## Essays on Technological Change, Skill Premia and Development


Kunst, D.M.



2020

# 6 Afterword

During conference small talks throughout my PhD, I have found it difficult to specify my own subfield within economics: more than once, the perhaps best fitting shorthand description "labor economics" created expectations about microdata and the exploitation of an exogenous policy reform, which I invariably failed to meet. I would hence like to use this afterword to offer some reflections on methodological conventions, recent trends in the economics literature, and how this thesis fits in.

On the one hand, it is in line with the trend of economic research becoming more empirical: in a nearly exhaustive sample of published economics papers between 1980 to 2015, Angrist et al. (2017) find that the journal influence-weighted share of papers classified as "empirical" (as opposed to "theoretical" or "econometrics") has increased from about one third to 55 percent. On the other hand, three of the core chapters of this thesis run against the trend of empirical work being design-based, exploiting some source of exogenous (quasi-)experimental variation; a trend which Angrist and Pischke (2010) refer to as the "credibility revolution".

In my view, there is a lot to be welcomed about the trends towards more, and more rigorous, empiricism in economics. In a famous essay, the philosopher Isaiah Berlin (1953) assigns important writers and thinkers to one of two archetypes: hedgehogs, who look at the world through the lens of a single grand idea that organizes their thinking, and foxes, whose thinking does not evolve around any overarching principle. The recent developments in economics can be viewed as an increasing number of empiricist-foxes poking holes into the theoretical edifices of their hedgehog-colleagues, who have been put on the defensive by the foxes´ improving ability to convincingly demonstrate that the hedgehog´s preferred model fails in important real-world settings.

Examples include the robust finding that moderate minimum wage increases do not cause significant employment losses (casting doubt on widely and strongly held[1] assumptions about labor markets being close to perfectly competitive), the many insights from behavioral economics (highlighting departures from the assumption of rationality), or the recent findings of much larger distributional effects of trade than previously acknowledged (pointing

[1]Perhaps best exemplified by nobel laureate James Buchanan´s response to the new evidence on the minimum wage: *"Just as no physicist would claim that "water runs uphill", no self-respecting economist would claim that increases in the minimum wage increase employment. Such a claim, if seriously advanced, becomes equivalent to a denial that there is even minimal scientific content in economics, and that, in consequence, economists can do nothing but write as advocates for ideological interests. Fortunately, only a handful of economists are willing to throw over the teaching of two centuries; we have not yet become a bevy of camp-following whores"* (Buchanan, 1996).

to significant frictions in “compensating the losers”, or in moving them to other industries). The new emphasis on rigorous empiricism has been successful at exposing economists who mistake their preferred model for *the* model (paraphrasing Rodrik 2015), paving the way for more informed discussions and better public policies.

However, this may have come at the expense of *generating* fewer hypotheses worth testing, hindering the progression of our knowledge in a different way: in an engaging statement of this argument in *The New Republic,* Schreiber (2007) quotes several prominent economists such as James Heckman, Raj Chetty and David Card who are concerned about a diversion of talent towards questions that are sufficiently small or specific to be answered with (quasi-)experimental methods, and away from investigating questions that are more important for human welfare. To fix thoughts about this possibility, Akerlof (2019) proposes a simple model in which a researcher selects from possible research topics that can be characterized along two dimensions: “hardness”, which captures the ease or difficulty of producing precise work on the topic, and “importance”. An emphasis on hardness may create “sins of omission” if important yet “soft” topics do not attract sufficient attention.

Akerlof argues that economics presently suffers from such a “hardness bias”, and presents three possible reasons: first, the pride many economists take in viewing their discipline as the most scientific social science, and a desire to keep their place in the perceived methodological pecking order. Second, an evaluation process in journals and hiring or promotion committees that favors precision–a relatively well-defined concept–over importance, which is easier to disagree on. And third, a strong emphasis on mathematics as opposed to the application of theory to the real world in graduate schools, leading to a self-selection of students with a preference for hardness relative to importance into PhD programs in economics.

