

The “task approach” to labor markets: an overview

David H. Autor

Accepted: 11 January 2013

© Institut für Arbeitsmarkt- und Berufsforschung 2013

Abstract An emerging literature argues that changes in the allocation of workplace “tasks” between capital and labor, and between domestic and foreign workers, has altered the structure of labor demand in industrialized countries and fostered employment polarization—that is, rising employment in the highest and lowest paid occupations. Analyzing this phenomenon within the canonical production function framework is challenging, however, because the assignment of tasks to labor and capital in the canonical model is essentially static. This essay sketches an alternative model of the assignment of skills to tasks based upon comparative advantage, reviews key conceptual and practical challenges that researchers face in bringing the “task approach” to the data, and cautions against two common pitfalls that pervade the growing task literature. I conclude with a cautiously optimistic forecast for the potential of the task approach to illuminate the interactions among skill supplies, technological capabilities, and trade and offshoring opportunities, in shaping the aggregate demand for skills, the assignment of skills to tasks, and the evolution of wages.

Der „TASKS-Ansatz“ für Arbeitsmärkte: ein Überblick

Zusammenfassung In einer wachsenden Literatur wird die Auffassung vertreten, dass Veränderungen in der Zuweisung von Arbeitsplatzaufgaben (tasks) zwischen Kapital und Arbeit und zwischen in- und ausländischen Arbeitskräften die

Struktur der Arbeitskräftenachfrage in den Industrieländern verändert und eine Polarisierung von Beschäftigung gefördert hat – d.h. steigende Beschäftigungszahlen in den best- und schlechtestbezahlten Berufen. Eine Analyse dieses Phänomens im Rahmen der gängigen Produktionsfunktion ist jedoch schwierig, da die Zuweisung von Aufgaben zu Arbeit und Kapital in diesem Modell im Wesentlichen statisch ist. Dieses Essay skizziert ein Alternativmodell zur Zuweisung von Kompetenzen zu Aufgaben basierend auf dem komparativem Vorteil, bespricht wichtige konzeptionelle und praktische Schwierigkeiten, vor denen Forscher stehen, wenn sie den „TASKS-Ansatz“ in Daten übertragen wollen und warnt vor zwei gängigen Fallstricken, die die immer stärker wachsende Literatur zu diesem Thema durchziehen. Ich schließe mit einer vorsichtig optimistischen Vorhersage für die Möglichkeiten des TASKS-Ansatzes zur Erläuterung der Interaktionen zwischen Angebot an Kompetenzen, technologischem Potenzial und Handels- und Offshoringmöglichkeiten in der Gestaltung der aggregierten Nachfrage nach Kompetenzen, der Zuweisung von Kompetenzen zu Aufgaben und der Entwicklung von Löhnen.

1 Introduction

The canonical production function found in economic models is anthropomorphic—it’s built to resemble humans. It features a role for labor and a role for machinery (capital), and, in general, these roles are distinct. This is not illogical—after all, these are separate inputs. But what precisely *is* distinctive about the role of each input is left opaque, and the nature of the interactions among them is highly constrained. In particular, capital is either a complement or a substitute for labor, different types of labor are either complements or substitutes for one another, and these

D.H. Autor (✉)
MIT Department of Economics, 50 Memorial Drive, E52-371,
Cambridge, MA 02142, USA
e-mail: dautor@mit.edu

D.H. Autor
NBER, IZA and IAB, Cambridge, MA, USA

roles are essentially fixed. Changes in the supply of each input can of course affect marginal products, but each factor's "purpose" in the production function is both distinct and static.

These restrictions stem from the fact that the canonical production function implicitly equates two distinct aspects of production. One aspect is *which* factors are used as inputs, e.g., capital, high skill labor, low skill labor. The other is *what* services these factors provide. In the canonical setup, a factor's identity and its role in the production function are synonymous. In reality, however, the boundary between "labor tasks" and "capital tasks" in production is permeable and shifting. We observe numerous instances in which technological advances allow machinery to substitute for labor in performing a specific set of tasks—for example, dispensing boarding passes and assigning seats at airport check-in counters in lieu of gate agents—while simultaneously complementing labor in other tasks, for example, allowing gate agents to rapidly identify alternate flights and issue new tickets. This evolving division of labor has a clear economic logic: novel tasks—those demanded by new products, techniques, or services—are often assigned first to workers because workers are flexible and adaptive. As these tasks are formalized and codified, they become fallow for automation since machinery typically has a cost advantage over human labor in rote execution of repetitive tasks.

A growing body of literature argues that the shifting allocation of tasks between capital and labor—and between domestic and foreign labor—has played a key role in reshaping the structure of labor demand in industrialized countries in recent decades. Analyzing this phenomenon is difficult within the canonical production function framework, however, because the assignment of tasks to labor and capital is, for most conceptual purposes, static. To overcome this conceptual hurdle, it is valuable to consider a setting where the equilibrium mapping between production tasks and productive factors responds dynamically to shifting economic and technological forces. The "task approach" to labor markets offers one promising, though far from complete, set of conceptual tools for studying these forces.

In this brief essay, I first sketch a formal model, developed by Acemoglu and Autor (2011), that offers a flexible and tractable approach for analyzing the interactions among skill supplies, technological capabilities, and potential trade and offshoring opportunities in shaping the demand for factors, the assignment of factors to tasks, and the productivity and hence the wage of each factor. This model removes what we view as an artificial set of distinctions typically made between labor, capital, and trade (or offshoring) in the canonical production model, where different factors of production play distinct and often incommensurate roles. In the this alternative model, the fundamental units of production are job tasks, which are combined to produce output. Tasks can be

supplied by domestic labor, foreign labor, or capital, the capabilities of which may change over time. Consequently, the assignment of factors to tasks in this model is determined by comparative advantage; in equilibrium, only the least-cost factor is assigned to any given task.

The second part of the essay briefly discusses some of the key conceptual and practical challenges that researchers face in bringing the "task approach" to the data. The literature is already grappling with these issues, as I note below, and I'm cautiously optimistic that progress will be rapid.

I end by highlighting some of the pitfalls that researchers may encounter in applying the task approach, most prominently, loose and inconsistent task definitions that tend to sow confusion, and the beguiling temptation for researchers to confidently forecast the "future of work" based only on the very modest understanding that we have developed to date.

2 Tasks in production¹

For the purposes of this discussion, it is essential to draw a distinction between skills and tasks. A task is a unit of work activity that produces output. A skill is a worker's *stock* of capabilities for performing various tasks. Workers apply their skills to tasks in exchange for wages. Canonical production functions draw an implicit equivalence between workers' skills and their job tasks, as noted above. Here, we emphasize instead that skills are applied to tasks to produce output—skills do not directly produce output. This distinction is of course inconsequential if workers of a given skill always perform the same set of tasks. It is relevant, however, when the assignment of skills to tasks is subject to change, either because shifts in market prices mandate reallocation of skills to tasks or because the set of tasks demanded in the economy is altered by technological developments, trade, or offshoring. In my view, we are currently in such an era.

If the boundary between "labor tasks" and "capital tasks" is fluid, what determines the division of labor—or, more precisely, the allocation of tasks—between these factors? At least two forces are central, one technological, the other economic. On the technological front, the boundary between labor and capital shifts primarily in one direction: capital typically takes over tasks formerly performed by labor; simultaneously, workers are typically assigned novel tasks before they are automated. This sequence of task allocation makes intuitive sense: when a task is unfamiliar or poses unexpected obstacles, workers can often draw on outside knowledge and problem-solving skills to devise work-arounds. By

¹This section draws extensively on Acemoglu and Autor (2011), and some sections of text are directly reproduced from that article.

contrast, few machines can improvise. Consequently, automating a task requires attaining a level of mastery beyond what is required for a worker to simply perform the task; it must be codified to the point where a relatively inflexible machine can perform the work semi-autonomously.

Even when a task is fully codified, however, this does not mean it *will* be automated. When Nissan Motor Company builds cars in Japan, it makes extensive use of industrial robots to reduce labor costs. When it assembles cars in India, it uses robots far more sparingly.² The key difference between production in India and Japan is not technology but cost: labor is comparatively cheap in India.

At the intersection of these two forces—technological feasibility and economic cost—lies the principle of comparative advantage. Comparative advantage in production means that the factor with the lowest economic cost of performing a task is assigned that task. Economic cost in turn reflects both a factor’s technological capability and its opportunity cost. It’s likely that the technological capabilities and economic costs of robots in Japan and India are quite similar. The opportunity cost of labor in India is far lower, however, and hence Nissan hires Indian workers to perform assembly tasks that are roboticized in Japan.

