RESEARCH STATEMENT

Raghav Malhotra (r.malhotra@warwick.ac.uk)

I am an applied microeconomic theorist, my core focus is demand theory but my work is wide-ranging and touches on a variety of topics. My work can be divided into three strands: (i) Demand Theory, (ii) Public Finance (iii) Labor economics. In the first strand, I try to understand what can be learnt about consumers from choice data. I also try to understand what can be said about the measures of welfare changes when we observe behaviour at differing prices. I believe it is important to understand the links between counterfactual estimates and our assumptions about behaviour. These are often untestable and may be pivotal to reaching our conclusions.

This theme extends to my work in public finance (the second strand), where I study how to better effect redistribution from the viewpoint of a state. Finally, we come to the third strand, labor economics. I try to see if there are testable empirical consequences of economic theory when looking at the effect of technological change on workers. I try to make sure that my work is of direct policy relevance and tries to communicate with the discipline outside just theory. Below I summarise my research, including my job market paper.

i. Demand Theory My Job Market Paper explores a classic issue, measuring the welfare impact of a price change. One way to do this is by using a long panel on individuals' consumption choices. One can use such data to estimate individual demand functions and preferences and, subsequently, the welfare impact. Unfortunately, the data typically available to measure such impacts often take the form of cross-sections, which leads us to study what can be learnt from cross-sectional distributions of demand.

In my job market paper "*Price Changes and Welfare Analysis: Measurement under Individual Heterogeneity*" (with Sebastiaan Maes), we show how to use the higher moments of demand to considerably improve estimates of the welfare effects of price changes. Our approach also enables inference on the distribution of welfare changes. We make a methodological contribution by deriving the relationship between the conditional moments of demand and the Slutsky equation. This allows us to use the information that cross-sectional data contains about income effects, which are essential for welfare calculations as they reveal how much individuals need to be compensated for a price change.

The results developed in this paper can also be used to test stochastic rationalizability. We can characterize rationalizable cross-sectional distributions of demand locally in the two-good case. Our method is computationally feasible, and we use it to construct a semi-decidable test of rationality. With more than two goods, we can test the negative semidefiniteness of compensated demand but not symmetry.

To demonstrate our approach's usefulness, we estimate the welfare effect of a 10% increase in transport prices on consumer welfare. We collect data on households' consumption bundles and income from 14 waves of the UK Household Budget Survey (2006-2019); and prices from the Office for National Statistics. Our results suggest that the representative agent approach significantly underestimates the welfare impact, as our estimate is 27.2% higher. Moreover, this bias is larger for individuals with low disposable income, who may well find it higher to substitute away from their consumption. We perform a similar exercise for food and housing and find even larger effects for the latter.

In "*Exact inference from finite market data* (with F. Kubler and H. Polemarchakis)" (R & R AER), we analyze inference from a finite number of observations in the following sense: potential observations are infinite, even dense in an appropriate domain, but we require inference to be attained after a finite number. We ask what one needs to know about preferences a priori in order

to be able to make non-trivial, exact inferences about the underlying data generating preference. First, we examine the case of observations on an individual's demand at different prices and incomes. Second, for the case of an exchange economy, we assume that only equilibrium prices and profiles of individual endowments are observable.

For both setups, we show that, given any two consumption bundles between which an individual is not indifferent, after sufficiently many observations, we can infer the individual's choices. A question that arises directly from this analysis is how these results extend to equilibrium comparative statics. The transfer paradox, introduced by Leontief (1936), makes it clear that knowledge of utility functions is necessary to identify even the direction of welfare effects of transfers. We show that, while predictions of exact comparative statics are generally impossible, we can predict approximate equilibrium prices from finite data.

The paper *(Functional) Characterizations vs (Finite) Tests: Partially Unifying Functional and Inequality-Based Approaches to Testing* deals with the following problem:- Consider a calculus-based restriction of a DM's demand function. Assume that demand functions that arise from utilities within our class of interest (e.g., separable preferences) satisfy this restriction. How and under what conditions can finite data be used to test a preference for a restriction?

I demonstrate that restrictions on demand can be tested, but only if they are characterizations. Weakening the standard assumptions needed for revealed preference characterizations, we construct a computationally efficient test. This is important for two reasons: first, several properties (e.g., weak separability) have very natural functional tests but no intuitive or simple finite tests. Second, our procedure lets an analyst test large datasets efficiently for properties where exhaustive restrictions are provably hard to check. This is an important theoretical contribution because it says that relaxing the if and only if condition on rationalizability can restore computational feasibility.

ii. Public Finance In *"Rethinking Distribution: Introducing Market Segmentation as a Policy Instrument"* (with Yatish Arya), we argue that policymakers must confront the increase recent increase in inequality by expanding the set of redistributive policy instruments traditionally available. We consider a stylized general equilibrium exchange economy where some goods are considered essential, whereas others are not and assume that only essential goods are relevant for distributional concerns. We find that under certain conditions, conventional redistributive instruments cease to achieve social welfare goals effectively. We then introduce a new policy instrument called Market Segmentation (MS henceforth) and demonstrate its superior performance under these conditions.