Among the possible consequences, a trend towards over-specialization is perhaps the most worrisome: generalists need to meet the standards of precision (empirical or theoretical) for multiple fields, whereas specialists only need to meet the standards of one. Hence, hardness bias encourages specialization, and Akerlof points to an increasing specialization of departments and seminar series and the proliferation of subfield journals as evidence that specialization in economics has indeed been increasing. This discourages research on topics at the intersection between subfields, and risks producing important sins of omission. For instance, Rajan (2011) argues that the global financial crisis of 2008 caught most economists by surprise even though all the individual pieces contributing to it had been understood: the problem was that these pieces could only have been put together by the rare economist with theoretical and institutional knowledge spanning the fields of macroeconomics, finance and real estate economics.[2]

At the same time, it appears to me that our discipline can become a more useful participant in discussions about the big societal challenges of our time–such as the rapid transforma-

[2] Moreover, the increasing focus on the number of publications in the “Top 5” journals for academic recruitments, which Heckman and Moktan (2018) show to be a problematic proxy for academic influence, may in part reflect the fact that economists in hiring committees are increasingly unable to judge the work of candidates on substantive grounds.

tion of labor markets, increasing market power, and global warming–if we tackle existing hardness-biases and their consequences. For the sake of concreteness, let me propose four examples: the first and most general point concerns the place of descriptive papers and case studies near the bottom of the pecking order when it comes to academic publishing.[3] While observational data to study topics such as the evolution of inequality or intergenerational mobility are becoming increasingly available (for instance, see the "World Inequality Database" (Alvaredo et al., 2018) or the data from the "Opportunity Insights" project (Chetty et al., 2017)), the uncertain prospects when it comes to publishing makes engaging with such data a risky business, especially for graduate students or junior faculty.

A personal experience perhaps illustrates this point: when I met with Remco to discuss potential thesis topics, he had been looking for a PhD student to analyze the extended "Occupational Wages around the World" database for some years already. Around the same time I was offered to write my thesis about the effects of the exogenous roll-out of internet on the performance of professional chess players, who could then practice online with much stronger opponents. For me, this was an easy choice–but I believe that the current methodological taste prevailing in hiring committees and editorial boards creates clear incentives for PhD students to go for the "harder" research design in such situations, sometimes at the cost of the relevance of the topic.

One may see this as a welcome development; perhaps causal evidence on a small question is often more useful than descriptive evidence on a big one (and some field experiments are able to address questions that are "big" in terms of their relevance for public policy, see Duflo (2017) for examples). But there is reason for skepticism: after decades of research on how to maximize the accuracy of real-world predictions, Tetlock and Gardner (2016) conclude that the best forecasters start from a base probability informed by historical precedent, and use additional information about the situation at hand to adjust this "prior" in a second step. The important point here is that making good predictions requires sound knowledge about real-world *correlations:* such correlations also pick up causal pathways that are not yet well understood or involve variables that are inherently hard to quantify, and which many design-based empirical studies would hence explicitly try to exclude from the identification. While the statement "correlation does not equal causation" is deeply ingrained in today´s empirical economists for good reason, an underestimation of the value of correlational evidence would therefore be a costly collateral damage when it comes to our ability to make accurate predictions.

For example, consider the finding of chapter 2 that there is a strong correlation between growing skill premia, increases in a country´s openness to trade, and decreases in union

[3] I simplify my argument by contrasting an archetypical design-based empirical study which exploits some form of exogenous variation to arrive at a causal estimate with research that does not center around such a design. Such research is often referred to as "descriptive", "observational" or "correlational" and usually involves survey- or administrative data. In practice, quasi-experimental methods (such as instrumental variable or double difference-techniques) blur this line, so that "hardness" in empirical research comes in more than two shades.

strength: these correlations are unlikely to be fully causal, as we for instance do not control for technology adoption of firms (likely to be correlated with trade openness) or changing pay norms (likely to be correlated with unionization trends). But can these findings serve as a useful prior to start from when predicting the effects of increasing trade, or decreasing union strength, in the future? Arguably, yes: absent information to the contrary, it appears reasonable to assume that what has been associated in the past will still be associated in the future. While researchers who dismiss correlational evidence often do so with the explicit motivation of *avoiding* misguided predictions, a research culture that discourages the conduct of interesting observational studies may hence actually make us collectively *worse* at predicting the effects of real-world policies by equipping us with less informative priors.

Moreover, the hurdles are even higher when it comes to topics where cross-national panel datasets are not available, such as the increasing ability of capital to escape taxation through the offshore finance-system (Shaxson, 2011): there may be a great deal to learn about the precise mechanisms and incentives at work and the context-specific economic significance from well-chosen case studies (Flyvbjerg, 2006)–yet such research currently stands little chance to be published in a prestigious economics journal, and consequently little chance to be undertaken by an academic economist in the first place.[4]

The near absence of case study-research fits into a broader pattern of emphasizing the "science" at the expense of the "craft" aspect of economics: Rodrik (2015) highlights that all knowledge generated by economics is context-dependent, so that the competent analysis of any particular situation hinges on the selection of the model that best captures its most relevant aspects. This idea goes back to Keynes, who described economics as "*a science of thinking in terms of models joined to the art of choosing models which are relevant to the contemporary world*". In other contexts, the importance of examples for learning this "art" is already well established: artificial intelligence researchers have largely abandoned the rule-based, "symbolic", approach in favour of the examples-based "neural network" approach (McAfee and Brynjolfsson, 2017), and psychologists agree that expert judgment results from the accumulation of context-specific detail (Kahneman and Klein, 2009).