Though I’ve drawn examples from the information age so far—electronic check-in counters, industrial robots—the ongoing reallocation of tasks between labor and capital is in no sense unique to the current era. One could argue broadly, for example, that the Industrial Revolution was an era in which artisanal tasks were rapidly reallocated from labor to capital while, simultaneously, new cognitive engineering and repair tasks were demanded of labor.³ The era in which we currently live also offers a particularly salient and economically important source of task reallocation: the rapid diffusion of computer technology, spurred by precipitous real price declines, has likely altered the tasks performed by workers at their jobs and ultimately the demand for human skills. In a 2003 paper, Frank Levy, Richard Murnane and I argued that:

“... (1) computer capital substitutes for workers in carrying out a limited and well-defined set of cognitive and manual activities, those that can be accomplished by following explicit rules (what we term “routine tasks”); and (2) that computer capital complements

workers in carrying out problem-solving and complex communication activities (“nonroutine” tasks). Provided that routine and nonroutine tasks are imperfect substitutes, these observations imply measurable changes in the task composition of jobs” (Autor et al. 2003).

The 2003 paper also offered a simple production function framework that showed in barebones terms how information technology might simultaneously displace labor from routine tasks while complementing labor in non-routine tasks.⁴ Subsequent work has greatly enriched this framework.

In this section, I lay out a simple task-based framework proposed in Acemoglu and Autor (2011), which in turn builds on Autor et al. (2003), Acemoglu and Zilibotti (2001), Autor (2006, 2008) and Costinot and Vogel (2010), among others.⁵ Distinct from the canonical model, this framework explicitly incorporates the potential for technological change or international trade and offshoring to cause tasks that were previously performed by domestic workers to either be replaced by capital or substituted by workers abroad. This provides a natural mechanism through which technological advances can lead to real non-monotone changes in the structure of employment by occupation as well as declines in the wages for certain groups of workers.

2.1 A task assignment model

Consider a static environment with a unique final good. For now, assume that the economy is closed and there is no trade in tasks (a possibility we allow for later). The unique final good is produced by combining a continuum of tasks represented by the unit interval, $[0, 1]$. Suppose, to simplify the analysis, that the technology combining the services of tasks is a constant elasticity of substitution aggregator, so that the final good is produced as

$$Y = \left[\int_0^1 y(i)^{\frac{\eta-1}{\eta}} di \right]^{\frac{\eta}{\eta-1}}, \quad (1)$$

where Y denotes the output of a unique final good, $y(i)$ is the “service” or production level of task i , and η is the elasticity of substitution between tasks.⁶

Suppose that there are three types of labor, high, medium and low skill workers, who inelastically supply H , M and L units of labor respectively. At any given point in time, a

²Fackler (2008). Fackler writes, “With Indian wages only about a tenth of wages in Japan, the Chennai [India] factory will rely more on human labor than Oppama [Japan], where robots staff many assembly lines.”

³See Goldin and Katz (1998), Katz and Goldin (2008) and Bessen (2011). Katz and Margo (2012) present evidence that the occupational distribution of U.S. manufacturing (though not the broader economy) “hollowed out” between 1850 and 1910 as new manufacturing capital replaced skilled artisans, thus increasing the occupational shares of both high skill professional, technical, managerial, clerical and sales workers, and, simultaneously, low skill operatives and laborers.

⁴In *The New Division of Labor*, Levy and Murnane (2004) provide a rich and nuanced discussion of the interaction between information technology and human labor in accomplishing workplace tasks.

⁵Goldin and Katz also explore related ideas in many recent writings. See, for example, Katz and Goldin (2008).

⁶Acemoglu and Autor (2011) studied the “Cobb-Douglas” case where $\eta = 1$, or equivalently, $\ln Y = \int_0^1 \ln y(i) di$. I also focus on this case here.

subset $\mathcal{I} \subset [0, 1]$ of the potentially feasible tasks is available (the remaining tasks cannot be produced). Each of the available tasks has the following production function

$$y(i) = A_L \alpha_L(i) l(i) + A_M \alpha_M(i) m(i) + A_H \alpha_H(i) h(i) + A_K \alpha_K(i) k(i), \quad (2)$$

where A terms represent factor-augmenting technology, and $\alpha_L(i)$, $\alpha_M(i)$ and $\alpha_H(i)$ are the task productivity schedules, designating the productivity of low, medium and high skill workers in different tasks. In particular, $\alpha_L(i)$ is the productivity of low skill workers in task i , and $l(i)$ is the number of low skill workers allocated to task i . The remaining terms are defined analogously. Similarly, $\alpha_K(i)$ is the task productivity schedule of capital and $k(i)$ is the amount of capital allocated to task i . Analogously to the canonical model, we can think of A_L , A_M and A_H as low, medium and high skill factor-augmenting technological changes. Distinct from the canonical model, however, a factor-augmenting technological change need not increase the wages of all factors in this setup, as discussed below.

Though each task can be performed by low, medium or high skill workers or by capital, the *comparative advantage* of skill groups differs across tasks, as captured by the α terms. These differences in comparative advantage are central to understanding the interplay of tasks and skills. In particular, the model imposes the following simple structure of comparative advantage: $\alpha_L(i)/\alpha_M(i)$ and $\alpha_M(i)/\alpha_H(i)$ are (strictly) decreasing. This assumption can be interpreted as stating that higher indices correspond to “more complex” tasks in which high skill workers are more productive than medium skill workers and medium skill workers are more productive than low skill workers. Though not very restrictive, this assumption ensures a particularly simple and tight characterization of equilibrium in this economy. We initially set $\alpha_K = 0$ so that capital is not a competing source of task supply, and hence all tasks are supplied by labor.⁷

Factor market clearing requires

$$\int_0^1 l(i) di \leq L, \quad \int_0^1 m(i) di \leq M \quad \text{and} \quad (3)$$

$$\int_0^1 h(i) di \leq H.$$

The structure of the model’s equilibrium is derived in Acemoglu and Autor (2011), and I provide only basic details here. Because of the simple nature of comparative advantage outlined above, the equilibrium of the model involves a partition of the continuum of tasks into three adjacent sets: the

least complex set of tasks, those on the interval $0 \leq i \leq I_L$, will be supplied by L workers; an intermediate set of tasks on the interval $I_L < i \leq I_H$ will be supplied by M workers; and the remaining highest skill set on the interval $I_H < i \leq 1$ will be supplied by H workers. Crucially, I_L and I_H , which are the cut-points of the task partition, are endogenously determined in the model.

Competitive labor markets require that the law of one price for skill applies. Each unit of L labor will receive a wage W_L , and similarly for units of M and H labor. Additionally, since each task can potentially be performed by any skill group, the allocation of tasks to skill groups is governed by a no arbitrage condition: for the marginal task located at I_L , the cost of performing this task must be identical in equilibrium whether it is supplied by L or M workers. Similarly, the cost of performing task I_H must, in equilibrium, be equated between M and H workers. For tasks on the interior of these sets ($i < I_L$, $I_L < i < I_H$, and $i > I_H$), however, workers of the relevant skill groups hold strict comparative advantage. Using these conditions, it is straightforward to demonstrate that there will exist a unique I_L^* and I_H^* that jointly satisfy the law of one price, the no arbitrage condition, and the market clearing condition in (3).

2.2 Implications for interpreting patterns of wages and employment

A central virtue of the model above is that it facilitates a more nuanced view of the nature of technological shifts, which, together with changes in the supply of skills, underlie the bulk of the changes in the earnings distribution over the last several decades. At a basic level, the distinction between high, middle and low skills adds an important degree of freedom to the model, allowing for non-monotone movements in wage levels and wage inequality as seen in the data.

A second distinction is that, while factor-augmenting technological improvements always increase the real earnings of both skilled and unskilled workers in the canonical model, this may not be so in the presence of endogenous allocation of workers to tasks. Acemoglu and Autor (2011) show that a factor-augmenting technological improvement (e.g., an increase in A_H) can reduce the wages of middle skill workers. In particular, this happens when new technologies, by increasing the productivity of high skill workers, encourage some of the tasks previously performed by middle skill workers to be shifted to high skill workers, but a corresponding shift of low skill tasks to middle skill workers is not profitable.⁸

Arguably the most important innovation offered by this task-based framework is that it can be used to investigate the

⁷The canonical model can be nested within this framework trivially by ignoring middle skill workers and assuming that $\alpha_L(i) = \alpha_L > 0$ and $\alpha_H(i) = 0$ for all $i \leq \bar{i}$ and $\alpha_L(i) = 0$ and $\alpha_H(i) = \alpha_H > 0$ for all $i > \bar{i}$. Acemoglu and Autor (2011) discuss other ways in which the model admits special cases similar to the canonical model.

⁸Loosely, this happens when I_H shifts down considerably while I_L does not change by much.

implications of capital (embodied in machines) directly displacing workers from tasks that they previously performed. In general, we would expect that tasks performed by all three skill groups are subject to machine displacement. Based on the patterns documented in the data above, as well as the general characterization of machine-task substitution offered by Autor et al. (2003), I believe that the set of tasks most subject to machine displacement in the current era are those that are routine or codifiable.⁹ Such tasks are primarily, though not exclusively, performed by medium skill workers (e.g., high school graduates and those with less than a four-year college degree). For this reason, let us suppose that there now exists a range of tasks $[I', I''] \subset [I_L, I_H]$ for which $\alpha_K(i)$ increases sufficiently (with fixed cost of capital r) so that they are now more economically performed by machines than by middle skill workers. For all the remaining tasks, i.e., for all $i \notin [I', I'']$, we continue to assume that $\alpha_K(i) = 0$.