Under MS, not every good type is freely tradeable for others. Concerning ourselves with distributional issues, we segment the market for essentials from non-essentials, i.e. they cannot be freely traded for each other. We formalize this by not allowing agents to spend the income generated by selling non-essential goods on essential goods above a certain threshold. Thus, we have one budget constraint with essential goods and one for non-essentials. In this setup, we compare the welfare effects of MS with three other policies, namely 1) Subsidies, 2) Direct transfers, and 3) Quantity rationing. We find that if the relative number of low-income individuals in the economy is high and "essentials" are consumed inelastically, MS outperforms transfers and subsidies.

iii, Labor Economics. The paper *"The Evolution of Skill Use Within and Between Jobs"* (With Costas Cavounidis, Vittoria Dicandia and Kevin Lang) integrates insights from both task-based and skill-biased approaches to model technical change. We model occupations as combining skills, akin to tasks in the tasks model, to produce intermediate goods, which are then embedded in a tractable GE model. Occupations utilize workers' skills in heterogeneous ways, and we impose

2

nearly no structure on their production technologies. Our model allows workers to choose first the skills they develop and then their occupations. The market prices intermediate goods. Crucially, unlike most of the literature, we do not assume that either occupations or the various skills they use are hierarchical.

We then exploit employment trends to uncover changes in skills' productivities. Our approach characterizes the kind of technological change that best explains shifts in unemployment. We use three DOT editions, 1960, 1970, and 1980 Censuses and the March CPS, to estimate changes in the relative productivities of skills. We find 'skill bias' - finger-dexterity productivity grew rapidly, while abstract-skill productivity lagged. With substitutability between abstract and routine inputs, these results also explain changing skill use within occupations.

Another paper "*Obsolescence Rents: Teamsters, Truckers, and Impending Innovations.* " (with Costas Cavounidis, Qingyuan Chai, and Kevin Lang) studies the behaviour of labour markets when a demand shock is anticipated but hasn't quite arrived. We consider large, permanent shocks to individual occupations whose arrival date is uncertain. We are motivated by the (nearly) inevitable arrival of self-driving trucks, which will dramatically reduce demand for truck drivers.

We develop an overlapping generations model in which the labor market is in steady-state before the announcement of a future technological change. The change will greatly reduce the demand for some occupation - say, widgeters. The exact arrival of the new technology is unknown but is believed to occur at a constant hazard. There are two phases of interest: the anticipatory stage before the shock arrives, and the adjustment stage after it arrives. The results regarding the post-shock aftermath stage are standard. The sudden decline in demand for widgeters reduces widgeters' wage and those who can, take other jobs. As older widgeters retire, the wage rises towards a new steady state (although not necessarily approaching it monotonically), with lower employment and a lower wage than the initial one. More importantly, in the anticipatory dread stage, young workers expect that the wage of widgeters may decline dramatically in their lifetime. With imperfect mobility, they must be compensated to become widgeters - receiving what we dub obsolescence rents.

We then apply the model to the effect of the arrival of motor trucks on teamsters, workers who drove vehicles drawn by teams of horses. These vehicles were the primary means for delivering goods locally. Long-distance freight was primarily handled by railroad. The analysis we conduct largely confirms our expectations.

A work in progress '*Skill Traps and Regional Divergence* '"(with Ivan Yotzov) considers the effect of the initial distribution of skills on commuting zones' responses to the China Shock. Regions with a more diverse set of industries rapidly see an uptick in non-manufacturing employment. In contrast, highly specialized regions saw no such uptick but rather experienced a large increase in the share of people outside the labor force. Next, we consider the effects of the China import shock on net job creation and net establishment entry rates. For both, we see a positive and significant effect in low specialization CZ and a negative effect in regions with high specialization, giving more evidence for differential adjustment. This suggests what we call a "skill trap".

Conclusion. In the future, I intend to continue working on each of the distinct strands. Concerning demand, I think extending both welfare measurement and revealed preference analysis to setups where individuals have behavioural biases is a fertile yet understudied area of research. My PF paper identifies a family of policy instruments that classical commodity taxation might have missed: taxation between groups of commodities. We exploit one such instrument. This would an area of future research which may be worth pursuing. Concerning labor, I believe that studying the implications of human capital acquisition on aggregate variables in an economy is an interesting avenue of future work.