

A Service of

ZBW

Leibniz-Informationszentrum Wirtschaft Leibniz Information Centre for Economics

Ocampo Díaz, Sergio

Working Paper A task-based theory of occupations with multidimensional heterogeneity

CHCP Working Paper, No. 2022-02

Provided in Cooperation with:

Centre for Human Capital & Productivity (CHCP), Department of Economics, University of Western Ontario

Suggested Citation: Ocampo Díaz, Sergio (2022) : A task-based theory of occupations with multidimensional heterogeneity, CHCP Working Paper, No. 2022-02, The University of Western Ontario, Centre for Human Capital and Productivity (CHCP), London (Ontario)

This Version is available at: https://hdl.handle.net/10419/286701

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



WWW.ECONSTOR.EU

Western University

Scholarship@Western

Centre for Human Capital and Productivity. CHCP Working Papers

Economics Working Papers Archive

2022

2022-2 A Task-Based Theory of Occupations with Multidimensional Heterogeneity

Sergio Ocampo

Follow this and additional works at: https://ir.lib.uwo.ca/economicscibc

Part of the Economics Commons

A Task-Based Theory of Occupations with Multidimensional Heterogeneity

by

Sergio Ocampo

Working Paper #2022-02

January 2022



Centre for Human Capital and Productivity (CHCP)

Working Paper Series

Department of Economics Social Science Centre Western University London, Ontario, N6A 5C2 Canada

A Task-Based Theory of Occupations with Multidimensional Heterogeneity^{*}

Sergio Ocampo[†]

Western University, Department of Economics

January 25, 2022

Abstract

I develop an assignment model of occupations with multidimensional heterogeneity in production tasks and worker skills. Tasks are distributed continuously in the skill space, whereas workers have a discrete distribution with a finite number of types. Occupations arise endogenously as bundles of tasks optimally assigned to a type of worker. The model allows us to study how occupations respond to changes in the economic environment, making it useful for analyzing the implications of automation, skill-biased technical change, offshoring, and worker training. Using the model, I characterize how wages, the marginal product of workers, the substitutability between worker types, and the labor share depend on the assignment of tasks to workers. I introduce automation as the choice of the optimal size and location of a mass of identical robots in the task space. Automation displaces workers by replacing them in the performance of tasks, generating a cascading effect on other workers as the boundaries of occupations are redrawn.

JEL: C78, E24, J23, J24, J31.

Key Words: Occupations, tasks, automation, assignment, skill mismatch.

^{*}I am most grateful to Fatih Guvenen, Jeremy Lise, Loukas Karabarbounis, and David Rahman for their guidance and continuous encouragement. I also want to thank Daron Acemoglu, Anmol Bhandari, Serdar Birinci, Jonathan Eaton, Emmanuel Farhi, Paweł Gola, Juan Herreño, Gueorgui Kambourov, Burhan Kuruscu, Ilse Lindenlaub, Lance Lochner, Keler Marku, Rory McGee, Ellen McGrattan, Emily Moschini, Pascual Restrepo, Baxter Robinson, Sergio Salgado, Kjetil Storesletten, Kurt See, Dominic Smith, Aaron Sojourner, Nancy Stokey, David Wiczer, and participants at seminars at various universities and the 2019 meeting of the Society for Economic Dynamics. All errors are my own.

[†]Email: socampod@uwo.ca; Web: https://sites.google.com/site/sergiocampod/.

Occupations undergo constant change, as evidenced by long-run changes in their skill content (e.g., Spitz-Oener, 2006; Atalay, Phongthiengtham, Sotelo, and Tannenbaum, 2020; Cavounidis, Dicandia, Lang, and Malhotra, 2021). These changes are felt by workers across the earnings distribution, from plant operators facing industrial robots and the secular decline in routine-based occupations to engineers facing the introduction of new software.¹ However, our understanding of how these changes affect the skill premia and workers' productivity and skill requirements is hampered by the lack of an endogenous theory of occupations, where the task composition of occupations evolves in response to changes in the economic environment—like the introduction of information technologies and automation.

In this paper, I develop a multidimensional assignment model of tasks to workers that explicitly allows for changes in the set of tasks that compose an occupation. In doing so, I endogenize the bundling of tasks into occupations, which determines the skill premia and workers' productivity and compensation. The endogenous response of occupations in a multidimensional setup allows me to study the implications of automation, skill-biased technical change, workers' skill upgrading, and other economic changes. By levering on results from optimal transport theory (Villani, 2009; Galichon, 2016), I keep the analysis tractable despite the challenges posed by incorporating multidimensional attributes of workers and tasks into the assignment model.²

In the model, production requires performing a collection of tasks, each generating a taskspecific output. The problem of a production unit (i.e., a plant) is to maximize production by assigning tasks to workers. Both workers and tasks are heterogeneous along multiple dimensions as in Lindenlaub (2017). Workers differ in the skills they possess (e.g., manual, cognitive, social), and tasks differ in the skills that are involved in performing them. In particular, workers' productivity depends on the degree of mismatch between their skills and

¹See Acemoglu and Restrepo (2020) on the effects of industrial robots, Autor, Levy, and Murnane (2003) and Caunedo, Keller, and Shin (2021) on routine-based occupations, and Martinez and Moen-Vorum (2021) on the automation potential of abstract occupations.

²The relevance of multiple types of skills in determining labor market outcomes of individuals has been long recognized. See Willis and Rosen (1979); Heckman and Sedlacek (1985); Autor, Katz, and Kearney (2006); Black and Spitz-Oener (2010); Deming (2017); Caunedo, Keller, and Shin (2021).



Figure 1: Assignment Example

the skills involved in the tasks they perform.³ The plant can also count with robots (or software) capable of performing tasks. Tasks assigned to them are automated.

To fix ideas, Figure 1 describes the setup faced by the manager of a plant with three types of workers who perform a set of tasks to produce a good. The workers and tasks differ in cognitive and manual skills. Each point in the plane characterizes a task with a different combination of skills. While some tasks are complex in terms of their cognitive skills and involve no manual skills, others use both types of skills, and so on. The distribution of tasks in the skill space describes the combination of tasks required for production. The main assumption I impose on the model is the discreteness of the distribution of workers.⁴ Workers are represented by points scattered in the skill space, defining given combinations of skills (e.g., x_1, x_2, x_3).

Note: The figure shows an example for an assignment in a two-dimensional skill space (cognitive and manual skills). There are three types of workers $\{x_1, x_2, x_3\}$, and tasks are continuously distributed over the unit square. The assignment partitions the space into three regions $\{\mathcal{Y}_1, \mathcal{Y}_2, \mathcal{Y}_3\}$ each of which is an occupation. x_n performs the tasks in occupation \mathcal{Y}_n . The assignment in this figure is not necessarily optimal.

³Multidimensional skill mismatch plays an important role for workers' occupational choices and earnings dynamics. See Poletaev and Robinson (2008), Kambourov and Manovskii (2009), Gathmann and Schönberg (2010), Yamaguchi (2012), Guvenen, Kuruscu, Tanaka, and Wiczer (2020), and Lise and Postel-Vinay (2020).

⁴The same assumption is used by Stokey (2018) to study the effects of task-biased technical change and by Adenbaum (2021) to study the gains from worker specialization. Discretizing worker types is also useful to study labor demand shocks and sorting (Fogel and Modenesi, 2021; Grigsby, 2021; Haanwinckel, 2021).

In the figure, the problem of the manager is to assign tasks to workers for production. The discreteness of the workers' distribution implies that a bundle of tasks is assigned to each worker, as is observed in plant-level data (Combemale, Whitefoot, Ales, and Fuchs, 2021a,b). The tasks in each bundle form the worker's occupation (Rosen, 1978). The figure shows the boundaries of occupations implied by an arbitrary assignment of tasks to workers that divides the space into regions (occupations) \mathcal{Y}_1 , \mathcal{Y}_2 , and \mathcal{Y}_3 . Alternative assignments need not divide the space into connected regions or can assign no tasks to some workers.