What are the implications of this type of technological change for the supply of different types of tasks and for wages? Acemoglu and Autor (2011) answer this question formally. Here, I provide the intuition. A key observation is that this form of technological change has the potential to generate the patterns of wage changes and polarization that have been widely discussed in the recent literature (Autor and Dorn *forthcoming*; Goos et al. 2011; Firpo et al. 2011; Autor et al. 2006, 2008; Goos and Manning 2007); that is, a task-replacing technological change can directly *reduce* wages of a skill group even as it raises total output. The reason, intuitively, is that a task-replacing technological change squeezes out the type of worker previously performing these tasks, thereby creating “excess supply.” These workers are therefore reallocated to tasks for which they have lower comparative advantage, which pushes their wages down.¹⁰ Simultaneously, by reducing the cost, and hence increasing the intensity of use, of the newly automated tasks, the task-replacing technological change complements each of the remaining tasks performed by labor. These countervailing effects imply that the real wage of the group that is directly displaced by technology in a subset of its original tasks

⁹Tasks with these attributes may also be particularly well suited to offshoring since their codifiability makes them readily tradable as the price of communications falls (also a technological change).

¹⁰Notice, however, that this reallocation process does not imply that these displaced *tasks* are obsolete—in fact, just the opposite. As the cost of performing routine tasks has declined by orders of magnitude, their use in production has grown explosively—think, for example, of the amount of processing power that goes into a single Google query. However, because these tasks are now performed by capital rather than labor, the consequences for the earnings power of workers who previously held comparative advantage in these tasks are at best ambiguous.

does not *have to* fall in real terms. But such an outcome is possible—perhaps even likely—in realistic cases.¹¹

The polarization of *occupational* employment also follows from this mechanism. As has been documented in the U.S. and many E.U. countries, the occupations that have contracted most rapidly as a share of total employment over the last three decades—in particular, clerical, administrative support, sales, production and operative positions—are reasonably well characterized as routine task-intensive: many of the core tasks of these occupations follow precise, carefully codified procedures. Because of exponential declines in the cost of computing power, these tasks are increasingly ripe for automation and hence are reassigned from labor to capital. As workers lose comparative advantage in routine-intensive activities, a greater mass of skills is reallocated towards the tails of the occupational distribution—both towards high skill analytic, reasoning and problem solving tasks and, ironically, towards traditionally low skill, in-person service tasks—thus leading to employment polarization.

3 Challenges in putting tasks to the test

The task model offers a potentially powerful framework for studying how changes in skill supplies, technologies, and trade and offshoring opportunities jointly shape the demand for labor, the assignment of skills to tasks and the real wages commanded by workers of different skill groups. There are substantial challenges, however, that confront a researcher in bringing these ideas to the data.

3.1 Measuring tasks

A first challenge is measurement. Conventional skill proxies, such as education and experience, are broadly collected by survey data sources and are widely accepted by applied economists as valid proxies for workers’ stocks of human capital. In the canonical setting where the “law of one price”

¹¹As detailed in Acemoglu and Autor (2011), there are also interesting general equilibrium effects on other skill groups. As the middle skill group is displaced, the task boundaries will change, encroaching on the other two skill groups. The relative degree of encroachment on L versus H workers depends on the shape of the comparative advantage schedules in the neighborhood of the initial set boundaries. If H workers have strong comparative advantage relative to M in their initial tasks, then the upper boundary will move little. If L workers have relatively weak comparative advantage in their initial tasks, then the predominant direction of task reassignment will be that tasks previously performed by L workers are reallocated to M workers. Although total output necessarily rises with the advent of task-replacing technological change, the net effect need not be positive even for groups indirectly affected. For example, if M and L workers are sufficiently close substitutes at the margin, the displacement of M workers into L tasks can reduce the wages of both groups.

for skill prevails, a worker's human capital is also a sufficient statistic for her productivity and hence her potential earnings. Thus, while one *could* ask what tasks a worker performs using her stock of human capital in the canonical setting, the question is not particularly interesting since job tasks—and more broadly, occupations—are not well-defined economic constructs in this framework.¹²

By contrast, a key premise of the task approach is that although skills are used to perform tasks, skills are *not* tasks. To understand how and why shifts in technology and trade affect skill demands and earnings, we therefore need to be able to measure the mapping between skills and tasks and observe the changes in this mapping over time. This presents a substantial measurement challenge since consistent information on job tasks is not commonly collected by representative data sources.

Researchers have taken at least three approaches to resolving this measurement problem. A first is to use occupations as proxies for job tasks. This has a natural appeal since occupations can readily be conceptualized as bundles of tasks that workers are required to perform. Because occupational classification schemes employed by national statistical agencies typically contain hundreds of distinct occupations, however, occupational codes are far too unwieldy in their raw form to serve as task measures. To make this problem manageable, it is necessary to boil down a large set of occupation codes to a lower dimensional object. One method is to aggregate many detailed occupations into a few broad categories, e.g., professional, technical, managerial, clerical, production, service, etc. This is typically straightforward because most occupational schemes are hierarchical. Hence, detailed codes are readily collapsed upward into their parent categories.

A limitation of this approach, however, is that it obscures any similarities in task content that cross broad occupational boundaries. For example, office clerical workers and assembly line machine operators have much in common from the perspective of the task framework: both make extensive use of routine tasks that have high potential for automation. Similarly, both truck drivers and food service workers engage intensively in non-routine manual tasks requiring detailed visual recognition and flexible adaptation to a changing physical environment, tasks that have proven extremely challeng-

ing to automate. Unfortunately, these overlaps among occupations in “task space” are in no way visible from standard occupational classification schemes that group occupations roughly according to the services that they provide (health services, production, analysis, etc.) rather than the tasks that they encompass. Thus, a researcher who wants to identify task commonalities that cross occupational boundaries is forced to make additional subjective judgments. Such judgment calls are generally undesirable in empirical work since they leave an uncomfortable amount of discretion for researchers to shape the data to fit their priors.¹³

To reduce the role of subjectivity in the task categorization step, an alternative approach to grouping occupations manually is to append a set of standardized job descriptors to each occupation. These descriptors can in turn be used to develop task measures. In the U.S., two major data sources, the Dictionary of Occupational Titles (DOT) and its successor, the Occupational Information Network (O*NET), offer sources for job descriptors, and both have been used frequently in empirical work on job tasks.

The primary strength of this method of classifying tasks is that its main inputs, externally supplied job content measures, can be carefully validated by the statistical agencies that supply them. The approach of assigning tasks to occupations using job task descriptors has two major limitations, however, one intrinsic and the other reflecting shortcomings in existing data collection efforts. Fundamentally, assigning task measures to occupations overlooks all heterogeneity in job tasks among individuals within an occupation. It's self-evident that individual worker skills and actual job tasks differ among workers within an occupation, and it seems likely that these within-occupation skill-task assignments are an important component of the overall equilibrium relationship between skills and tasks. Thus, at best, occupation level task measures provide a rough approximation to the microeconomic assignment process.

Perhaps even more importantly, task measures assigned at the occupational level offer a static view of the tasks an occupation comprises. If, as the task framework suggests, job tasks assignments are subject to continual re-optimization, this static perspective is problematic. Under the reasonable assumption that task changes within occupations (the ‘intensive margin’) tend to move in the same direction as task changes made visible by changes in the relative sizes of occupations (the ‘extensive margin’), then measuring job tasks using static measures of occupational content will systematically understate the extent of the task reallocations taking place. Fortunately, this problem can be overcome if occupational job descriptor databases are regularly

¹²Indeed, one simple sense in which the task approach has been productive is by spurring economists to focus attention on shifts in occupational composition—most notably, in the case of employment polarization—rather than maintaining an exclusive focus on the wage returns to skills. It seems self-evident that the changing distribution of occupational assignments—that is, shifts in what people *do* at their jobs—contains information about the evolving nature of production in the economy. But perhaps because occupations do not have a theoretical analogue in the canonical production framework, they have until recently received limited empirical attention from economists (though see Goldin 2002 for a notable exception).

¹³On the other hand, workers' occupational classifications are generally hard-coded in survey and administrative data sets. Hence, using these codes directly is empirically transparent despite its other disadvantages.

refreshed. In the case of the U.S. Dictionary of Occupational Titles, updates were infrequent and appeared to suffer from status quo bias (Miller et al. 1980). In the case of O*NET, the sheer number of distinct occupations—approximately 1,000 in O*NET release 14.0—and the vast quantity of unique scales make it extremely costly to refresh the entire database at high frequency.

