below.

Proposition 8 (Social Welfare Maximizing Allocations) *Let $\psi^k(\theta')$ be a vector of Social Welfare Weights. Define the Social Welfare Function as:*

$$\mathcal{W}_\psi = \sum_{\theta'} \int_k \psi^k(\theta') \int_{\boldsymbol{\xi}} \sum_j x_j^k(\boldsymbol{\xi}, \theta') \left\{ u[c^k(\boldsymbol{\xi}, \theta'), q_j^k(\boldsymbol{\xi}, \theta'); j, \theta'] + \xi_j^k \right\} g(\boldsymbol{\xi}, \theta') d\boldsymbol{\xi} di \quad (\text{D12})$$

Then a **Social Welfare Maximizing allocation** is consumption allocations and commodity choice rules and $c_j^k(\boldsymbol{\xi}, \theta')$, $q_j^k(\boldsymbol{\xi}, \theta')$, $x_j^k(\boldsymbol{\xi}, \theta')$ for all k, θ' to maximize \mathcal{W}_ψ subject to resource constraints on all goods. The following conditions characterize the allocation:

1. For all k, θ' , consumption allocations must satisfy:

$$\psi^k(\theta') u_c[c^k(\theta'), q_j^k(\theta'); j, \theta'] = \Lambda_o \quad \text{and} \quad \psi^k(\theta') u_{q_j}[c^k(\theta'), q_j^k(\theta'); j, \theta'] = \Lambda_j. \quad (\text{D13})$$

2. For all k, θ' the commodity choice rule is

$$x_j^k(\boldsymbol{\xi}, \theta') = \begin{cases} 1, & \text{if } u[c^k(j, \theta'), q_j^k(\theta'); j, \theta'] + \xi_j^k - \frac{1}{\psi^k(\theta')} \left[\Lambda_o c^k(j, \theta') + \Lambda_j q_j^k(\theta') \right] \geq \\ & \max_{j'} \left\{ u[c^k(j', \theta'), q_{j'}^k(\theta'); j', \theta'] + \xi_{j'}^k - \frac{1}{\psi^k(\theta')} \left[\Lambda_o c^k(j', \theta') + \Lambda_{j'} q_{j'}^k(\theta') \right] \right\} \\ 0, & \text{otherwise} \end{cases} \quad (\text{D14})$$

This allocation is a Pareto Efficient Allocation and coincides with a Complete Markets Allocation under some ex-ante transfers.

E. Appendix: Multiplicative Case

In the main text we consider the canonical case of the linear additive random utility model (LRUM). In its simplest case—without an outside good, type differences in utility, or heterogeneous fixed characteristics of goods—utility from a choice j is given by:

$$u(q_j^i(\boldsymbol{\xi})) + \xi_j^i$$

Here we consider the multiplicative case:

$$u(q_j^i(\boldsymbol{\xi}) \xi_j^i)$$

Summary of results. Our main results are:

1. **Equivalences** - The core results in the main text are (i) that the incomplete markets allocation is inefficient, and that efficiency is obtained with complete markets, (ii) the efficient discrete choice policy incorporates the private / social marginal cost of the choice, (iii) the efficient consumption policy equates marginal rates of substitution across choices to relative prices. All of these features also hold in the multiplicative case.
2. **Arrow vouchers** - The only key departure is that *Arrow Vouchers*—which pay off only conditional on choices—are insufficient to decentralize the efficient allocation. Outside of $u(q) = \log q$, consumption $q_j^i(\boldsymbol{\xi})$ depends on the realization of the shock.
3. **Aggregation to CES** - In the special case of u being power utility, and ξ_j^i being Frechet distributed, aggregation yields an aggregate demand system that is equivalent to one that would be obtained by a representative household maximizing a CES utility function and consuming all goods. This is true under both complete and incomplete markets.
4. **CES utility** - However, the preferences of the representative CES consumer are different depending on market structure. Under incomplete markets the elasticity of substitution of the representative agent depends on preference dispersion. Under complete markets the elasticity of substitution is one.
5. **Efficiency \neq CES Rep. Consumer** - This provides a strong result that there is no equivalence between (a) efficiency, (b) the existence of the representative CES consumer.