The optimal assignment of tasks to workers seeks to maximize production by minimizing skill mismatch, subject to the limited supply of workers of each type. Technology determines the directions in which mismatch is more harmful. For example, cognitive mismatch may affect productivity more than manual mismatch, or being over-qualified for a task may be less harmful than being under-qualified. The direction of mismatch determines the shape of the occupation boundaries, while the supply of workers' skills determines their location.

Using the model, I show that the key outcomes are fully characterized by the boundaries that define occupations. First, I show that marginal products depend only on the productivity of workers in tasks along the boundaries of their occupations. Similarly, the skill premia and wages are determined the workers' skill mismatch in boundary tasks. Marginal products capture how much production increases if additional tasks were reassigned to a worker. These additional (marginal) tasks come from the displacement of neighboring workers along the boundaries of the worker's occupation. Because of this, only boundary tasks are relevant for the skill premia and wages.

I then show that the productivity differences in boundary tasks induce an endogenous ranking of workers. Because of the multidimensional nature of skills, there is not an a priori ranking of low- or high-skilled workers; rather, the ranking reflects their marginal products. The displacement of workers when reassigning tasks along the boundaries of occupations reveals this ranking. It is only optimal to displace workers if they are less productive along the boundary than their neighbors, that is, if they have a lower marginal product. This ranking changes in response to changes in technology, task demand, automation, etc., causing tasks to be reassigned across workers. Rank changes across occupations are an important feature of U.S. wage dynamics as shown in Tan (2021).

After establishing the assignment's properties, I show how to use the model to study worker-replacing technologies like automation and offshoring.⁵ These technologies are directed toward replacing workers in specific tasks, taking away some, but not all, of the tasks of an occupation (e.g., industrial robots taking spots in the assembly line),⁶ and therefore they are more likely to transform rather than eliminate occupations. In the model, occupations are transformed directly by losing tasks to robots or software and indirectly through tasks being reassigned across workers.

I model automation as the introduction of a mass of identical "robots." Automation is directed through the choice over the robots' location in the skill space and the tasks assigned to them. This works in the same way as choosing which tasks to offshore. Automation is optimally directed toward regions that exhibit high skill mismatch between workers and tasks, regions around the occupation boundaries and not necessarily tied to low- and high-ranked occupations. In this way, automation can affect low-ranked occupations—like cashiers with the introduction of self-checkout machines—or high-ranked occupations—like stockbrokers and accountants with trading and accounting software that takes over some of their tasks.

Automation generates task displacement, triggering a reassignment of tasks across occupations. I show that workers who previously performed the automated tasks are not the only ones affected; less productive workers are also displaced as the boundaries of occupations shift to preserve the employment of more productive workers. Whether or not wages decrease overall depends on how productive robots are at the tasks they take over

⁵For manufacturing, Acemoglu and Restrepo (2020) estimate that industrial robots have displaced 756,000 workers between 1993 and 2007. Simultaneously, advances in software and AI have made it possible to automate tasks of clerical occupations and of more specialized workers like accountants. Offshoring, another worker-replacing technology, operates through the same channels (Blinder, 2009; Blinder and Krueger, 2013).

 $^{^{6}}$ McKinsey Global Institute (2017) reports that while 50% of tasks are automatable using currently available technology, less than 5% of occupations are fully automatable.

(Acemoglu and Restrepo, 2018a). The model predicts that workers with higher task displacement have larger reductions in wages, as documented by Acemoglu and Restrepo (2021), and that the worker rank only changes in response to large task displacement.

In addition to automation, I consider two forms of directed worker-enhancing changes in the form of worker training and skill-biased technical change, and I establish their effect on workers. First, I show that worker training increases the worker's marginal product and wage by reducing the mismatch in the tasks they perform. Unlike automation, training does not displace workers, although occupations do change in response to the new skill distribution of workers. Second, I find that skill-biased technical change is optimally directed toward skills at which the workforce is most adept. Productivity increases by complementing the skills with the lowest mismatch. This contrasts with automation, where productivity increases by replacing workers in tasks they are not well suited for, those with the highest mismatch.

Finally, I extend the model by allowing tasks to be left unassigned. In this scenario, tasks are only performed if workers are productive enough relative to their wages. I show how skill accumulation by workers changes, and potentially expands, the set of tasks performed in the economy. One important consequence of allowing tasks to be left unassigned is that automation ceases to be a pure worker-replacing technology. Automation can now complement workers by taking over tasks that are either not worthwhile for workers to perform or are too specialized given their current skills.

Related literature I adopt a task approach to production as in the single-dimensional models of Rosen (1978), Autor, Levy, and Murnane (2003), and Acemoglu and Autor (2011). I complement this literature by incorporating multidimensional heterogeneity in tasks and workers as in Lindenlaub (2017), but I depart from her model in making the distribution of workers discrete. Because of this, the assignment generates many-to-one rather than one-to-one-matches, which allows me to study the bundling of tasks into occupations. Relative to single-dimensional models, the complexity level of a task is endogenously determined and

worker specific within my multidimensional framework. This complexity level serves a singledimensional index for tasks. Underlying skill differences across workers mean that the same task may have a high rank for some workers and a low rank for others. Moreover, this rank responds to changes in technology that affect workers' productivity across tasks.

My work also relates to applications of assignment models building on the classic work of Roy (1951) and Sattinger (1975, 1993), most recently to the work of Gola (2021), who studies the assignment of workers to firms across different sectors. I model occupations as the outcome of an assignment from a continuum of tasks to a discrete number of workers. This setup has been used by Feenstra and Levinsohn (1995) in the context of a continuum of buyers choosing from a discrete set of products. I generalize their results using tools from optimal transport theory (Villani, 2009; Galichon, 2016) and extend the model to consider applications to technical change, unemployment, automation, and worker training.⁷

Finally, I provide a theoretical framework for analyzing the direction and consequences of automation. This contributes to the literature on automation and other worker-replacing technologies (e.g., Acemoglu and Restrepo, 2018b, 2020, 2021; Aghion, Jones, and Jones, 2019; Hemous and Olsen, 2021). I explicitly model the multidimensional nature of skill heterogeneity, which is important for determining the automatability of tasks, as shown recently by Frey and Osborne (2017), Webb (2020), and Martinez and Moen-Vorum (2021).

1 Task assignment model

I present a model where occupations arise endogenously as bundles of tasks assigned to workers. In this model, the bundling of tasks into occupations reacts to changes in technology (e.g., automation, skill-biased technical change) and demographics (e.g., the distribution of workers' skills), implying changes in the productivity, substitutability, and compensation of workers. I characterize workers with the vector of skills they possess and tasks with the

⁷In the setup of Feenstra and Levinsohn (1995), the techniques I develop can be applied to the problem of designing a new product (defined by a vector of characteristics) given a distribution of consumers.

vector of the skills involved in performing them. These skills can be general traits, such as cognitive, manual, and social ability, or more specific, such as typing, programming, lifting, and analytical writing.

The model describes the problem of organizing production at a single production unit (i.e., a plant). Production requires workers to complete a continuum of tasks. The production unit has finitely many types of workers at its disposal to be assigned to tasks. A single type of worker can then perform various tasks. I refer to the set of tasks performed by a worker as their occupation, as in studies using plant-level production data (Combemale et al., 2021a,b). A worker's productivity on a given task depends on how well their skills match the skills used in performing the task (i.e., the degree of mismatch). In what follows, I describe the role of workers, tasks, and the production technology.

Workers Workers are characterized by a vector of skills $x \in S \subset \mathbb{R}^d$, where S is the space of skills and $d \ge 1$ is the number of skills. There are N types of workers in the economy: $\{x_1, \ldots, x_N\} \equiv \mathcal{X}$. There is a mass p_n of workers of type x_n . Each worker is endowed with one unit of time and supplies their time inelastically so that workers of type n supply a total of p_n units of time.⁸ The workers' outside option, if they are not assigned to any task, has a value of $\underline{w} \ge 0.^9$ When studying the effects of automation in Sections 4.1 and 5.2, I change the set of workers by introducing robots that can also perform tasks.