Alongside these intrinsic limitations of occupation-level task measures, there are practical difficulties stemming from the design of the job content descriptors provided by the major U.S. data sources (DOT and O*NET). In both data sources, job content measures are often vague, repetitive, and constructed using ambiguous and value-laden scales that are likely to confuse respondents.¹⁴ The DOT, which was first issued by the U.S. Department of Labor in 1939 and last revised in 1991, contains 44 objective and subjective job content scales. When the DOT was replaced by the O*NET in 1998, the complexity of the database increased by an order of magnitude. Version 14.0 of the O*NET database, released in June of 2009, contained 400 separate rating scales, which is almost half as many scales as the number of occupations coded by O*NET. These scales were not of course developed with “task measurement” in mind since the design of DOT and O*NET preceded recent research interest in job tasks. In practice, this means that researchers who wish to use these databases as sources for task measures are essentially *required* to pick and choose among the plethora of scales available, a problem that is much more severe for O*NET than for DOT. Researcher discretion again becomes paramount in this data construction process, and some transparency is inevitably lost. While I have found that task measures distilled from DOT and O*NET can serve as powerful proxies for occupational tasks, I am at best only moderately comfortable with these tools because their complexity and opacity places little discipline on how they are applied and interpreted.

A third approach to task measurement that avoids some of the pitfalls above is to collect job task information directly

¹⁴For example, Item 30 of the Skills questionnaire of O*NET release 14.0 asks respondents to rate the importance of Systems Evaluation, defined as “Identifying measures or indicators of system performance and the actions needed to improve or correct performance, relative to the goals of the system” of Labor Employment and Administration (2009). The O*NET survey anchors for this question include: Level 2, determining why a coworker was unable to complete a task on time; Level 4, understanding why a client is unhappy with a product; and Level 6, evaluating the performance of a computer system of Labor Employment and Administration (2009). As a 2010 National Research Council (Tippins and Hilton, 2010, p. 178) report noted, “it is difficult... to conceive of a well-defined metric that would unambiguously place the level of systems evaluation required for “evaluating the performance of a computer system” either above or below the level of systems evaluation required for “understanding why a client is unhappy with a product.”

from survey respondents alongside other demographic, employment and wage data. Adding task measures directly to worker surveys places no restrictions on the variability of tasks within as well as across occupations, does not inadvertently impose the assumption that occupational tasks are static, and allows task measures to be designed for testing specific hypotheses. There have been several successful efforts in this direction to date. Most significantly, the IAB/BIBB labor force data, collected in 1979, 1984/85, 1991/92, 1998/99 and 2005/06, offer detailed self-reported data on workers’ primary activities at their jobs.¹⁵ To my knowledge, the task measures in the IAB/BIBB were first brought to prominence among economists worldwide by a sardonic article by DiNardo and Pischke (1997), who purported to assess whether pencil use had changed the U.S. wage structure during the 1980s. Subsequent work by Spitz-Oener (2006), Dustmann et al. (2009), Antonczyk et al. (2009), Black and Spitz-Oener (2010), and Gathmann and Schönberg (2010), along with other recent studies that are too numerous to list here, used the job activity questions in the IAB/BIBB data to measure job tasks and explore links between technological change, changes in task inputs, and shifts in wage structure. One striking finding from this body of work is that changes in the task structure of employment in Germany between 1979 and 1999 are primarily accounted for by changes in task input *within* detailed occupations rather than shifts in employment across occupational categories.

The IAB/BIBB data were not, to my knowledge, purpose-built to measure economy-wide levels or over-time changes in work activities (i.e., job tasks). Several more recent data collection efforts explicitly tackle these goals. The British Skills Survey, fielded in 1986, 1992, 1997, 2001 and 2006 by Francis Green and collaborators, has sought to provide consistent measures of skills used in the workplace by surveying workers about their work activities, training requirements, use of technologies, and freedom to exercise discretion in performing their jobs (see Felstead et al. 2007 for a summary of recent findings from this body of work). The survey of Skills, Technology, and Management Practices (STAMP) fielded by Michael Handel provides a detailed cross-sectional view of work activities in the U.S., and has the potential to provide a valuable time-series view of the evolution of these activities provided that the survey is repeated (Handel *forthcoming*). In addition, a survey conducted by the Princeton Data Improvement Initiative (PDII) in 2008 collected a series of experimental workplace task measures for a representative sample of U.S. households

¹⁵Unfortunately, the set of job activity questions used varies substantially across the different survey years. This almost certainly reduces the reliability of the IAB/BIBB data as a source for tracking the evolution of job task inputs in aggregate.

(see Autor and Handel [forthcoming](#)). While this survey will not offer a time-series component, it serves as a useful testbed for exploring the value-added of person level task measures alongside occupation level measures. Most recently, researchers at IAB have fielded a new survey instrument that seeks to rigorously operationalize five major task constructs: analytic tasks, interactive tasks, manual tasks, routine tasks, and autonomy-demanding tasks (Matthes et al. 2012).

One lesson these recent initiatives offer is that designing survey questions that successfully measure routine, codifiable tasks is itself a challenging task. The heart of the difficulty is that what is routine from the perspective of human labor is generally not routine from the perspective of machine automation and vice versa. For example, mopping floors is a mundane repetitive task from the perspective of janitors. The fact that the visual recognition and environmental adaptability required to perform this task pose daunting technical challenges for computer science is a characteristic of current computer technology, *not* a characteristic of the work from a human labor perspective. Consequently, asking workers to assess whether their job tasks are codifiable from the perspective of machine execution—roughly akin to asking truck drivers to compare their skills to the current capabilities of Google’s driverless car—is unlikely to provide reliable information.¹⁶ A more promising approach may be to ask workers concretely about which of many commonplace tasks they regularly perform, and then to apply outside expertise to classify whether these tasks are routine or non-routine from a machine execution perspective. Of course, we should expect that the range of tasks that can be automated will continue to expand outward from the most “routine” core—performing basic mathematical operations, storing and retrieving data—towards tasks requiring ever greater adaptability, judgment and, perhaps eventually, creativity.

¹⁶For example, among the task questions that Michael Handel and I contributed to the PDII survey, we included the following item intended to measure the frequency of routine tasks, “How much of your workday (involves/involved) carrying out short, repetitive tasks? [almost all of the time, more than half the time, less than half the time, almost none of the time, don’t know]” Consistent with expectations, managers and professionals report performing relatively few routine tasks. To our surprise, however, workers in service occupations (e.g., food service, cleaning, security) report spending about as much time performing routine tasks as clerical workers, and transportation workers report substantially more. In retrospect, we should have anticipated exactly this pattern of responses. To the typical delivery driver, security guard or food service worker, it is almost surely the case that their primary job tasks appear quite repetitive—though these tasks don’t appear as such to computer scientists. Thus, the notion of “routineness” that we had in mind in writing the survey question is simply alien to how most survey respondents think about it. (In our partial defense, Handel and I were invited to join the PDII effort shortly before the final survey went into the field, and hence, we did not have the opportunity to pilot the task questions. I nevertheless view our initial effort as naive.)

3.2 The counterfactual problem

Alongside the practical challenge of measurement, the task framework poses two central conceptual challenges that the literature is now beginning to tackle. A first concerns the intrinsic difficulty of identifying credible counterfactuals in a setting where the allocation of labor to sectors (occupations, tasks) is determined by comparative advantage, as in a Roy and Donald (1951) model.¹⁷ To see the problem, consider the question of what is the “market wage” of a worker of a given skill level in the task model above. At one level, the answer is self-evident. In the *equilibrium* of the model, a “law of one price” for skill prevails: there is one wage W_H , W_M and W_L for each of the three skill levels.¹⁸ However, these equilibrium wages are contingent on the *assignment* of skills to tasks. In reality, each worker in the task economy could perform any of the tasks available, and would earn a different wage in each of the three segments of the task space. Moreover, if a non-trivial subset of workers were to reallocate their labor across task categories, this would alter the aggregate wage offered to each skill group, since it would shift the effective supply of skill.

These notional reallocations are out-of-equilibrium thought experiments and, superficially, have little practical relevance.¹⁹ They become relevant, however, when shocks to the economy shift the allocation of skills to tasks—for example, the case discussed above where a task-replacing technological change displaces middle skill workers from a set of tasks they’d previously performed. In a canonical production model with fixed roles for each input, this technological advance would be equivalent to an increase in the

¹⁷See Heckman and Honore (1990), Heckman and Sedlacek (1985), and Heckman and Scheinkman (1987).

¹⁸In an extension to the model, Acemoglu and Autor (2011) consider a setting in which each worker i has an endowment of each of the three skills, l_i, m_i and h_i , only one of which can be used at a time and where the distribution of endowments varies over the population. Each worker therefore chooses which skill to apply to which task to maximize earnings, and hence the wage of worker i is equal to $w_i = \max[W_L l_i, W_M m_i, W_H h_i]$, where W_L, W_M and W_H are in turn determined by the supply of skills to tasks (that is, aggregate supply to each skill group is jointly determined with the wage level). This setting gives rise to both greater wage heterogeneity and richer wage dynamics than the baseline case, and also captures, in my view, an important additional element of realism. When a skilled production worker loses his job due to plant shutdown and takes employment instead as a food service worker, it is plausible to think that he has not only changed job tasks but also changed the skill set used to perform these tasks—for example, switching from using m skills for production tasks to l skills for service tasks. This extension to the model is also consistent with the findings in Neal (1999) and Gathmann and Schönberg (2010) that job changers tend to switch into occupations that are similar in task content to their prior jobs, and when they move further afield by changing industries or by substantially changing job tasks, they tend to experience a wage penalty.