We establish these under a specialized case where (i) utility functions, distributions of tastes and endowments are the same across types, θ , (ii) utility only depends on the choice of the quantity of the discrete good $q_j^i(\boldsymbol{\xi})$ and not on other characteristics of j , (iii) there is no outside

good. That is, expected utility as in equation (3) in the main text becomes:

$$V = \int_{\xi} \sum_j x_j(\xi) u(q_j(\xi) \xi_j) g(\xi) d\xi$$

The conclusion that the incomplete markets and complete markets allocations are, respectively, inefficient and efficient, does not require any functional form assumptions on u or distributional assumptions of g . We use the case of power utility, and draws of ξ_j^i from independent Fréchet distributions to derive our closed form aggregation results (points 3, 4, 5, above). That is, u is:

$$u(q\xi) = \frac{(q\xi)^{1-\sigma}}{1-\sigma}.$$

Note that in the limit as $\sigma \rightarrow 1$, this reduces to additive log preferences: $\log q + \log \xi$. With Fréchet distributed idiosyncratic preferences, this is identical to the case studied by ADT. In this special case the results from the main text apply, that is the incomplete markets allocation is efficient. Away from $\sigma = 1$, we show that aggregation to a CES representative agent is possible with incomplete markets and complete markets, but the former is *inefficient*.

E.1. Incomplete markets

The standard allocation that obtains under incomplete markets is solved as in the text

$$\max_{q_j(\xi), x_j(\xi)} \int_{\xi} \sum_j x_j(\xi) \{u(\xi_j q_j(\xi)) g(\xi)\} d\xi$$

subject to a budget constraint for each vector of preferences ξ :

$$[\lambda(\xi)] : \sum_j x_j(\xi) [p_j q_j(\xi) - W] \leq 0, \quad \text{for all } \xi$$

where $W = \sum_j p_j Y_j$ is the value of the individual's endowment. Since all individuals are identical and there are a measure one of them, the individual endowment y_j equals the aggregate endowment Y_j .

We proceed as in the main text to solve this problem. The individuals' problem has the following optimality conditions. First, the consumption allocation must satisfy

$$\begin{aligned} u'(q_j(\xi) \xi_j) \xi_j &= \frac{\lambda(\xi)}{g(\xi)} p_j \\ q_j(\xi) &= \frac{W}{p_j} \end{aligned}$$

If we write down the Lagrangian, fix a ξ , and then compare terms featuring $x_j(\xi)$, we then

obtain the optimal $x_j(\boldsymbol{\xi})$ follows:

$$x_j(\boldsymbol{\xi}) = \begin{cases} 1 & , \text{ if } u(q_j(\boldsymbol{\xi}) \xi_j) - \lambda(\boldsymbol{\xi}) [p_j q_j(\boldsymbol{\xi}) - W] \geq \max_{j'} \{u(q_{j'}(\boldsymbol{\xi}) \xi_{j'}) - \lambda(\boldsymbol{\xi}) [p_{j'} q_{j'}(\boldsymbol{\xi}) - W]\} \\ 0 & , \text{ otherwise} \end{cases}$$

Since the budget constraint binds for each choice, then the term in $[\cdot]$ is zero regardless of the j chosen, and hence we have

$$x_j(\boldsymbol{\xi}) = \begin{cases} 1 & , \text{ if } u(q_j(\boldsymbol{\xi}) \xi_j) \geq \max_{j'} \{u(q_{j'}(\boldsymbol{\xi}) \xi_{j'})\} \\ 0 & , \text{ otherwise} \end{cases}$$

E.1.1. Special case - Power utility - Frechet

The level of consumption $q_j(\boldsymbol{\xi}) = q_j$ is independent of the level of ξ_j and determined by the budget constraint. The *incomplete markets choice probability* then obtains from integrating the choice rule

$$\begin{aligned} x_j(\boldsymbol{\xi}) &= \arg \max_j \left\{ \frac{1}{1-\sigma} (q_j \xi_j)^{1-\sigma} \right\} \\ x_j(\boldsymbol{\xi}) &= \arg \max_j \left\{ \frac{1}{1-\sigma} \left(\frac{W}{p_j} \xi_j \right)^{1-\sigma} \right\} \\ x_j(\boldsymbol{\xi}) &= \arg \max_j \{ -\log p_j + \log \xi_j \} \end{aligned}$$