Tasks Production requires completing a set of tasks $\mathcal{Y} \subseteq \mathcal{T}$. To facilitate exposition, I assume that $\mathcal{T} \equiv S$ so that tasks and workers are elements of the same space S. This assumption does not affect any of the results (see the general setup in Appendix B). Tasks $y \in \mathcal{Y}$ differ in the skills involved in performing them and how many times they must be performed. One unit of time is required to perform a task once. Tasks are continuously

⁸An alternative interpretation is that of N different workers with effective productivity $\{p_1, \ldots, p_N\}$.

⁹In general, the outside options of workers depend on their type, so a type n worker has an outside option of \overline{w}_n . This does not change the optimal assignment described in Section 2 but plays a role when tasks can be left unassigned in Section 5.

distributed on \mathcal{Y} . I denote the density of tasks used in production by $g : \mathcal{Y} \to \mathbb{R}_+$ and maintain the following assumption throughout:

Assumption 1. The set and distribution of tasks satisfy the following properties:

- *i.* $g: \mathcal{Y} \to \mathbb{R}_+$ is an absolutely continuous probability density function with an associated absolutely continuous measure G on \mathcal{Y} ;
- ii. there are enough workers to complete all tasks; i.e., $G(\mathcal{Y}) = \int_{\mathcal{Y}} g(y) \, dy \leq \sum_{n=1}^{N} p_n$;
- iii. and the set of tasks \mathcal{Y} is compact.

Task output Workers vary in their productivity across tasks depending on the degree of mismatch between the skills they possess (x) and the skills involved in performing the task (y). The function $q : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}$ gives the worker-task-specific output generated by a worker with skills x performing task y. No output is generated by unassigned tasks. Abusing notation, I use \emptyset to denote unassignment so that $q(\emptyset, y) = 0$ for all $y \in \mathcal{Y}$. In general, the degree of mismatch between the worker's and task's skills is captured by some notion of the distance between the worker and the task in the skill space.

The optimal assignment balances the desire to minimize the mismatch between workers and the tasks they perform, with the capacity constraints imposed by the limited availability of workers. The assignment's exact shape depends on how q captures mismatch but the general properties of the assignment do not. I return to this in Section 2, where I discuss the optimal assignment, and in Section 4.2, where I discuss changes in technology.

Finally, it will prove useful to impose a functional form on q to derive further results:¹⁰

Assumption 2. The task output production function satisfies

$$q(x,y) = \exp\left(a'_{x}x + a'_{y}y - (x-y)'A(x-y)\right),$$
(1)

¹⁰The functional form in Assumption 2 corresponds to the payoff function of Tinbergen (1956), the utility function of Feenstra and Levinsohn (1995), and the production technology of Lindenlaub (2017).

where A is symmetric and positive definite.

Under (1), the mismatch between a worker and a task is measured by the weighted quadratic distance between the worker's and task's skills. Matrix A controls the weights of each skill in the mismatch. The terms $a'_x x$ and $a'_y y$ capture more skilled workers and tasks involving higher skill levels generating more output.

The assignment of tasks to workers is described by a function $T: \mathcal{Y} \to \mathcal{X}$ Assignment so that task y is performed by worker $T(y) \in \mathcal{X}$. The set of tasks performed by a type of worker forms that worker's occupation. The occupation of type x_n workers is¹¹

$$\mathcal{Y}_n \equiv T^{-1}(x_n) = \left\{ y \in \mathcal{Y} \, | \, x_n = T(y) \right\}. \tag{2}$$

Figure 1 in the introduction shows an example of an assignment that partitions the space of tasks into three occupations, corresponding to three worker types.

An assignment is *feasible* if workers can supply the time demanded by their occupation. The demand for workers x_n corresponds to the time required to perform their assigned tasks:

$$D_n \equiv \int_{\mathcal{Y}_n} dG. \tag{3}$$

Definition 1. An assignment T is feasible if $D_n \leq p_n$ for all $n \in \{1, \ldots, N\}$.

Production Technology Production at the plant aggregates the output from all worker/task pairs through a Cobb-Douglas technology given an assignment $T{:}^{12}$

$$F(T) = \exp\left(\int_{\mathcal{Y}} \ln q\left(T\left(y\right), y\right) \, dG\right). \tag{4}$$

¹¹It is possible that $\mathcal{Y}_n = \emptyset$ so that no task is assigned to worker x_n . ¹²The aggregator does not need to be of the Cobb-Douglas type. The results hold for aggregators of the CES family: $F(T) = \left(\int (q(T(y), y))^{\frac{\sigma-1}{\sigma}} dG(y)\right)^{\frac{\sigma}{\sigma-1}}$, with $\sigma > 1$. See Appendix B.

Under this technology, production only occurs if all tasks are assigned and performed. Recall that task output is zero if a task is left unassigned; $q(\emptyset, y) = 0.^{13}$ I extend the technology to allow for unassigned tasks in Section 5.

Single-dimensional representation It is possible to map the underlying multidimensional heterogeneity into worker-specific single-dimensional indices of task complexity. Doing so facilitates comparison with single-dimensional task assignment models and makes explicit the role of multidimensional skills in the model.¹⁴ These indices are endogenously determined by the model's primitives (here the production function q and the distribution of tasks G). Moreover, rather than having a common ordering for tasks (where workers rank tasks in the same way), the ordering is worker specific, reflecting the underlying differences in skills and the mismatch between workers and tasks.

The single-dimensional complexity index corresponds to worker-specific task rankings, with the higher-ranked tasks being more productive. The rank is worker specific because different workers have different productivity in a given task depending on their underlying (multidimensional) skill mismatch. Crucially, changes in the distribution of tasks or the production technology affect the rankings and can do so in different ways for different workers. The rank of a task y for worker x_n is

$$r_n(y) \equiv \Pr\left(q\left(x_n, Y\right) \le q\left(x_n, y\right)\right),\tag{5}$$

where the probability is taken with respect to the distribution of tasks, G. The ranks are

¹³The production technology resembles a continuous version of Kremer (1993)'s O-ring production function. To make the comparison precise, it is necessary to change the interpretation of q. Consider a continuous production line indexed by $y \in \mathcal{Y}$. A fatal error can occur at each point in the production line that terminates the production process in failure. The arrival rate of an error is given by $\ln q(x, y) \ge 0$ and depends on the point in the production process (y) and the worker assigned to that point (x). The probability that no error arrives at the end of the whole process is given by (4). Thus, F(T) can also be interpreted as expected output given an assignment T. See Sobel (1992) for another application of this idea.

¹⁴The approach is akin to the construction of the "canonical formulation" of the multisector model in Gola (2021). Adopting the single-dimensional representation is also useful for studying how changes in the overall demand for skills affect the assignment. See Appendix C.

worker independent when q satisfies the separability assumption of Chiappori, Oreffice, and Quintana-Domeque (2012, p. 665).

The model's representation is completed by defining an alternative production function, $Q_n: [0,1] \to \mathbb{R}_+$, that maps rankings into task output:

$$Q_n\left(r_n\left(y\right)\right) \equiv q\left(x_n, y\right). \tag{6}$$

 $Q_n(r)$ gives the output of worker x_n when performing a task with rank r. It follows from the definition of the rank in (6) that Q_n is increasing in r. For a given set of primitives $\{\mathcal{X}, \mathcal{Y}, q, G\}$, it is then possible to solve for the assignment in terms of $\{r_n, Q_n, \}_{n=1}^N$. However, the rankings' dependence on the model's primitives means that the mapping must be recomputed in response to any change in the economic environment. Because of this, I focus on the solution to the multidimensional setup in the rest of the paper.