¹⁹Unless, of course, one posits a set of plausible frictions in the economy that lead to some degree of misallocation at any given moment.

relative supply of middle skill labor, which would reduce the middle skill wage and, in most cases, raise the wage of other factors of production—low and high skill labor, in particular—through q-complementarity. In a setting with comparative advantage, however, the effects are more nuanced and contingent. The substitution of capital for labor in performing middle skill tasks reduces the total set of tasks performed by M labor and causes the threshold tasks, I_L^* and I_H^* , to shift outward (falling in the case of I_L^* and rising in the case of I_H^*).²⁰ These shifts impact W_M through three channels. One is the standard diminishing marginal productivity effect seen in the canonical model. A second is a q-complementarity effect, here operating directly on M labor: reflecting their lower cost, the newly automated tasks are used more intensively in production, which in turn raises the marginal product of *all* other tasks, including those performed by M labor. The third operative channel is reallocation: as I_L^* and I_H^* are shifted outward, M labor is reassigned to a set of tasks in which it was initially less productive, thus reducing its marginal product. In short, how a task-replacing technological change (or an increase in own labor supply) affects the earnings of M workers depends importantly on how productive M labor is in alternative uses—that is, in the tasks to which it is *not* initially assigned.

These same forces also determine how the task-replacing technological change affects the relative earnings of H and L workers. We noted above that an M -labor replacing technological change causes I_L^* and I_H^* to shift outward, thus encroaching on the tasks previously performed by H and L labor. What determines the degree of encroachment is the degree of substitutability between M labor and H or L labor at the respective task thresholds, I_L^* and I_H^* . If H labor has strong comparative advantage relative to M labor in performing the tasks just above I_H^* and L labor has comparatively weak comparative advantage relative to M labor in performing tasks just below the I_L^* threshold, then M labor will primarily be allocated *downward* as the M -labor displacing technology is adopted, meaning that W_L falls relative to W_H and wage inequality between high and low skill workers rises. If the pattern of comparative advantage is reversed, then W_L rises relative to W_H and wage inequality falls. Thus, comparative advantage shapes the impact of a shift in technology or skill supplies anywhere in the task distribution on the task assignment, productivity and wage of all other factors.

²⁰Recall that in equilibrium, the cost of performing the I_L^* threshold task with either M or L labor must be equated, and similarly for M and H labor at the I_H^* threshold task. By squeezing M workers into a smaller set of tasks, a middle skill task-replacing technical change initially reduces their marginal productivity, hence W_M . Because it is now cheaper to perform the threshold tasks with M labor than with either L or H labor, I_L^* must shift downward and I_H^* upward until the no arbitrage conditions are again satisfied.

These features add to the realism of the model, but they pose challenges for empirical implementation because the productivity of factors in performing tasks to which they are not currently assigned is inherently hard to estimate. For example, an ambitious study by Costinot and Donaldson (2012) studies how reductions over the course of 160 years in the cost of moving agricultural goods across the U.S. have affected consumer welfare through three distinct mechanisms: direct reductions in transport costs; gains from price convergence across markets; and, most relevant here, gains from reallocation of land parcels to crops that are more valuable when distant markets become accessible (e.g., avocados are now grown in California for consumption in Manhattan). The innovation that makes this last crucial step feasible is that the authors exploit agronomic data to estimate the potential productivity of each parcel of land for producing *all* food crops, not simply those with which it is currently planted. This ingenious approach is not necessarily broadly feasible in other settings, unfortunately. In a recent paper, for example, Hsieh et al. (2012) seek to assess how much the reductions in discrimination against women and minorities over six decades have increased aggregate productivity by increasing allocative efficiency in the labor market (in our terminology, the assignment of skills to tasks). Since no labor market data equivalent to agronomic data are available for estimating counterfactual task productivities by gender and race, the Hsieh et al. estimation is necessarily more dependent on distributional assumptions than is the Costinot-Donaldson study.²¹

3.3 On the “wage return” to tasks²²

A related conceptual and empirical challenge concerns the appropriate use of task measures in wage regressions. A key implication of the theoretical framework above is that holding the task efficiencies (that is, the $\alpha(\cdot)$'s) constant, changes in the market value of tasks should affect the evolution of wages by skill group. In particular, the model predicts that if the relative market price of the tasks in which a skill group holds comparative advantage declines, the relative wage of that skill group should also decline—even if the group reallocates its labor to a different set of tasks (i.e., due to the change in its comparative advantage).

Critical to this prediction is the distinction made between the wages paid to a skill group and the wages paid to a given

²¹Over the short and medium term, changes in supplies and technologies will generally lead to evolutionary rather than revolutionary reallocations of factors. If so, it may be sufficient for many leading empirical questions to be able to estimate the distribution of potential factor productivities in a *neighborhood* nearby to the current equilibrium allocation rather than across all possible scenarios.

²²This discussion of the use of task measures in wage regressions draws from Acemoglu and Autor (2011), and some text is repeated from that source.

task—a distinction that is meaningful because the assignment of skills to tasks is endogenous. To see the implications of this distinction, consider a technological change that raises the productivity of high skill workers in all tasks (e.g., an increase in A_H). The model implies that this would expand the set of tasks performed by high skill workers (i.e., lowering I_H), so that some tasks formerly performed by medium skill workers would now be performed by high skill workers instead. Thus, relative wages paid to workers performing these (formerly) middle skill tasks would actually increase since they are now performed by the more productive high skill workers.²³ But crucially, the analysis in Acemoglu and Autor (2011) also proves that the relative wage of medium skill workers, who were formerly performing these tasks, would fall.

One important implication of this set of observations is that the intuitively appealing approach of regressing wages on current tasks performed by workers to estimate the “returns to tasks” may generate potentially misleading results. The appeal of estimating task-wage regressions is almost reflexive for labor economists steeped in the Mincerian tradition of regressing log wages on years of schooling and potential experience. At the heart of the Mincerian approach is the observation that a worker’s human capital, proxied by schooling, is a quasi-fixed stock variable, determined prior to labor market entry. In this setting, a log earnings regression can recover the market rental rate of human capital (i.e., the return to education)—at least under certain stringent assumptions.

By contrast, the set of tasks that a worker performs on the job is an endogenous state variable that is simultaneously determined by the worker’s stock of human capital and the contemporaneous productivity of the tasks that human capital could accomplish. This implies that task assignments are themselves a function of the current wage distribution, and hence a regression of wages on job tasks (rather than on worker skills) will be inherently difficult to interpret. As an extreme example of this pitfall, note as above that the relative wage paid to a task can move in the *opposite* direction from the relative wage paid to the skill group that initially performed that task. By contrast, the relative wage paid to a given skill group always moves in the same direction as its comparative advantage—that is, a technological change that increases the productivity of a skill group necessarily raises its relative wage. This suggests that even in a task setting, the Mincerian approach of regressing wages on skills

may be more informative about the wage structure than the alternative of regressing wages on tasks.

A growing literature is now grappling with these challenges. A recent working paper by Firpo et al. (2011), for example, develops an innovative method for measuring the impact of changing task prices on wage structure. Using a simple statistical model of occupational wage setting, their work predicts that occupations that are specialized in tasks that have declining market value should see a reduction in both mean occupational wages and the *variance* of occupational wages, and vice versa for tasks with rising prices. This latter (variance) effect stems from the interaction between a falling task price and a fixed distribution of task efficiencies within an occupation; as the market value of a given task falls, the range of wages paid to workers with differing productivities in that task compresses along with it.

One issue that needs further study in the Firpo, Fortin and Lemieux approach is that changes in task prices will presumably lead to changes in the self-selection of workers into occupations—as is implied by the comparative advantage model and, more generally, by any model in which workers make maximizing choices. These endogenous changes in occupational entry and exit should *also* affect occupational wage means and variances. At present, the approach taken by Firpo, Fortin and Lemieux is similar to the Mincer regression in that it implicitly treats the skill allocations (here, the assignment of skills to tasks) as predetermined. This simplifies the predictions of the statistical model but imposes an unknown cost in terms of economic realism.

4 Some concluding concerns

Revealed preference demonstrates that I am optimistic about the potential of the “task approach” to labor markets to offer insights into the interactions among skill supplies, technological capabilities, and potentially trade and offshoring opportunities, in shaping the demand for skills, the assignment of skills to tasks, and the evolution of wages. Rather than recapitulate the sources of my enthusiasm, let me conclude by offering three recommendations to researchers working in this literature.