Under the additional assumption that $\xi_j \sim F(\xi_j)$ is Frechet with shape parameter η , we obtain the choice probabilities

$$\begin{aligned} \rho_j &= \frac{\exp \{ \eta \log p_j \}}{\sum_m \exp \{ \eta \log p_k \}} \\ \rho_j &= \frac{p_j^\eta}{\sum_k p_k^\eta} \end{aligned}$$

Aggregation. Aggregation follows the same steps as in ADT. Aggregation is simple since $q_j = W/p_j$ is independent of ξ :

$$Q_j = \rho_j q_j = \left(\frac{p_j^\eta}{\sum_k p_k^\eta} \right) \times \left(\frac{W}{p_j} \right)$$

Define a price index $P = [\sum_k p_k^{-\eta}]^{-\frac{1}{\eta}}$, which then gives:

$$Q_j = \frac{p_j^{-\eta}}{P^{-\eta}} \times \frac{W}{p_j}$$

Define Q by $PQ = W$. Then we have

$$Q_j = \left(\frac{p_j}{P}\right)^{-(1+\eta)} Q.$$

Working with this we can solve for Q :

$$Q = \left[\sum_j Q_j^{\frac{\eta}{1+\eta}} \right]^{\frac{1+\eta}{\eta}}$$

Representative agent. It is clear from this that a *CES representative agent* exists, under which the demand curves are computed under

$$\max_{Q_j} \frac{Q^{1-\sigma}}{1-\sigma}, \quad Q = \left[\sum_j Q_j^{\frac{\eta}{1+\eta}} \right]^{\frac{1+\eta}{\eta}} \quad \text{subject to} \quad \sum_j p_j q_j = W.$$

Relationship to ADT Note that the demand system is identical to that which is obtained under $\sigma = 1$. That is, other than being isoelastic, the functional form of u does not enter the demand curves, and demand is obtained *as if* the individual has $u(q\xi) = \log(q\xi)$. Given our prior results, it may be tempting to conclude that the equilibrium allocation is efficient, however below we show that this is not the case.

Efficiency? Generically, the incomplete markets allocation is *inefficient*. We will observe this easily below, by noting that—outside of the case of $\sigma = 1$ —with complete markets $q_j(\xi_j)$ depends on ξ_j , whereas this is not the case with incomplete markets. Nonetheless a representative agent exists with CES utility. This provides a counter-example to claims such as the following: “the efficiency property holds for any idiosyncratic preference structure that aggregates to yield CES preferences that satisfy the independence of irrelevant alternatives”. This is an example for which all of these hold, yet the allocation is inefficient.

E.2. Complete markets

The complete markets allocation is solved as in the text. The individual maximizes expected utility, V :

$$V = \max_{q_j(\boldsymbol{\xi}), x_j(\boldsymbol{\xi})} \int_{\boldsymbol{\xi}} \sum_j x_j(\boldsymbol{\xi}) \{u(q_j(\boldsymbol{\xi}) \xi_j) g(\boldsymbol{\xi})\} d\boldsymbol{\xi}$$

subject to a consolidated budget constraint that aggregates over all preference draws:

$$[\Lambda] : \int_{\boldsymbol{\xi}} \left[\sum_j x_j(\boldsymbol{\xi}) p_j q_j(\boldsymbol{\xi}) \right] \varphi(\boldsymbol{\xi}) d\boldsymbol{\xi} \leq W.$$

The price of securities $\varphi(\boldsymbol{\xi})$ will equal $g(\boldsymbol{\xi})$ in equilibrium, and hence we make that substitution immediately.

The complete market problem has the following optimality conditions. First, the consumption allocation satisfies the following first order condition:

$$u'(q_j(\boldsymbol{\xi}) \xi_j) \xi_j = \Lambda p_j.$$

This implies that $q_j(\boldsymbol{\xi})$ depends on $\boldsymbol{\xi}$ only through ξ_j . Note that under log utility $q_j(\xi_j)$ is independent of ξ_j . This is clearly necessary for the complete markets and incomplete markets allocations to coincide, which we should expect under log utility given our results on the efficiency of the allocation in ADT and the fact that when u is log, the ADT economy is obtained.