Also the "credibility revolution" in economics has provided ample support for this perspective by challenging claims to universality by any particular model–but somewhat paradoxically, it has at the same time discouraged economists from engaging in the craft of model selection by pushing the creative use of observational data and case studies to study important real-world problems to the fringes of the methodological toolbox of academic economists.[5]

[4] Interestingly, case studies do get cited in economics papers when it comes to making the results of aggregate empirical analyses sound more plausible. However, to take chapter 3 of this thesis as an example, the cited case studies on craftsmen in the printing and textile industry and the deskilling aspects of recent technology adoption in China all appeared either in bottom-tier economics or other social science journals. I would expect this pattern to be quite general–and if this is correct, economists implicitly acknowledge the value of detailed case study research in addition to empirical studies at a more aggregate level, while at the same time providing little incentives for the production of such research.

[5] For an example of a creative use of observational data to identify the binding constraints to economic development, see the "growth diagnostics" approach of Hausmann et al. (2008).

It appears to me that a more problem-driven economics discipline that held the craft of model selection in higher regard would be a better guide for public policy.

A second and more specific example of arguably "missing" research because of hardness bias and over-specialization is the analysis of labor market power: although the idea that firms may have wage-setting power dates back to Adam Smith and was formalized already in Robinson (1933), there is a surprising lack of published research on the phenomenon.[6] Kruger (2018) argues that increasing labor market power due to declining union strength, collective bargaining and minimum wages, a proliferation of outsourcing (which facilitates the excercise of monopsony power), non-compete and no-poaching clauses in labor contracts, as well as increasing employer concentration can account for slow wage growth in the US in spite of low unemployment. At least some of these phenomena have also been documented for many other countries, and may have contributed to the phenomenon of globally declining labor shares since the 1980s (Karabarbounis and Neiman, 2014).

Investigating the link between the labor market power of firms, individual wages, and the aggregate labor share calls for research at the intersections of labor economics, macroeconomics and industrial organization–three subfields with very different methodological tastes and publication standards. While analyses in labor economics are often of partial equilibrium nature, macroeconomics emphasises general equilibrium effects, and while reduced-form analyses are common in labor economics, publication norms in industrial organization typically require structural analyses. Research on labor market power calls for the application of concepts from industrial organization to labor markets– yet Angrist and Pischke (2010) single out this subfield as being one of the few to resist the "credibility revolution", and argue that this leads to assumption-driven results and a scarcity of useful empirical evidence. In any case, it probably leads to little understanding of, or interest in, concepts from industrial organization by most labor economists. Given the hurdles presenting themselves to researchers trying to publish at the intersection of these subfields, it is perhaps not surprising that labor market power has not attracted more attention in the past–a lack of attention which has carried over to policymakers, as Azar et al. (2019) point out.

A third example is the dismissal of measures of subjective well-being as too "soft" to merit serious analysis by most economists: although the shortcomings of changes in the gross domestic product as a proxy for changes in well-being are obvious in particular for high income-countries (Stiglitz et al., 2009), the topic still does not attract much attention beyond the fringes of economics. Frey and Stutzer (2002) give examples how research on self-reported "happiness" as a summary measure of subjective well-being can improve economic theory and policy–but as of 2019, the not-so-original but "hard" argument that happiness is not to be taken seriously because it is ordinal (Bond and Lang, 2019) is more likely to be published in a "Top 5"-journal than a paper that makes use of the increasingly available

[6] Among the exceptions is Staiger et al. (2010), which documents monopsony power in the labor market for US nurses. A working paper version of the study was published already in 1999, hinting at an unusually rocky path to publication even by today´s standards.

happiness data.

Since 2012, John Helliwell, Richard Layard, Jeffrey Sachs and their collaborators publish the annual "World Happiness Report", which draws on standardized happiness surveys from more than 150 countries to shed light on policy questions and address methodological concerns alongside. For instance, the 2018 report studies the relationship between happiness and international as well as rural-urban migration, which is both topical and helps to answer the question whether happiness depends more on where you were born or on where you live (to which the short answer seems to be: on where you live, see Helliwell et al. 2018). Given the attention economists pay to the (ordinal) concept of utility and the maximization of welfare (which requires interpersonal comparisons of utility), it appears to me that our discipline has little excuse not to, and much to gain from, engaging seriously with research on happiness (and measures of well-being more generally).[7]