A first concerns task definitions. Following Autor et al. (2003), researchers have often used a rough taxonomy of three broad task groups: routine cognitive and manual tasks, which are readily codifiable for purposes of machine performance; abstract analytical and managerial tasks, which may require creativity, hypothesis formation, problem solving or persuasion; and non-routine manual tasks (or, simply, manual tasks), which may require physical flexibility and adaptability, visual recognition, or non-scripted communications. Commencing with the influential work of Blinder (2009), economists have also focused on another dimension

²³Nor is this notion far-fetched. Skill levels in production and clerical occupations, as measured by the college employment or wage-bill share, have risen as employment in these occupations has declined (Autor et al. 2003). A plausible interpretation of this pattern is that educated workers have comparative advantage in the set of non-routine tasks in these occupations that remain.

of task content: “offshorability,” describing tasks that may be performed in a remote location without substantial quality degradation.²⁴

While these four task attributes—routine, abstract, manual, offshorable—are broadly distinct, there are important overlaps among them. For example, many routine codifiable tasks (e.g., performing calculations, checking a document for spelling errors) can potentially be automated or offshored. On the other hand, many tasks that are eminently offshorable are clearly non-routine—for example, providing technical support to frustrated computer users, interpreting medical x-rays, or taking drive-through orders at roadside McDonalds.²⁵

How should these overlaps affect task classification? A rudimentary rule of thumb is that task classification schemes should provide separate scales for distinct task attributes. In practice, however, researchers frequently employ broad classification schemes that collapse distinctions among attributes. For example, in constructing measures of task “offshorability,” it is commonplace for researchers to classify tasks as offshorable if they are highly codifiable, and hence can be performed without direct supervision, *or* if they do not require face-to-face contact and hence suffer little quality degradation when performed remotely. This broad definition of offshorability is problematic. As Blinder and coauthors have emphasized in their discussions of offshorability (cf. Blinder 2009; Blinder and Krueger forthcoming), the key attribute that makes a task potentially “offshorable” is that it does not depend on face-to-face worker contact or close proximity between worker and customer.²⁶ This requirement does not *exclude* many non-routine tasks—such as interpreting x-rays, or providing technical support—nor does it *include* all routine tasks, such as generating web pages dynamically in response to user queries. While many routine task-intensive activities appear to be potentially suited for offshoring, what makes them suitable is that

they are not dependent on face-to-face interactions, not that they are routine per se.²⁷

In a similar vein, researchers sometimes classify tasks as *non-offshorable* because they require workers to exercise judgment, think creatively, or engage in problem-solving. Yet, it cannot be the case that overseas workers are unable to exercise judgment or perform higher cognitive functions! Clearly, what makes these particular non-routine tasks difficult to offshore is that they are typically performed by workers in management, leadership or decision-making roles—and these roles in turn demand face-to-face interactions. Here again, the offshorability question turns on quality degradation due to distance, not on a task’s routineness or non-routineness.

These observations raise the concern that the emerging task literature may be significantly inside the frontier of what is feasible in terms of precise terminology and consistent measurement. Addressing these shortcomings should therefore be a high priority on the research agenda since it is difficult for a field to advance when researchers lack shared definitions of core constructs or consistent measures of these constructs.

To advance this agenda, my second recommendation to researchers is to use, re-use, recycle, replicate, repeatedly apply and attempt to converge upon a shared and standardized set of task measures.²⁸ Although empiricists may have a preference for their own purpose-built task measures, social science will be better served if, at a minimum, researchers use “off the shelf” measures alongside these alternatives. This practice is critical for illuminating when the potentially divergent conclusions reached by competing studies reflect disagreements about substance or instead discrepancies in measurement. While there is surely room for improvement on both fronts, a lack of clarity on either hinders intellectual progress.

As one small step towards speeding the process of convergence, I have placed all task measures used in my published papers—along with associated crosswalks and supporting documents—on my MIT homepage for download.²⁹ If other researchers are similarly interested in contributing to this public good, I am pleased to host and maintain a web page that provides access to data sets and documentation developed by various task projects for use by the broader research community.

Finally, let me remark on the difficulties on extrapolating from theory to policy. One of the ironies of the computer era

²⁴One may ask what distinguishes offshorability from traditional trade in goods. Perhaps the best answer is that offshoring refers primarily to service tasks rather than physical goods. Many of these services are non-storable and hence need to be produced in real time at the point at which they are used (e.g., teaching, technical support, person-to-person sales). Thus, a valid measure of “offshorability” describes which of these tasks must be performed in the location where they are consumed and which can be delivered from a distance.

²⁵Surprisingly, the voice heard through the loudspeaker when you place your drive-through order may not belong to a worker in the restaurant.

²⁶How could a task not require face-to-face contact but still need to be performed nearby to a customer site? Package delivery is one such example. United Parcel Service (UPS) does not require face-to-face contact with package recipients, but it does need to drive up to their houses to deliver packages.

²⁷Having made these distinctions, researchers may draw different conclusions about whether the widely observed phenomenon of employment polarization is due to automation, offshoring or some third set of factors. But this is a question of cause-and-effect, not measurement.

²⁸It is regrettably the case that there are almost as many distinct task classifications as there are papers in the task literature.

²⁹<http://web.mit.edu/dautor/www>.

is that information technology has increased labor's comparative advantage in traditionally low skill service tasks, such as food preparation and personal care, relative to its value in traditional middle skill tasks such as computation, information processing, and repetitive production activities. One inference that economists and policymakers are tempted to draw from this trend is that advanced countries should give up on "middle skill" education because there is no future for middle skill jobs. This conclusion is mistaken for multiple reasons. A first is that education is cumulative; students cannot attain high skills (e.g., proving theorems) without first mastering middle skills (e.g., arithmetic). Since it is difficult to forecast which students might ultimately succeed in the higher echelons of education, there is substantial option value in investing universally in students' middle skills. The efficiency case for these investments is complemented by an equity case. Choosing to *not* universally invest in students' "middle skills" would imply foreclosing the economic horizons of many citizens at an early age—an idea that few democratic societies would want to embrace.

Second, my reading of the evidence is that middle skill jobs are not slated to disappear. While many middle skill *tasks* are susceptible to automation, many middle skill *jobs* demand a mixture of tasks from across the skill spectrum. To take one prominent example, medical paraprofessional positions—radiology technicians, phlebotomists, nurse technicians, etc.—are a numerically significant and rapidly growing category of relatively well-remunerated, middle skill occupations. While these para-professions do not require a college degree, they do demand one to two years of post-secondary vocational training. Significantly, mastery of "middle skill" mathematics, life sciences, and analytical reasoning is indispensable for success in this training.³⁰

Why are these middle skill jobs likely to persist and, potentially, to grow? My conjecture is that many of the tasks currently bundled into these jobs cannot readily be unbundled—with machines performing the middle skill tasks and workers performing the residual—without a substantial drop in quality. Consider, for example, the com-

monplace frustration of calling a software firm for technical support only to discover that the support technician knows nothing more than what is on his or her computer screen—that is, the technician is a mouthpiece, not a problem solver. This example captures one feasible division of labor: machines performing routine technical tasks, such as looking up known issues in a support database, and workers performing the manual task of making polite conversation while reading aloud from a script. But this is not generally a productive form of work organization because it fails to harness the complementarities between technical and interpersonal skills. Stated in positive terms, routine and non-routine tasks will generally coexist within an occupation to the degree that they are complements—that is, the quality of the service improves when the worker combines technical expertise and human flexibility.³¹

This reasoning suggests that many of the middle skill jobs that persist in the future will combine routine technical tasks with the set of non-routine tasks in which workers hold comparative advantage—interpersonal interaction, flexibility, adaptability and problem-solving.³² Medical para-professions are one leading example of this virtuous combination, but this example is not a singularity. This broad description also fits numerous skilled trade and repair occupations—plumbers, builders, electricians, HVAC installers, automotive technicians—marketing occupations, and even modern clerical occupations that provide coordination and decision-making functions rather than simply typing and filing. Indeed, even as some formerly middle skill occupations are stripped of their routine technical tasks and arguably deskilled—for example the stockbroking occupation—other formerly high-end technical occupations are made accessible to workers with less esoteric technical mastery, for example, the nurse practitioner occupation that increasingly performs diagnosing and prescribing tasks in lieu of physicians. I expect that a significant stratum of middle skill, non-college jobs combining specific vocational skills with foundational middle skills—literacy, numeracy, adaptability, problem-solving and common sense—will persist in coming decades.

The economics profession is very far from a full understanding of the interactions among rising worker skills, advancing technology, improvements in offshoring and trade opportunities, and shifting consumer demands in determining the division of labor, the growth of aggregate productivity, and the level and inequality of earnings within and

³⁰ Simultaneously, it becomes more critical than ever to improve workers' capabilities to add value in manual and service tasks. While one may legitimately be skeptical that education can do much to improve productivity in labor-intensive, technologically-lagging tasks, it is worth noting that workers with post-secondary education appear to earn more in essentially *all* walks of life—including mundane service occupations—than do those with a high school or lesser education (Carnevale et al. 2011). Barring the unlikely case that these wage differences entirely reflect self-selection on unobserved ability into post-secondary education rather than post-secondary education's value added, these earnings differentials suggest that increasing the skills and capabilities of workers that will perform the lower skill tasks in the economy may be productive in aggregate and may have the secondary benefit of containing the growth of inequality.

³¹ Lawrence Katz memorably titles workers who virtuously combine technical and interpersonal tasks as "the new artisans" (see Friedman 2010).

³² In general, these same demands for interaction frequently privilege face-to-face interactions over remote performance, meaning that these same middle skill occupations may have relatively low susceptibility to offshoring.

between skill groups. The “task approach” to labor markets does not come close to offering a solution to this vast intellectual puzzle. It may however serve to put a few puzzle pieces into their rightful locations. If so, this will mark progress towards a worthy goal.