Second, the choice rule for a given draw $\boldsymbol{\xi}$ satisfies the variational inequality for $x_j(\boldsymbol{\xi})$:

$$x_j(\boldsymbol{\xi}) = \begin{cases} 1 & , \text{ if } u(q_j(\xi_j) \xi_j) - \Lambda p_j q_j(\xi_j) \geq \max_{j'} \{u(q_{j'}(\xi_{j'}) \xi_{j'}) - \Lambda p_{j'} q_{j'}(\xi_{j'})\} \\ 0 & , \text{ otherwise} \end{cases}.$$

The complete markets allocation is efficient. This can be observed readily applying the proofs for efficiency that we apply in the main text.

Comparing the complete markets allocation to the incomplete markets allocation, we observe that the individual achieves two additional outcomes with complete markets. First, *within a good* j , with incomplete markets consumption $q_j(\boldsymbol{\xi}) = W/p_j$ is *independent of* $\boldsymbol{\xi}$. With complete markets, the individual is able to move resources to equate marginal utilities of consumption $u'(q_j(\boldsymbol{\xi}) \xi_j) \xi_j$ across all $\boldsymbol{\xi}$ for which the individual chooses j . That is $q_j(\xi_j)$ depends on ξ_j . Note that this implies that *Arrow Vouchers* are no longer applicable in the multiplicative case, as payoffs of contracts require full indexation by individual preference shocks. Second, *across goods* j , as in the case of additive random utility, the individual is able to equalize the price adjusted marginal utility of consumption $u'(q_j(\xi_j)) \xi_j/p_j$. Without further functional form assumptions,

the optimality condition for $q_j(\boldsymbol{\xi})$ can be used in the choice rule:

$$\max_j \{u(q_j(\xi_j)\xi_j) - u'(q_j(\xi_j)\xi_j)\xi_j q_j(\xi_j)\}.$$

Note that this has the same properties as the efficient / complete markets choice rule in the additive random utility model studied in the main text. The term in $\{\cdot\}$ is the utility of consuming good j , minus the private marginal cost which is due to reducing resources in the budget constraint and equal to the multiplier Λ times expenditure.

E.2.1. Special case - Power utility - Frechet

Our key result here is that *aggregation to a Representative Agent CES* formulation is possible, in that it can be solved to deliver the demands $Q_j(p_1, \dots, p_j, \dots, p_J)$. However it is clear that the allocation that is obtained does not give the full details of the efficient allocation, since $q_j(\xi_j)$ depends also on ξ_j . This is readily observable from the optimality condition for consumption under power utility:

$$q_j(\boldsymbol{\xi}) = q_j(\xi_j) = \left(\frac{1}{\Lambda}\right)^{\frac{1}{\sigma}} p_j^{-\frac{1}{\sigma}} \xi_j^{\frac{1-\sigma}{\sigma}}$$

This shows clearly that, conditional on choosing j , consumption q_j is increasing in the preference value.

Using this first order condition in the choice rule, the choice rule becomes

$$\max_j \{u(q_j(\xi_j)\xi_j) - u'(q_j(\xi_j)\xi_j)\xi_j q_j(\xi_j)\} = \max_j \left\{ \frac{\sigma}{1-\sigma} \left(\frac{1}{\Lambda}\right)^{\frac{1-\sigma}{\sigma}} p_j^{-\frac{1-\sigma}{\sigma}} \xi_j^{\frac{1-\sigma}{\sigma}} \right\} = \max_j \{-\log p_j + \log \xi_j\}$$

Note that absent any *distributional assumptions* on tastes, we obtain the result that under power utility, the efficient choice probabilities themselves are equivalent to those under incomplete markets. However the quantities consumed conditional on realizations of the shocks are different. As above, if we assume that F is the product of independent Frechet draws with tail parameter η , we obtain:

$$\rho_j = \frac{p_j^{-\eta}}{\sum_k p_k^{-\eta}}.$$

Aggregation. Aggregation is more complicated in the complete markets case due to the dependence of q_j on the preference shock ξ_j . Total demand for good j is:

$$Q_j = \int_{\xi} x_j(\xi) q_j(\xi_j) g(\xi) d\xi \quad , \quad q_j(\xi_j) = \left(\frac{1}{\Lambda}\right)^{\frac{1}{\sigma}} p_j^{-\frac{1}{\sigma}} \xi_j^{\frac{1-\sigma}{\sigma}}.$$