A final example relates to policies to limit global warming, and the endogeneity of (green) technological change. There is broad consensus among economists on the basic problem and the first-best solution: greenhouse gas emissions are a negative externality, and a substantial part of global emission reductions will likely need to come from technological innovations. An appropriate global carbon tax could both make us internalize the social cost of emissions in the short run, and steer technological change in the right direction in the long run.[8]

However, the public good-characteristics of the global climate make climate policy inherently second-best, and a sustainable climate policy also has to balance diverging notions of fairness, time horizons and risk attitudes across and within countries. Some climate economists have little patience for the resulting complexities of real-world climate policy: for instance, the prominent VU climate economist Richard Tol argues that any climate policy beyond a carbon tax reflects for the most part rent-seeking on the part of politicians, and an innate desire of bureaucrats to increase bureaucracy (Tol, 2017, 2019). From this perspective, there is little left to do for economists when it comes to climate policy debates, other than repeating their call for a global carbon tax–and perhaps this perception contributes to explaining the lack of research on global warming by economists highlighted by Oswald and Stern (2019).

But I wonder whether a discipline that paid greater respect to the craft-aspect of economics would also need to stop there: the endogenous nature of technological change is well understood theoretically (Acemoglu, 2002, 2007), as is the case for public sponsoring of research that involves positive externalities or creates socially desirable new markets as opposed to merely fixing existing ones (Mazzucato, 2013). The question is what follows from this: (green) industrial policy still does not have a good name among many economists because

[7] Also see Sperling (2020) for an engaging attempt to define "economic dignity" as a concept that covers the blind spots of widely-used metrics such as GDP.

[8] See for instance the poll by the IGM Economic Experts Panel (2011), which highlights the popularity of a carbon tax among prominent economists. However, note that uncertainty about the elasticity of emissions with respect to a carbon tax, and economies of scale in the discovery and production of green alternatives, provide strong arguments for a first-best climate policy that goes beyond a carbon tax (Stern, 2015).

of the view that it depends on the ability to “pick winners”, although Rodrik (2019) argues that it instead depends on the weaker condition of “letting losers go” in an appropriately context-specific institutional framework of public research funding.

Evidence from large-scale empirical studies is unlikely to be able to settle this debate (Rodrik, 2012), even if the studies find a source of exogenous policy variation to exploit.[9] By contrast, academic economist´s understanding of markets and incentives is likely to generate useful insights about when and how public innovation funding works when applied to the greater detail which case studies provide. Unfortunately, there are presently few incentives for academic economists to take a closer look at specific contexts, as resulting papers will likely not appear “hard” enough for either empiricists or theorists to merit publication: particular industrial or innovation policies tend to be not exogenous enough for empiricists, and not general enough for theorists.

If these examples of possible consequences of hardness bias appear too speculative to some readers, this may relate to the point I am trying to make: perhaps, allowing ourselves a bit more informed speculation in a spirit of good faith and intellectual humility will be what it takes to make our discipline even more useful. Significant insights on how to address today´s big challenges are unlikely to see the light of day in a polished form, and a culture of only voicing and publishing “hard” evidence appears prone to creating the sins of that Akerlof is concerned about. By contrast, a research culture that attaches a higher value to the craft of model selection in real-world contexts will also automatically be one in which researchers will see the need to cross the boundaries of subfields more often, and fostering this intellectual exchange may also fertilize the equally important basic research within subfields. Other readers may argue that we should leave the “softer” research to policy research institutes, central banks and think tanks. However, the potential sins of omission concern topics of first-order importance to foundational economic and societal problems, and I do not see a good justification for maintaining a research culture that steers the best academic economists away from them.

Fortunately, priorities may have started to shift already: interest in inequality and labor market power has been surging over the past few years, and Suresh Naidu, Dani Rodrik, and Gabriel Zucman have recently launched the “Economics for Inclusive Prosperity” network, which invites academic economists to submit policy briefs that combine “big picture” thinking on current challenges that is firmly grounded in economic theory with a discussion of concrete policy options. I hope that also the culture in the core business of academic economists, academic publishing, will create more space for what may be the one great quality that hedgehogs in Isaiah Berlin´s fable possess: not their dogmatism, but their interest in how their work contributes to some bigger objective–which should arguably be to inform

[9]For instance, Criscuolo et al. (2019) exploit arguably exogenous variation in investment subsidies from area-specific European rules and find that subsidies increase manufacturing employment. Rodrik (2019) points out that while such results are useful as a proof of principle, they leave open the question how “endogenous” real-world industrial policies will play out: they may be more or less effective than the “exogenous allocation”-benchmark, and much will depend on the details of their design.

public policies, and ultimately improve welfare.