2 The optimal assignment of tasks to workers

The problem is finding a feasible assignment that maximizes output:

$$\max_{T} F(T) \qquad \text{s.t. } \forall_{n} D_{n} \le p_{n}.$$
(7)

This can be seen as the problem of a manager organizing production, taking as given the type and quantity of workers at their disposal. If all firms are identical, the allocation coincides with that of firms hiring workers in a decentralized market and assigning them to tasks once hired, with wages determined by the workers' marginal products (see Section 3.2).¹⁵

It is possible to guarantee the existence and uniqueness of a solution by imposing conditions only on the production technology q. Proposition 1 makes this precise:

¹⁵The model can also be interpreted as a multisector assignment model where heterogeneous workers sort into a finite set of sectors. See Appendix C.

Proposition 1. Consider the optimal assignment problem in (7). If q is such that

- *i.* every worker/task pair is productive: q(x, y) > 0 for all pairs $(x, y) \in \mathcal{X} \times \mathcal{Y}$;
- ii. $q(x, \cdot)$ is upper-semicontinuous in y given $x \in \mathcal{X}$; and
- *iii.* q discriminates across workers: for all $x_n \neq x_\ell$, $q(x_n, y) \neq q(x_\ell, y)$ G-a.e.

Then, there exists a (G-) unique solution T^* to the problem in (7). Moreover, there exists a unique $\lambda^* \in \mathbb{R}^N$ with $\min \lambda_n^* = 0$ such that T^* is characterized as

$$T^{\star}(y) = \underset{x \in \mathcal{X}}{\operatorname{argmax}} \left\{ \ln q\left(x, y\right) - \lambda_{n(x)}^{\star} \right\},$$
(8)

where n(x) gives the index of a type of worker $x \in \mathcal{X}$.

Proof. The result is established by expressing the problem in (7) as an optimal transport problem. The proof is divided into three lemmas that relax the problem by allowing for non-deterministic assignments and then constructing a solution from the dual of the relaxed problem. A non-deterministic assignment is a joint measure over worker/task pairs: $\pi : \mathcal{X} \times \mathcal{B}(\mathcal{Y}) \to \mathbb{R}_+$, where $\mathcal{B}(\mathcal{Y})$ denotes the Borel sets of \mathcal{Y} . In contrast to a deterministic assignment T, a non-deterministic assignment π allows for mixing in the matching of workers and tasks, so that the same task can be performed by more than one type of worker. The relaxed problem, in terms of π , is linear in the choice of the assignment after a suitable transformation of the objective function:

$$\max_{\pi \in \Pi(P,G)} \sum_{n=1}^{N} \int_{\mathcal{Y}} \ln q \left(x_n, y \right) \, d\pi \left(x_n, y \right). \tag{9}$$

The dual to (9) expresses the problem in terms of the multipliers (or potentials) λ and ν :

$$\max_{\pi \in \Pi} \sum_{n=1}^{N} \int_{\mathcal{Y}} \ln q \left(x_n, y \right) \, d\pi \left(x_n, y \right) = \inf_{\substack{(\lambda, \nu) \in \mathbb{R}^N \times L^1(G)\\\lambda_n + \nu(y) \ge \ln q(x_n, y)}} \sum_{n=1}^{N} \lambda_n p_n + \int_{\mathcal{Y}} \nu \left(y \right) \, dG \qquad (10)$$
$$= \inf_{\nu^x \in \mathbb{R}^N} \sum_{n=1}^{N} \nu_n^x p_n + \int_{\mathcal{Y}} \max_n \left\{ \ln q \left(x_n, y \right) - \lambda_n \right\} \, dG.$$

Under the conditions above, this problem admits a unique solution that characterizes a deterministic assignment T^* according to equation (8).

Lemmas B.1, B.2, and B.3 formalize these arguments and establish the properties of the solution to the primal and dual problems. The results follow from Theorems 5.10 and 5.30 in Villani (2009), summarized in Theorem A.1 of Appendix A. All lemmas are stated and proven in Appendix B.

Before discussing the characterization of the optimal assignment, I briefly discuss the role of the three conditions in Proposition 1. The first condition ensures that F(T) is

defined for any assignment, the second guarantees that duality applies to the problem so that the equality in (10) holds, and the third plays a crucial in role in establishing the uniqueness of an optimal assignment function T^* . The injectivity of q in x given y makes it possible to distinguish between workers in each task. Without this injectivity, there would not be a unique deterministic assignment; instead, optimal assignments would be non-deterministic but only differ in the assignment of workers across tasks in which they are equally productive.¹⁶ The discreteness of the problem on the worker side considerably relaxes the standard differentiability conditions necessary to obtain a deterministic assignment. The injectivity plays the same role as the "twist condition" of Carlier (2003) and the condition for positive assortative matching in Lindenlaub (2017).¹⁷

The characterization of the optimal assignment in (8) allows me to more explicitly characterize the occupations in terms of the production technology q:

$$\mathcal{Y}_n = \left\{ y \in \mathcal{Y} \,|\, \forall_\ell \, \ln q \, (x_n, y) - \lambda_n^\star \ge \ln q \, (x_\ell, y) - \lambda_\ell^\star \right\}. \tag{11}$$

Tasks are optimally assigned to workers who are more productive at performing them (lower skill mismatch), subject to the penalty captured by λ^* that balances the demand for that type of worker with the limited supply of workers (p_n) .

The occupation boundaries play a central role in determining the assignment's properties, like the the workers' marginal product, compensation, and substitutability. Intuitively, the boundaries are formed by a worker's marginal tasks, a concept that I expand on in Section 3.1. Formally, the boundaries are formed by task $y \in \partial \mathcal{Y}_n$ for which the inequality in (11) is

¹⁶This assumption precludes q to take a Leontieff form, i.e., requiring a minimum level of skill to perform tasks, where skills in excess of the requirement do not affect output. In this case there would be areas of the task space that can be arbitrarily divided among (qualified) workers, generating the same overall output.

¹⁷When the distribution of workers is continuous, it is necessary to distinguish between workers who are arbitrarily close (in the skill space). This idea is captured by the derivative of the objective function, which is required to be injective, often referred to as the twist condition. These conditions underlie the local existence results for the assignment in Villani (2009, ch. 10) used across many applications of optimal transport. In technical terms, both the condition (*iii*) in Proposition 1 and the conditions on the injectivity of the derivative of the objective function guarantee that the q-subdifferential of (the dual potential) ν is single valued: $\partial^q \nu(y) = \{x \in \mathcal{X} \mid \lambda(x) + \nu(y) = \ln q(x, y)\}$. The q-subdifferential of ν , $\partial^q \nu(y)$, corresponds to the solution of the optimal assignment problem as in (8). See Lemma B.3 in Appendix B.



Figure 2: Assignment Example: Mismatch Loss

Note: The figure shows the optimal assignments in a two-dimensional skill space (cognitive and manual skills) with three types of workers $\{x_1, x_2, x_3\}$ with mass $P = \{0.4, 0.3, 0.3\}$. Tasks are uniformly distributed over the unit square; i.e., $\mathcal{Y} = [0,1]^2$ and g(y) = 1. The production function q is given by $\ln q(x, y) = a'_x x + a'_y y - d(x, y)$, where d(x, y) is a distance measure capturing mismatch. Each sub-figure considers a different distance of the L-p distance family: $d(x, y) = \left(\sum_{i=1}^d |x_i - y_i|^p\right)^{1/p}$. A higher p puts more weight on the dimension with the highest mismatch. The Chebyshev distance is obtained as $\lim_{p\to\infty} \left(\sum_{i=1}^d |x_i - y_i|^p\right)^{1/p} = \max_i |x_i - y_i|$.

met with equality for some ℓ :

$$\partial \mathcal{Y}_{n} = \left\{ y \in \mathcal{Y} | \exists_{\ell \neq n} \ln q \left(x_{n}, y \right) - \lambda_{n}^{\star} = \ln q \left(x_{\ell}, y \right) - \lambda_{\ell}^{\star} \right.$$

$$\wedge \forall_{m \neq \ell, n} \ln q \left(x_{n}, y \right) - \lambda_{n}^{\star} \ge \ln q \left(x_{m}, y \right) - \lambda_{m}^{\star} \right\}.$$

$$(12)$$

The assignment's key feature is that the mismatch between workers (pairs) is constant across the occupation boundaries, as in equation (12). The difference in mismatch is reflected by the difference in the productivity between workers $(\ln q (x_n, y) - \ln q (x_{\ell}, y))$ and is captured by the multipliers as $(\lambda_n^{\star} - \lambda_{\ell}^{\star})$.