Simplifying this:

$$Q_j = \left(\frac{1}{\Lambda}\right)^{\frac{1}{\sigma}} p_j^{-\frac{1}{\sigma}} \mathbb{E} \left[\xi_j^{\frac{1-\sigma}{\sigma}} \mid \text{Choose } j \right].$$

We can aggregate to obtain an expression for the multiplier Λ . Starting with the first order condition for $q_j(\xi_j)$, and then aggregating:

$$\begin{aligned} u'(q_j(\xi_j) \xi_j) q_j(\xi_j) \xi_j &= \Lambda p_j q_j(\xi_j) \\ \int \sum_j x_j(\xi_j) (q_j(\xi_j) \xi_j)^{1-\sigma} dG(\xi) &= \Lambda \int \sum_j x_j(\xi_j) p_j q_j(\xi_j) \varphi(\xi) d\xi \\ \int \sum_j x_j(\xi_j) \left(\frac{1}{\Lambda}\right)^{\frac{1-\sigma}{\sigma}} p_j^{-\frac{1-\sigma}{\sigma}} \xi_j^{\frac{1-\sigma}{\sigma}} dG(\xi) &= \Lambda W \\ \left(\frac{1}{\Lambda}\right)^{\frac{1}{\sigma}} &= \frac{W}{\sum_j p_j^{-\frac{1-\sigma}{\sigma}} \mathbb{E} \left[\xi_j^{\frac{1-\sigma}{\sigma}} \mid \text{Choose } j \right]}. \end{aligned}$$

Using this in the demand curve for Q_j :

$$Q_j = \left(\frac{p_j^{-\frac{1-\sigma}{\sigma}} \mathbb{E} \left[\xi_j^{\frac{1-\sigma}{\sigma}} \mid \text{Choose } j \right]}{\sum_j p_j^{-\frac{1-\sigma}{\sigma}} \mathbb{E} \left[\xi_j^{\frac{1-\sigma}{\sigma}} \mid \text{Choose } j \right]} \right) \frac{W}{p_j}$$

Under the assumption that ξ_j are distributed Frechet, then we have the result that

$$\mathbb{E} \left[\xi_j^{\frac{1-\sigma}{\sigma}} \mid \text{Choose } j \right] = \Gamma \left(\frac{\eta - \left(\frac{1-\sigma}{\sigma}\right)}{\eta} \right) \rho_j^{-\frac{1-\sigma}{\eta\sigma}}.$$

where $\Gamma(\cdot)$ is the cdf of the gamma distribution. Substituting in the expression for ρ_j :

$$\mathbb{E} \left[\xi_j^{\frac{1-\sigma}{\sigma}} \mid \text{Choose } j \right] = \left\langle \frac{\bar{\Gamma}}{(\sum_k p_k^{-\eta})^{-\frac{1-\sigma}{\eta\sigma}}} \right\rangle p_j^{\frac{1-\sigma}{\sigma}}.$$

Finally, using this in the expression for Q_j , the term in $\langle \cdot \rangle$ cancels from numerator and denominator, to give:

$$Q_j = \left(\frac{p_j^{-\frac{1-\sigma}{\sigma}} p_j^{\frac{1-\sigma}{\sigma}}}{\sum_j p_j^{-\frac{1-\sigma}{\sigma}} p_j^{\frac{1-\sigma}{\sigma}}} \right) \frac{W}{p_j}$$

$$Q_j = \left(\frac{1}{J} \right) \frac{W}{p_j}$$

We can also return to express $q_j(\xi_j)$ in closed form. Let's define $P = \left[\sum_j p_j^{-\eta} \right]^{-\frac{1}{\eta}}$. Using this:

$$q_j(\xi_j) = \frac{1}{\Gamma} \xi_j^{\frac{1-\sigma}{\sigma}} \left(\frac{p_j}{P} \right)^{-\frac{1-\sigma}{\sigma}} Q_j$$