Executive summaries

The canonical production function found in economic models is anthropomorphic—it’s built to resemble humans. It features a role for labor and a role for machinery (capital), and, in general, these roles are distinct. This is not illogical—after all, these are separate inputs. But what precisely is distinctive about the role of each input is left opaque, and the nature of the interactions among them is highly constrained. In particular, capital is either a complement or a substitute for labor, different types of labor are either complements or substitutes for one another, and these roles are essentially fixed. Changes in the supply of each input can of course affect marginal products, but each factor’s “purpose” in the production function is both distinct and static.

These restrictions stem from the fact that the canonical production function implicitly equates two distinct aspects of production. One aspect is which factors are used as inputs, e.g., capital, high skill labor, low skill labor. The other is what services these factors provide. In the canonical setup, a factor’s identity and its role in the production function are synonymous. In reality, however, the boundary between “labor tasks” and “capital tasks” in production is permeable and shifting. We observe numerous instances in which technological advances allow machinery to substitute for labor in performing a specific set of tasks—for example, dispensing boarding passes and assigning seats at airport check-in counters in lieu of gate agents—while simultaneously complementing labor in other tasks, for example, allowing gate agents to rapidly identify alternate flights and issue new tickets. This evolving division of labor has a clear economic logic: novel tasks—those demanded by new products, techniques, or services—are often assigned first to workers because workers are flexible and adaptive. As these tasks are formalized and codified, they become fallow for automation since machinery typically has a cost advantage over human labor in rote execution of repetitive tasks.

A growing body of literature argues that the shifting allocation of tasks between capital and labor—and between domestic and foreign labor—has played a key role in reshaping the structure of labor demand in industrialized countries in recent decades. Analyzing this phenomenon is difficult within the canonical production function framework, however, because the assignment of tasks to labor and capital is, for most conceptual purposes, static. To overcome this conceptual hurdle, it is valuable to consider a setting where the

equilibrium mapping between production tasks and productive factors responds dynamically to shifting economic and technological forces. The “task approach” to labor markets offers one promising, though far from complete, set of conceptual tools for studying these forces.

In this brief essay, I first sketch a formal model, developed by Acemoglu and Autor (2011), that offers a flexible and tractable approach for analyzing the interactions among skill supplies, technological capabilities, and potential trade and offshoring opportunities in shaping the demand for factors, the assignment of factors to tasks, and the productivity and hence the wage of each factor. This model removes what we view as an artificial set of distinctions typically made between labor, capital, and trade (or offshoring) in the canonical production model, where different factors of production play distinct and often incommensurate roles. In this alternative model, the fundamental units of production are job tasks, which are combined to produce output. Tasks can be supplied by domestic labor, foreign labor, or capital, the capabilities of which may change over time. Consequently, the assignment of factors to tasks in this model is determined by comparative advantage; in equilibrium, only the least-cost factor is assigned to any given task.

The second part of the essay briefly discusses some of the key conceptual and practical challenges that researchers face in bringing the “task approach” to the data. The literature is already grappling with these issues, as I note below, and I’m cautiously optimistic that progress will be rapid. I end by highlighting some of the pitfalls that researchers may encounter in applying the task approach, most prominently, loose and inconsistent task definitions that tend to sow confusion, and the beguiling temptation for researchers to confidently forecast the “future of work” based only on the very modest understanding that we have developed to date.

Kurzfassungen

Die gängige Produktionsfunktion, die in ökonomischen Modellen angewandt wird, ist anthropomorph – sie ist so konstruiert, dass sie menschliches Verhalten abbildet. Sie weist sowohl der Arbeit als auch dem Kapital eine Rolle zu – und im Allgemeinen sind diese Rollen verschieden. Das ist nicht unlogisch, da es sich um verschiedene Inputfaktoren handelt. Jedoch bleibt unklar, was genau die jeweilige Rolle der beiden Inputfaktoren ausmacht. Zudem sind die Interaktionen zwischen beiden in diesem Modell hochgradig eingeschränkt. Insbesondere ist Kapital entweder ein Komplement oder ein Substitut für Arbeit, und unterschiedliche Arbeitstypen sind entweder Komplemente oder Substitute füreinander, wobei diese Rollen grundsätzlich festgelegt sind. Natürlich können Änderungen

des Angebots jedes Inputfaktors die Grenzprodukte beeinflussen, aber der „Zweck“ jedes Inputfaktors in der Produktionsfunktion ist sowohl distinkt als auch statisch.

Diese Restriktionen leiten sich aus der Tatsache ab, dass die gängige Produktionsfunktion implizit zwei verschiedene Aspekte der Produktion gleichsetzt. Ein Aspekt ist, welche Faktoren als Input eingesetzt werden, z.B. Kapital, hochqualifizierte Arbeit, gering-qualifizierte Arbeit. Der andere ist, welchen Nutzen diese Faktoren stiften. Nach dem gängigen Ansatz sind die Faktor-Identität und seine Rolle in der Produktionsfunktion synonym. In der Realität jedoch ist die Grenze zwischen „Arbeitsfunktionen“ (*labor tasks*) und „Kapitalfunktionen“ (*capital tasks*) in der Produktion durchlässig und veränderlich. Wir beobachten unzählige Beispiele, wo bei bestimmten Tätigkeiten Arbeit aufgrund des technischen Fortschritts durch Kapital ersetzt wird – zum Beispiel bei der Ausgabe von Flugtickets und von Sitzplätzen an Check-in-Schaltern in Flughäfen – während gleichzeitig bei anderen Aufgaben und Tätigkeiten ein komplementäres Verhältnis zwischen Arbeit und Kapital entsteht, etwa indem das Flughafenpersonal durch die Technik in die Lage versetzt wird, rasch Alternativflüge zu finden und neue Tickets auszustellen. Diese fortschreitende Zweiteilung der Arbeit hat eine klare ökonomische Logik: Neuartige Aufgaben – wie sie neue Produkte, Technologien oder Dienstleistungen erfordern – werden oftmals zuerst Arbeitnehmern zugeteilt, da Menschen flexibel und anpassungsfähig sind. Sobald sie stärker formalisiert und kodifiziert sind, sind sie reif für die Automatisierung, da Maschinen bei routinemäßiger Ausführung sich wiederholender Tätigkeiten typischerweise Kostenvorteile gegenüber menschlicher Arbeit haben.

Ein wachsender Teil der Literatur argumentiert, dass die veränderliche Allokation der Aufgaben zwischen Kapital und Arbeit – sowie zwischen heimischer und ausländischer Arbeit – eine Schlüsselrolle in der Neuordnung der Struktur der Arbeitsnachfrage der industrialisierten Länder in den letzten Jahrzehnten gespielt hat. Jedoch ist die Analyse dieses Phänomens im Rahmen der gängigen Produktionsfunktion schwierig, weil dort die Aufgabenverteilung auf Arbeit und Kapital, aus überwiegend konzeptionellen Gründen, statisch ist. Um dieses konzeptionelle Hindernis zu überwinden ist es wichtig, eine Situation zu betrachten, wo die Gleichgewichtsfunktion zwischen Produktionssaufgaben und Produktionsfaktoren dynamisch auf veränderliche ökonomische und technologische Kräfte reagiert. Der „TASK-Ansatz“ für Arbeitsmärkte bietet ein vielversprechendes (wenn auch keineswegs vollständiges) Set von konzeptionellen Werkzeugen, um diese Phänomene zu studieren.

In diesem knappen Essay beschreibe ich zuerst ein formales Modell, entwickelt von Acemoglu and Autor (2011), welches einen flexiblen und handhabbaren Ansatz

bietet, um die Interaktionen zwischen Qualifikationsangebot, technologischen Fähigkeiten, etwaigen Handels- und „Offshoring“-Möglichkeiten bei der Faktornachfrage, der Faktorzuzuweisung auf die Aufgaben, sowie der Produktivität und damit der Entlohnung jedes Faktors zu analysieren. Die künstlichen Unterscheidungen, welche typischerweise zwischen Arbeit, Kapital und Handel (oder „Offshoring“) im gängigen Produktionsmodell gemacht werden, in dem verschiedene Produktionsfaktoren unterschiedliche und oftmals inkommensurable Funktionen übernehmen, werden durch dieses alternative Modell beseitigt. In diesem alternativen Modell sind die grundlegenden Einheiten der Produktion Arbeitsanforderungen (*job tasks*), die kombiniert werden, um einen Output zu produzieren. Diese Aufgaben können durch heimische Arbeit, ausländische Arbeit oder Kapital erledigt werden, wobei sich die jeweiligen Fähigkeiten über die Zeit verändern können. Daraus folgt, dass in diesem Modell die Faktorverteilung auf die Aufgaben vom komparativen Vorteil bestimmt wird; im Gleichgewicht ist einer gegebenen Aufgabe nur der günstigste Faktor zugeteilt.