The shape of the boundaries in the skills space depends on how the mismatch measure implicit in q weighs the distance between workers and tasks. However, the assignment follows the common principle of minimizing mismatch regardless of the functional form chosen for q. Figure 2 exemplifies this by plotting the optimal assignment under four different measures of distance for the plane. I consider distances belonging to the L-p family: $d(x,y) = \left(\sum_{i=1}^{d} |x_i - y_i|^p\right)^{1/p}$. A larger value for p makes the distance measure more sensitive to the dimension of highest mismatch, which affects the curvature of the boundaries.¹⁸

Imposing Assumption 2 simplifies the characterization of occupations. When q is as in equation (1), the occupation boundaries take the form of hyperplanes whose normal vectors depend on matrix A and the difference in skills between neighboring workers. The boundary between the occupations of workers x_n and x_ℓ is

$$y \in \mathcal{Y}_n \cap \mathcal{Y}_\ell \longleftrightarrow 0 = y' \underbrace{A(x_\ell - x_n)}_{\text{Normal Vector}} - \frac{1}{2} \underbrace{\left(x'_\ell A x_\ell - x'_n A x_n + a'_x (x_\ell - x_n) + \lambda'_\ell - \lambda'_n\right)}_{\text{Intercept}}.$$
 (13)

Equation (13) reveals an equivalence between the optimal assignment and the partition

¹⁸Alternative measures of mismatch can reduce the dimensionality of the problem by encoding the differences between workers and tasks into a single dimension. This is the case when measuring mismatch with the cosine similarity between the skill vectors of workers and tasks. The assignment problem reduces to finding cutoff values for angles $\theta \in [0, \pi/2]$ that characterize rays partitioning the space. The cosine similarity has been used to measure the distance between occupations by Gathmann and Schönberg (2010) and Baley, Figueiredo, and Ulbricht (2021). I provide details for the cosine similarity mismatch in Appendix D.



Figure 3: Assignment Example: Quadratic Mismatch Loss

Note: The figure shows the optimal assignment in a two-dimensional skill space (cognitive and manual skills) with three types of workers $\{x_1, x_2, x_3\}$ with mass $P = \{0.4, 0.3, 0.3\}$. Tasks are uniformly distributed over the unit square; i.e., $\mathcal{Y} = [0, 1]^2$ and g(y) = 1. The production function q is as in (1) with $A = I_2$.

induced by a power diagram, which allows me to use tools from computational geometry to analyze the assignment's properties.¹⁹ The optimal assignment partitions the space into convex polyhedra defined by hyperplanes as in Aurenhammer, Hoffmann, and Aronov (1998) and Galichon (2016, ch. 5). Figure 3 exemplifies this by plotting the optimal assignment under the task output function of Assumption 2.

Indirect production function The optimal assignment implies an indirect production function that transforms inputs (workers) into goods. Crucially, the amount of an input (a type of worker) used in production and what that input is used for are not the same (Autor, 2013). Consequently, the relationship between inputs and output depends on how tasks are assigned to workers and how the assignment itself changes as the amount of inputs varies.

¹⁹A power diagram partitions a space into cells that minimize the power between a node (x) associated with the cell and the points y in the cell. The power function between two points is pow $(x, y) = d(x, y)^2 - \mu$, where d(x, y) is a distance and $\mu \in \mathbb{R}$.

The role of workers in production is captured by the value of the assignment problem (7):

$$V(p_1, \dots, p_N) \equiv \max_T F(T) \qquad \text{s.t. } \forall_n D_n \le p_n.$$
(14)

Function $V : \mathbb{R}^N_+ \to \mathbb{R}_+$ describes how production changes when the workforce composition changes, allowing for workers to be reassigned optimally across tasks.

3 Properties of the assignment

I now examine the properties of the assignment. The multipliers λ^* that characterize the assignment in (8) play a central role by determining the boundaries of occupations and therefore workers' marginal product and compensation. I show that λ^* also implies an ordering of workers based on their relative productivity in boundary tasks. This ranking plays an important role in shaping the effects of automation and other worker-replacing technologies studied in Section 4.

3.1 Marginal products and worker ranks

The marginal product of type n workers is obtained from the indirect product function V as the change in output if the supply of type n workers (p_n) were to increase. In this way, the marginal product takes into account how the assignment changes optimally in response to an increase in the supply of type n workers.²⁰

Proposition 2. The marginal product of type n workers is

$$MP_n \equiv \frac{\partial V\left(p_1, \dots, p_N\right)}{\partial p_n} = F\left(T^\star\right) \lambda_n^\star.$$
(15)

²⁰It is also possible to define a measure of the marginal product of a type n worker at a given task y, given some arbitrary assignment T. See Appendix B.2.

Proof. The result follows from the relationship between the multiplier λ of the relaxed problem (9) and its dual (10) and the multipliers of the original problem (7). The relationship is obtained by applying the envelope theorem (Milgrom and Segal, 2002) and the definition of the indirect production function in (14). See Lemma B.4 in Appendix B for a detailed derivation of the result.

The relationship between the value of λ^* and the workers' marginal product comes from how the assignment responds to an increase worker supply. When p_n increases, the additional workers increase output only if tasks are reassigned to them from other workers. The first tasks to be reassigned are those in the occupation boundaries. Consider the occupations of two types of workers, n and ℓ . Tasks in the boundary, $y \in \mathcal{Y}_n \cap \mathcal{Y}_\ell$, satisfy

$$\lambda_n^{\star} - \lambda_\ell^{\star} = \ln q \left(x_n, y \right) - \ln q \left(x_\ell, y \right). \tag{16}$$

Then, the difference in the multipliers λ_n^* and λ_ℓ^* is given by the log difference of output in boundary tasks, that is, the percentage increase (or decrease) in output if the tasks along the boundary are reassigned from ℓ to n.²¹ Thus, it is only optimal to use the additional supply of workers if output increases along the boundary of their occupation. In this sense, these are a worker's marginal tasks, as in Figure 4a.

The reassignment process that originates in the increase in the supply of workers creates a cascading effect that affects other workers' assignments (Figure 4b). As tasks are reassigned toward type n workers, workers along the \mathcal{Y}_3 boundaries are displaced. This process generates an excess supply of workers of other types, giving rise to a new round of reassignment along the boundaries. Following the cascading process reveals an ordering of workers by productivity, with the least productive worker being displaced by increases in the supply of more productive ones. As a result, the least productive worker has zero marginal product. Increases in the supply of that type of workers do not increase output because the additional

²¹It is useful to consider an example with finitely many tasks, say $\{y_1, y_2\}$, one assigned to worker x_n and the other to worker x_ℓ . Then, total output is $F(T) = q_1(x_n) q_2(x_\ell)$. If the assignment changes by having worker x_n perform both tasks, the new output is $F\left(T'\right) = \frac{q_2(x_n)}{q_2(x_\ell)}F(T)$. Then, $\ln \frac{F\left(T'\right)}{F(T)} = \ln \frac{q_2(x_n)}{q_2(x_\ell)} = \lambda_n - \lambda_\ell$ so that output increases by approximately $100 (\lambda_n - \lambda_\ell) \%$.