Collecting results, we have

$$\rho_j = \left(\frac{p_j}{P} \right)^{-\eta}$$

$$P = \left[\sum_j p_j^{-\eta} \right]^{-\frac{1}{\eta}}$$

$$Q_j = \frac{1}{J} \frac{W}{p_j}$$

$$q_j(\xi_j) = \frac{1}{\Gamma} \xi_j^{\frac{1-\sigma}{\sigma}} \left(\frac{p_j}{P} \right)^{-\frac{1-\sigma}{\sigma}} Q_j.$$

Representative agent. The aggregate demands for each good j are as if they are derived from a representative agent facing the problem

$$\max_{Q_j} \frac{Q^{1-\sigma}}{1-\sigma} \quad Q = \prod_j Q_j$$

$$\sum_j p_j Q_j = W$$

That is, the representative agent has CES utility with a unit elasticity, i.e. Cobb-Douglas. Note that $\sum_j p_j Q_j \neq PQ$. The *ideal price index* for the representative agent is that which is usually obtained under Cobb-Douglas:

$$\tilde{P} = \prod \left(\frac{p_j}{J} \right)^{1/J}.$$

E.3. Conclusions

We can now collect results across incomplete and complete markets

1. Complete markets

- (a) Aggregate demand Q_j can be obtained from a representative agent with CES preferences with elasticity of substitution of 1

$$(Q_1^*, \dots, Q_J^*) = \arg \max_{Q_j} \frac{Q^{1-\sigma}}{1-\sigma}, \quad Q = \prod Q_j \quad \text{subject to} \quad \sum_j p_j q_j = W.$$

$$\text{Demand:} \quad Q_j^* = \frac{1}{J} \left(\frac{p_j}{1} \right)^{-1} \left(\frac{W}{1} \right)$$

- (b) Individual demand $q(\xi_j)$ of individuals buying good j differ by ξ_j :

$$\rho_j = \left(\frac{p_j}{P_j} \right)^{-\eta}, \quad q_j(\xi_j) = \frac{1}{\Gamma} \xi_j^{\frac{1-\sigma}{\sigma}} \left(\frac{p_j}{P} \right)^{-\frac{1-\sigma}{\sigma}} Q_j^*, \quad \text{where } P = \left[\sum_j p_j^{-\eta} \right]^{-\frac{1}{\eta}}$$

2. Incomplete markets

- (a) Aggregate demand Q_j can be obtained from a representative agent with CES preferences with elasticity of substitution of $1 + \eta$

$$(Q_1^*, \dots, Q_J^*) = \arg \max_{Q_j} \frac{Q^{1-\sigma}}{1-\sigma}, \quad Q = \left[\sum_j Q_j^{\frac{\eta}{1+\eta}} \right]^{\frac{1+\eta}{\eta}} \quad \text{subject to} \quad \sum_j p_j q_j = W.$$

$$\text{Demand:} \quad Q_j^* = \left(\frac{p_j}{P} \right)^{-(1+\eta)} \left(\frac{W}{P} \right), \quad \text{where } P = \left[\sum_j p_j^{-\eta} \right]^{-\frac{1}{\eta}}$$

- (b) Individual demand $q(\xi_j)$ of individuals buying good j are independent of ξ_j :

$$\rho_j = \left(\frac{p_j}{P_j} \right)^{-\eta}, \quad q_j(\xi_j) = \frac{W}{p_j}, \quad \text{where } P = \left[\sum_j p_j^{-\eta} \right]^{-\frac{1}{\eta}}$$

Combined, these results establish a number of conclusions

1. Under incomplete markets, we have generalized the aggregation results of ADT to the multiplicative case with power utility. As in ADT, aggregation to a CES representative agent is possible.
2. The fact that “aggregate demands for each good can be constructed from the optimization problem of a representative agent with CES utility” does not imply that “the equilibrium allocation is efficient”.

- This is clear from the fact that with complete and incomplete markets, aggregate demands for each good can be constructed from the optimization problem of a representative agent with CES utility, but the incomplete markets allocation is inefficient.
3. The ADT case is nested under the limit of $\sigma \rightarrow 1$, and as per our results in the main text, is efficient.