Der zweite Teil des Essays diskutiert kurz einige der wichtigsten konzeptionellen und praktischen Herausforderungen, die Forscher dabei haben, den TASKS-Ansatz empirisch zu überprüfen. Die Literatur setzt sich wie bereits erwähnt mit diesen Fragen auseinander und ich bin verhalten optimistisch, dass sie schnelle Fortschritte machen wird.

Abschließend hebe ich einige der Fallstricke hervor, denen sich Forscher bei der Anwendung des TASKS-Ansatzes gegenüber sehen. Dies gilt insbesondere für die vagen und inkonsistenten Aufgabendefinitionen, die tendenziell Verwirrung stiften sowie die trügerische Versuchung für Forscher, die „Zukunft der Arbeit“ lediglich aufgrund der sehr bescheidenen Erkenntnisse, die wir bis heute entwickelt haben, vorhersagen zu wollen.

Acknowledgements Autor acknowledges funding from the National Science Foundation (grant SES-1227334). I thank Joachim Möller and Brendan Price for valuable comments that improved this essay.

References

- Acemoglu, D., Autor, D.H.: Skills, tasks and technologies: implications for employment and earnings. In: Handbook of Labor Economics, vol. 4B, pp. 1043–1171 (2011)
- Acemoglu, D., Zilibotti, F.: Productivity differences. *Q. J. Econ.* **116**(2), 563–606 (2001)
- Antonczyk, D., DeLeire, T., Fitzenberger, B.: Polarization and rising wage inequality: comparing the U.S. and Germany. *IZA Discuss. Pap.* 4842 (2010)
- Autor, D.H., Dorn, D.: The growth of low skill service jobs and the polarization of the U.S. labor market. *Am. Econ. Rev.*, forthcoming
- Autor, D.H., Handel, M.J.: Putting tasks to the test: human capital, job tasks and wages. *J. Labor Econ.*, forthcoming
- Autor, D.H., Levy, F., Murnane, R.J.: The skill content of recent technological change: an empirical exploration. *Q. J. Econ.* **118**(4), 1279–1333 (2003)

- Autor, D.H., Katz, L.F., Kearney, M.S.: The polarization of the U.S. labor market. *Am. Econ. Rev.* **96**(2), 189–194 (2006)
- Autor, D.H., Katz, L.F., Kearney, M.S.: Trends in U.S. wage inequality: revising the revisionists. *Rev. Econ. Stat.* **90**(2), 300–323 (2008)
- Bessen, J.E.: Was mechanization de-skilling? The origins of task-biased technical change. Boston University School of Law Working Paper, pp. 11–13 (2011)
- Black, S.E., Spitz-Oener, A.: Explaining women’s success: technological change and the skill content of women’s work. *Rev. Econ. Stat.* **92**(1), 187–194 (2010)
- Blinder, A.S.: How many US jobs might be offshorable? *World Econ.* **10**(2), 41–78 (2009)
- Blinder, A.S., Krueger, A.B.: Alternative measures of offshorability: a survey approach. *J. Labor Econ.* (forthcoming)
- Carnevale, A.P., Rose, S.J., Cheah, B.: The college payoff: education, occupations, lifetime earnings. Technical report, Georgetown University Center on Education and the Workforce (2011)
- Costinot, A., Vogel, J.: Matching and inequality in the World Economy. *J. Polit. Econ.* **118**(4), 747–786 (2010)
- Costinot, A., Donaldson, D.: How Large Are the Gains from Economic Integration? Theory and Evidence from U.S. Agriculture, 1880–2002 (2012)
- DiNardo, J.E., Pischke, J.-S.: The returns to computer use revisited: have pencils changed the wage structure too? *Q. J. Econ.* **112**(1), 291–303 (1997)
- Dustmann, C., Ludsteck, J., Schönberg, U.: Revisiting the German wage structure. *Q. J. Econ.* **124**(2), 843–881 (2009)
- Fackler, M.: Nissan and others add factories in emerging markets. *New York Times*, 8 June (2008)
- Felstead, A., Gallie, D., Green, F., Zhou, Y.: Skills at Work, 1986 to 2006. ESRC Centre on Skills, Knowledge and Organisational Performance, Cambridge (2007)
- Firpo, S., Fortin, N.M., Lemieux, T.: Occupational tasks and changes in the wage structure. *IZA Discuss. Pap.* 5542 (2011)
- Friedman, T.L.: The election that wasn’t. *New York Times*, 23 Oct. (2010)
- Gathmann, C., Schönberg, U.: How general is human capital? A task-based approach. *J. Labor Econ.* **28**(1), 1–49 (2010)
- Goldin, C.: A pollution theory of discrimination: male and female differences in occupations and earnings. NBER Work. Pap. 8985 (2002)
- Goldin, C., Katz, L.F.: The origins of technology-skill complementarity. *Q. J. Econ.* **113**(3), 693–732 (1998)
- Goos, M., Manning, A.: Lousy and lovely jobs: the rising polarization of work in Britain. *Rev. Econ. Stat.* **89**(1), 118–133 (2007)
- Goos, M., Manning, A., Salomons, A.: Explaining job polarization in Europe: the roles of technology and globalization. *CEP Discuss. Pap.* dp1026 (2011)
- Handel, M.J.: What do people do at work: a profile of U.S. jobs from the survey of workplace skills, technology, and management practices (STAMP). Technical report, OECD, forthcoming
- Heckman, J.J., Honore, B.E.: The empirical content of the Roy model. *Econometrica* **58**(5), 1121–1149 (1990)
- Heckman, J.J., Sedlacek, G.: Heterogeneity, aggregation, and market wage functions: an empirical model of self-selection in the labor market. *J. Polit. Econ.* **93**(6), 1077–1125 (1985)
- Heckman, J.J., Scheinkman, J.: The importance of bundling in a Gorman-Lancaster model of earnings. *Rev. Econ. Stud.* **54**(2), 243–255 (1987)
- Hsieh, C.-T., Hurst, E., Jones, C.I., Klenow, P.J.: The Allocation of Talent and U.S. Economic Growth (2012)
- Katz, L.F., Goldin, C.: *The Race Between Education and Technology*. Harvard University Press, Cambridge (2008). Number March 2007
- Katz, L.F., Margo, R.A.: *Technical Change and the Relative Demand for Skilled Labor: the United States in Historical Perspective* (2012)
- Levy, F., Murnane, R.J.: *The New Division of Labor: How Computers Are Creating the Next Job Market*. Princeton University Press, Princeton (2004)
- Matthes, B., Christoph, B., Janik, F., Ruland, M.: Collecting information on job tasks an instrument to measure tasks required at the workplace in a multi-topic survey. Technical report, November, IAB 2012
- Miller, A.R., Treiman, D.J., Cain, P.S., Roose, P.A. (eds.): *Work, Jobs and Occupations: A Critical Review of the Dictionary of Occupational Titles*. National Academy Press, Washington (1980)
- Neal, D.: The complexity of job mobility among young men. *J. Labor Econ.* **17**(2), 237–261 (1999)
- U.S. Department and Training Administration of Labor Employment Data Dictionary: O*NET 14.0 data dictionary. Technical report, National Center for *NET Development (2009)
- Roy, A.D.: Some thoughts on the distribution of earnings. *Oxf. Econ. Pap.* **3**(2), 135–146 (1951)
- Spitz-Oener, A.: Technical change, job tasks, and rising educational demands: looking outside the wage structure. *J. Labor Econ.* **24**(2), 235–270 (2006)
- Tippins, N.T., Hilton, M.L. (eds.): *A Database for a Changing Economy: Review of the Occupational Information Network (O*NET)*, committee ed. The National Academies Press, Washington (2010)

David H. Autor is a Professor and Associate Department Head of the Massachusetts Institute of Technology Department of Economics. He is also a Faculty Research Associate of the National Bureau of Economic Research and Editor in Chief of the *Journal of Economic Perspectives* (published by the American Economic Association), and has served on the Board of Editors at the *American Economic Journal*, *Applied Economics* and the *Journal of Labor Economics*. Autor received a B.A. in Psychology from Tufts University in 1989 and a Ph.D. in Public Policy at Harvard University’s Kennedy School of Government in 1999. His current fields of specialization include labor market impacts of technological change and globalization, earnings inequality, disability insurance and labor supply, and temporary help and other intermediated work arrangements. Autor is the recipient of an NSF CAREER award for his research on labor market intermediation, the Alfred P. Sloan Foundation Fellowship, and the Sherwin Rosen Prize in 2008 for outstanding contributions in the field of Labor Economics. In 2012, he was elected to the American Academy of Arts and Science, and he is also a Fellow of the Society of Labor Economists and was a recipient of both the John T. Dunlop Outstanding Scholar Award given by the Labor and Employment Relations Association (2006) and MIT Undergraduate Economics Association Teaching Award (2005). Autor is also Co-director of the MIT School Effectiveness and Inequality Initiative (SEII), and Associate Director of the NBER Disability Research Center, which is funded by the Social Security Administration. He is also a Research Associate of the IZA in Bonn, the IAB in Nuremberg, and the Institute for Research on Poverty at the University of Wisconsin, Madison. Prior to obtaining his Ph.D., Autor spent three years directing efforts in San Francisco and South Africa to teach computer skills to economically disadvantaged children and adults.