Note: The figure shows the assignment in a two-dimensional skill space (cognitive and manual skills). Five types of workers are considered $\{x_1, \ldots, x_5\}$ with mass $P = \{0.3, 0.2, 0.3, 0.1, 0.1\}$. Tasks are uniformly distributed over the unit square; i.e., $\mathcal{Y} = [0, 1]^2$ and g(y) = 1. The production function q is as in (1) with $A = I_2$. The left panel highlights the marginal tasks of worker x_3 . The right panel shows the reassignment of tasks following an increase in the supply of x_3 workers. The \mathcal{Y}_3 boundaries are pushed outwards as well as the shared boundary of \mathcal{Y}_2 and \mathcal{Y}_1 , following the cascading effect of the reassignment.

workers are left unassigned.²²

The total gain in output from the initial increase in the supply of type n workers takes into account the increase in output from all the reassignments. Using the relation in (16), and recalling that min $\lambda_n^* = 0$, we get a total increase in output of λ_n^* as in Lemma 2.

The value of λ^* depends on the differences in skills and skill mismatch of workers relative to the least productive worker. Under Assumption 2, it is possible to make this precise. Manipulating equation (13) gives

$$\lambda_{n}^{\star} = \underbrace{a_{x'}(x_{n} - \underline{x})}_{\text{Difference in Skills}} - \underbrace{(x_{n} - y_{n})'A(x_{n} - y_{n})}_{x_{n} \text{ mismatch at boundary}} - \underbrace{(\underline{x} - \underline{y})'A(\underline{x} - \underline{y})}_{\underline{x} \text{ mismatch at boundary}},$$
(17)
Difference in Mismatch

²²The least productive worker has zero marginal product because of the capacity constraint on the set of tasks captured by G. This motivates the normalization of the multiplier λ in Proposition 1.

where \underline{x} are the skills of the least productive worker (the worker with $\lambda_m^* = 0$) and y_n and \underline{y} are boundary tasks of workers x_n and \underline{x} , respectively. This makes it clear that the contribution of each skill to the worker' marginal product depends on the whole assignment. Therefore, skills are not individually priced independently of the occupations, similar to the results of Mandelbrot (1962) and Heckman and Scheinkman (1987).

Finally, it is worth stressing the role of the multipliers λ^* in determining worker ranks. An a priori ranking of workers does not exist because they are heterogeneous along multiple dimensions. However, the optimal assignment implies a ranking based on the relative productivity of workers in their boundary tasks. This ranking is completely captured by λ^* and has direct implications for how a reassignment happens following a change in the environment. I will return to this in Section 4.1 when I introduce worker-replacing technologies like automation.

3.2 Worker compensation and the skill premia

The marginal product determines worker's compensation. Consider the problem of a pricetaking firm seeking to maximize profits. The firm's problem is to choose both the demand for each type of worker and the assignment of tasks to workers:

$$\max_{T} F(T) - \sum_{n=1}^{N} w_n D_n(T), \qquad (18)$$

where w_n is the wage paid to type *n* workers and D_n is their demand (equation 3). This problem results in the same assignment as the optimal assignment problem in (7) if the workers' wages correspond to their marginal products plus a constant guaranteeing that they all receive at least their outside option:²³

$$w_n = F(T^*) \lambda_n^* + \kappa, \quad \text{where } \kappa \ge \underline{w}.$$
 (19)

The level of wages is not pinned down because only the difference in wages affects the assignment (see equation 11).²⁴ Nevertheless, wage differences are determined by marginal products and are informative about the role of skill mismatch in determining skill premia. In particular, the wage of a worker relative to the lowest paid worker gives their marginal product, which depends on the worker's mismatch and skill level as in (16) and (17).

The marginal product of workers also pins down the labor share. Workers do not appropriate all the output they produce, only the output they generate in marginal (boundary) tasks plus the compensation for their outside option. The labor share is²⁵

$$LS \equiv \frac{\sum_{n=1}^{N} w_n D_n}{F(T)} = \sum_{n=1}^{N} \lambda_n^* p_n + \frac{\kappa G(\mathcal{Y})}{F(T)}.$$
(20)

3.3 Substitutability across workers

The substitutability of workers in production, as measured by their elasticity of substitution, plays an important role in policy analysis and in determining the effects of changes to technology. Intuitively, workers performing similar tasks are more substitutable, as are workers with similar skills. I make these results precise by computing the elasticity

²³These results apply if all production units in the economy operate with the same technology (i.e., G and q) or if markets are segmented. It is also possible to determine wages in an economy with various technologies described by different combinations of tasks (or blueprints), captured by different G distributions. Doing so becomes intractable in the multidimensional setting I am describing. See Haanwinckel (2021) for an application of these ideas to the dispersion of income in a single-dimensional framework.

²⁴The indeterminacy of the level of wages is a common feature of assignment models (Sattinger, 1993). This result is not a consequence of the discreteness in the distribution of workers; see Teulings (1995) and Lindenlaub (2017). The level of wages is pinned down, for example, when there is an excess supply of workers $(P > G(\mathcal{Y}))$ so that (at least) one type of worker will be partially unassigned (unemployed), driving down the wage for that type of worker to w. This will be the case when I introduce automation in Section 4.1.

 $^{^{25}}$ Adenbaum (2021) exploits the first order condition of this problem and the relationship of marginal products with the labor share to determine productivity gains from worker specialization.

of substitution under the optimal assignment. To do this, I use the Morishima elasticity of substitution (Blackorby and Russell, 1981, 1989), which is the appropriate measure of substitutability in a setup with more than two types of workers.²⁶ It measures the change on the ratio of demands for two inputs (in this case two workers, D_{ℓ}/D_n) after the marginal products (MP_n, MP_{\ell}) change.

Definition 2. The elasticity of substitution between type n and ℓ workers is

$$M_{\ell n} \equiv \frac{\partial \ln D_n / D_\ell}{\partial \ln M P_n} = \mathcal{E}_{\ell n} - \mathcal{E}_{nn}, \qquad (21)$$

where $\mathcal{E}_{\ell n} \equiv \frac{\mathrm{MP}_n}{D_\ell} \frac{\partial D_\ell}{\partial \mathrm{MP}_n}$ is the cross-elasticity of demand for worker k with respect to a change in worker n's marginal product.²⁷

A direct implication of the assignment is that workers are gross substitutes in the sense of Kelso and Crawford (1982): the elasticity of substitution is always positive. An increase in worker x_n 's marginal product, captured by an increase in λ_n^* , results in a decrease of their demand, D_n , and a weak increase in the demand for other workers (see equation 11). It follows that $\mathcal{E}_{nn} < 0$ and $\mathcal{E}_{n\ell} \ge 0$. From the point of view of worker compensation, increasing λ_n^* raises the cost of worker n, causing the firm to substitute them for other workers.

Moreover, only neighbors are directly substitutable for one another. The cross-elasticity is zero between workers who do not share a boundary. Only worker x_n 's neighbors are directly affected when λ_n^* changes. The sensitivity of the boundaries depends on the slope of the production function q evaluated at the boundary tasks; see (16). To obtain the magnitude of the cross-elasticities, $\mathcal{E}_{n\ell}$, it is necessary to determine this change, which depends on q.

The geometric structure gained by imposing Assumption 2 makes it possible to further characterize the change in demand following a change in λ^* . In this case, the change in demand is given by the area of a (hyper)trapezoid, formed as the plane that defines the

²⁶See Baqaee and Farhi (2019) for a recent application of the Morishima elasticity.

²⁷With more than two inputs, the direction of the change in the ratio of marginal products matters because the demands for inputs changes differently depending on whether MP_n or MP_l vary (Blackorby and Russell, 1989, p. 885). Because of this, the elasticity is in general asymmetric so that $M_{\ell n} \neq M_{n\ell}$.



Figure 5: Assignment Example: Change in worker demand

Note: The figure shows the assignment in a two-dimensional skill space (cognitive and manual skills). Five types of workers are considered $\{x_1, \ldots, x_5\}$ with mass $P = \{0.3, 0.2, 0.3, 0.1, 0.1\}$. Tasks are uniformly distributed over the unit square; i.e., $\mathcal{Y} = [0,1]^2$ and g(y) = 1. The production function q is as in (1) with $A = I_2$.

boundary between occupations moves. Figure 5b illustrates this by increasing the value of λ_3^{\star} . The figure shows that the boundaries of \mathcal{Y}_3 shift "inward," reducing the demand for x_3 and increasing the demand for their neighbors. In contrast, the boundary of \mathcal{Y}_2 , which does not share a boundary with \mathcal{Y}_3 , does not change.

The following Proposition formalizes the above results:

Proposition 3. Let $\lambda^* \in \mathbb{R}^N$ characterize the optimal assignment as in (8). If q is continuous, then D_n is continuously differentiable with respect to λ and

i. if $\mathcal{Y}_n \cap \mathcal{Y}_\ell = \emptyset$, then $\frac{\partial D_n}{\partial \lambda_\ell} = \frac{\partial D_\ell}{\partial \lambda_n} = 0$ and $\mathcal{E}_{n\ell} = \mathcal{E}_{\ell n} = 0$; *ii.* $\frac{\partial D_n}{\partial \lambda_n} = -\sum_{\ell \neq n} \frac{\partial D_\ell}{\partial \lambda_n} < 0$.

If Assumption 2 holds, then

iii.
$$\forall_{\ell \neq n} \quad \frac{\partial D_n}{\partial \lambda_\ell} = \frac{\operatorname{area}(\mathcal{Y}_n \cap \mathcal{Y}_\ell)}{2\sqrt{(x_n - x_\ell)'A'A(x_n - x_\ell)}} = \frac{\int_{\mathcal{Y}_n \cap \mathcal{Y}_\ell} dG}{2\sqrt{(x_n - x_\ell)'A'A(x_n - x_\ell)}} \ge 0.$$

Proof. The first result follows immediately from the characterization of occupations in (11). The second result follows from the feasibility of the assignment, which ensures that $\sum_{n=1}^{N} D_n = \int_{\mathcal{Y}} dG$ so that the sum of demands is constant. Then,

$$\frac{\partial D_n}{\partial w_n} + \sum_{m \neq n} \frac{\partial D_m}{\partial w_n} = 0.$$
(22)

The third result extends the findings of Feenstra and Levinsohn (1995) for arbitrary configurations of workers (x) by applying Reynolds' transport theorem (see Theorem A.2 in Appendix A). I present the complete proof in Lemma B.5 of Appendix B.3.

The first result in Proposition 3 formalizes the idea that only neighbors are directly substitutable, while the second result exploits the feasibility of the assignment to find a relationship between the change in demand across workers. This relationship gives rise to a sharper formula for the elasticity of substitution that takes into account the cascading effect tied to task reassignment:

Corollary 1. The Morishima elasticity of substitution between workers x_n and x_ℓ is a weighted average of the cross-elasticities of demand of all workers, with the weights given by the demand of each worker type relative to worker n's demand.

$$M_{\ell n} = \left(1 + \frac{D_{\ell}}{D_n}\right) \mathcal{E}_{\ell n} + \sum_{m \neq n, \ell} \frac{D_m}{D_n} \mathcal{E}_{m n}.$$
(23)

The last result of Proposition 3 exploits the geometric structure induced by Assumption 2 to compute closed-form expressions for the derivatives of demand. The cross-derivative of demand depends on how exposed two workers are to one another, measured by the length of their boundary, and how similar their skills are, measured by the weighted distance between their skills (x_n and x_ℓ). This captures the idea that workers with more common tasks, i.e., with longer boundaries, are more substitutable. How much the boundary reacts to a change in demand depends on how similar workers are at performing tasks, as captured by the denominator. The closer their skills, the more substitutable they are along their boundary.

4 Directed technical change

Changes in technology are a major factor in shaping the way in which tasks are assigned to workers. For instance, the increase of information technology (IT) in the workplace has shifted focus away from manual skills and changed the distribution of tasks across occupations. Alternatively, automation technologies and offshoring have replaced workers in performing certain tasks across manufacturing jobs, customer services, accounting, etc.

In this section, I consider two forms of technical change and study how they affect the division of tasks into occupations. Innovation in worker-replacing technologies lead to certain tasks being automated and to (the remaining) tasks being reassigned to workers. Innovation in skill-enhancing technology, such as IT in the modern workplace, or the power loom in the industrial revolution, changes worker productivity, inducing task reassignments to reduce mismatch across occupations.

4.1 Worker-replacing technologies

I introduce worker-replacing technologies in the form of robots that can replace workers at performing tasks. The robot is a flexible technology that can be adapted to perform different types of tasks, capturing the flexibility of current technologies like industrial robots or AI programs (Frey and Osborne, 2017; Acemoglu and Restrepo, 2020). It also relates to other work-replacing technologies like offshoring (Blinder, 2009; Blinder and Krueger, 2013).

The automation problem consists of designing a robot and assigning tasks to the workers and the robot to maximize production. The tasks assigned to the robot are automated. I denote by $r \in \mathcal{S}$ the robot's skills and by $p_r \geq 0$ its supply. The automation technology is embodied by a cost function $\Omega : \mathcal{S} \times \mathbb{R}_+ \to \mathbb{R}$ so that the cost of producing a mass p_r of a robot with skills r is given by $\Omega(r, p_r)$. Once the robot is designed, the set of available workers is expanded to include it: $\mathcal{X}_R \equiv \{x_1, \ldots, x_N, r\}$. Accordingly, the assignment is now described by a function $T_R : \mathcal{Y} \to \mathcal{X}_R$. I denote by $q_R : \mathcal{S} \times \mathcal{Y} \to \mathbb{R}$ the production technology of the robot so that a robot with skills r performing task y produces $q_R(r, y)$.

When tasks are automated, the total demand for labor decreases, leaving some workers unassigned.²⁸ As tasks are assigned to the robot, the workers who would have performed those tasks are directly displaced. However, these workers do not necessarily remain unassigned because they can take over other workers' tasks. The end result of this process depends on their relative productivity as captured by λ^* . As in Section 3.1, workers with the lowest marginal product will become displaced even if the tasks in their occupation are not directly affected by automation.

However, a large enough displacement of tasks can change worker ranks, reflecting the increased mismatch of displaced workers after being reassigned. In this way, low levels of task displacement reduce wages, but larger levels can also change the worker ranking. These patterns are in line with findings for the U.S. from Acemoglu and Restrepo (2021).

The automation problem is to choose jointly the robot's skills and mass (r, p_r) and the new assignment (T_R) to maximize output, net of the automation cost (Ω) :

$$\max_{\{r,p_r,T_R\}} \quad F_R(T_R) - \Omega(r,p_r) \qquad \text{s.t. } \forall_n D_n \le p_n \qquad D_R \le p_r, \tag{24}$$

where

$$F_{R}(T_{R}) = \exp\left(\int_{\mathcal{Y}\setminus\mathcal{Y}_{R}} \ln q \left(T_{R}(y), y\right) \, dG + \int_{\mathcal{Y}_{R}} \ln q_{R}(r, y) \, dG\right)$$
(25)

and

$$\mathcal{Y}_R = T_R^{-1}(r), \qquad D_R = \int_{\mathcal{Y}_R} dG. \tag{26}$$

It is convenient to think of the problem in two steps, first solving for an optimal assignment given a set of workers and a robot and then choosing the robot's optimal skills and mass. The problem of finding an optimal assignment can be simplified by using the results in Proposition 1. Taking as given the robot's skills and mass (r, p_r) , the optimal

²⁸This is a consequence of the assumption that the set of tasks to be performed (\mathcal{Y}) is fixed, as is the distribution of tasks (G). I relax this assumption in Section 5.

assignment is characterized by a vector $\mu^{\star} \in \mathbb{R}^{N+1}$:²⁹

$$T_{R}^{\star}(y) = x_{n} \longleftrightarrow \forall_{\ell} \ln q (x_{n}, y) - \mu_{n}^{\star} \ge \ln q (x_{\ell}, y) - \mu_{\ell}^{\star}$$

$$\wedge \quad \ln q (x_{n}, y) - \mu_{n}^{\star} \ge \ln q_{R}(r, y) - \mu_{R}^{\star}.$$
(27)

The problem then becomes

$$\max_{\{r,p_r\}} \quad V_R(r,p_r) - \Omega(r,p_r), \qquad (28)$$

where $V_R(r, p_r) \equiv F_R(T_R^*)$ takes into account how the optimal assignment reacts to changes in the robot's skills and mass.

I obtain the first order conditions of the problem using the envelope theorem of Milgrom and Segal (2002) and Reynolds' transport theorem.³⁰ I first focus on the derivative of output with respect to the robot's skills:

$$\frac{\partial \left(V_R\left(r, p_r\right) - \Omega\left(r, p_r\right)\right)}{\partial r} = V_R\left(r, p_r\right) \int_{\mathcal{Y}_R} \frac{\partial \ln q_R\left(r, y\right)}{\partial r} dG - \frac{\partial \Omega\left(r, p_r\right)}{\partial r} = 0_{d \times 1}.$$
 (29)

The first term in (29) accounts for the change in output across all tasks assigned to the robot. It gives the net gain in output from a change in the robot's skills as skill mismatch changes with r across tasks. Unlike previous results, all of the tasks assigned to the robot matter, not only those in the boundary of the automated region.

It is convenient to use an explicit functional form for q_R to fix ideas. Assuming that $q_R(r, y)$ is as in (1), the first order condition becomes

$$\frac{\partial \left(F_R\left(\mu^{\star}\left(r, p_r\right), r\right) - \Omega\left(r, p_r\right)\right)}{\partial r} = 2F_R D_R\left(\frac{a_x}{2} - A\left(r - b_R\right)\right) - \frac{\partial \Omega\left(r, p_r\right)}{\partial r} = 0_{d \times 1}, \quad (30)$$

 $^{^{29}}$ This strategy has been exploited by the optimal sensor placement literature under quadratic loss functions; see Aurenhammer et al. (1998, Thm. 1) and Xin et al. (2016, Thm. 1).

 $^{^{30}}$ See Xin et al. (2016) for further applications of power diagrams with capacity constraints. Proposition 5 in Appendix B provides an alternative derivation for the result based on de Goes et al. (2012). The alternative proof is lengthier but more explicit, making it clear how changing the robot's skills affects output.

where $b_R = \frac{\int_{\mathcal{Y}_R} y dG}{D_R}$ is the centroid (or barycenter) of the automated area. Absent other considerations, it is optimal to set the robot's skills to the centroid of the automated region as this minimizes the (quadratic) loss from skill mismatch, thus maximizing the robot's output.³¹ The robot's skills deviate from the centroid to account for gains from having higher skills (a_x) and for the cost of automation ($\partial \Omega(r, p_r)/\partial r$).

The first order condition with respect to p_r takes the usual form of equating marginal product to marginal cost. As in (15), the marginal product is $MP_R = F_R \mu_R^*$:

$$F_R(T_R^{\star})\,\mu_R^{\star} - \frac{\partial\Omega\left(r, p_r\right)}{\partial p_r} = 0. \tag{31}$$

The first order condition is descriptive of the properties that the robot's skills must satisfy relative to the automated region, but it does not pin down the set of tasks to be automated. The automation problem in (24) is not concave in r, and thus condition (29) is only necessary and not sufficient (Urschel, 2017). Regardless, the problem can be solved numerically using a version of Lloyd's algorithm (Lloyd, 1982). This algorithm has been proven to converge monotonically to a local minimum of the objective function (Du et al., 2010). Urschel (2017) gives sufficient conditions for convergence to a global minimum.³²

Figure 6 presents the solution to the automation problem, assuming that q and q_R satisfy Assumption 2 and the automation cost is $\Omega(r) = r'A_R r$. Both figure panels differ only on the weights of cognitive and manual skills in the automation cost function. The two examples in the figure capture a general feature of the automation problem: robots are designed to replace workers in performing tasks where skill mismatch is high. These tasks are located along the boundaries of occupations, and thus the "core" tasks are less likely to be automated because the worker performing those tasks is best suited to do so.

 $^{^{31}}$ This result is shared by the literature on the optimality of centroidal Voronoi diagrams. It is also used in K-means and other vector quantization methods.

 $^{^{32}}$ In practice there are only finitely many candidates for a global minimum, making the selection of the solution simple. It is optimal to automate tasks around one of the vertices of the partition induced by the initial assignment (without automation). Aurenhammer (1987) shows there at most 2n–5 of these vertices in a diagram when the production function is quadratic in x and y, d = 2 and $n \ge 3$.



Figure 6: Directed Automation Example: Quadratic Automation Cost

Note: The figure shows the result of the automation problem, taking the robot's mass as given in a two-dimensional skill space (cognitive and manual skills). Three types of workers are considered $\{x_1, x_2, x_3\}$ with mass $P = \{0.5, 0.3, 0.2\}$. Tasks are uniformly distributed over the unit square; i.e., $\mathcal{Y} = [0,1]^2$ and g(y) = 1. The production function q is as in (1) with $A = I_2$, $a_x = [0.2, 0.1]'$, and $a_y = [0,0]'$. The automation cost function is $\Omega(r) = r'A_R r$, with A_R diagonal. The mass of the robot is fixed at $p_r = 0.05$. The assignment without the robot is presented in gray.

Panel 6a assumes symmetric weights, and thus it is optimal to automate the tasks around the center vertex of the original assignment. However, because endowing the robot with high cognitive and manual skills is costly, it is not optimal to place it in the automated area. The introduction of the robot displaces all three workers with respect to the original assignment, but only worker x_1 is displaced after reassignment. This shows how it can be optimal to design robots that are themselves mismatched for the tasks they perform. Because of this, output increases but mismatch remains.

Panel 6b assumes a higher weight on automating cognitive tasks. It is now optimal to automate the tasks around the vertex formed by \mathcal{Y}_1 , \mathcal{Y}_2 , and the boundary of the task space. These tasks involve less cognitive skills, so it is possible to locate the robot's skills closer to the centroid of the automated region. As in panel 6a, automation induces a reassignment of tasks along the \mathcal{Y}_1 boundary toward more productive workers. Wages and the labor share The effect of automation on wages is ambiguous. First, automation causes tasks to be reassigned. As a result, the workers previously performing the automated tasks are not the only ones affected. The reassignment weakly increases the mismatch between workers and tasks. Introducing the robot relaxes the assignment problem and weakly decreases the value of the multipliers associated with each worker; i.e., $\mu_n^* \leq \lambda_n^*$ for $n = \{1, \ldots, N\}$. This increase in mismatch reduces marginal products and wages. Second, automation reduces the skill mismatch for the tasks being automated, increasing overall output. This then increases the marginal product and wages of all workers.

Whether or not wages decrease depends on how productive robots are at the tasks they overtake (Acemoglu and Restrepo, 2018a). A major increase in productivity due to automation can increase workers' marginal product, increasing wages, while moderate increases in output from the automated tasks can be dominated by the higher mismatch experienced by workers, ultimately reducing their wages.

Regardless of the change in wages, the labor share decreases as the value of the multipliers, μ^* , and the total demand for labor, $G(\mathcal{Y} \setminus \mathcal{Y}_R)$, decrease.

Lemma 1. The labor share $LS = \sum_{n=1}^{N} w_n D_n / F(T_R)$ decreases with automation.

Proof. The decrease in the total demand for labor, $G(\mathcal{Y} \setminus \mathcal{Y}_R) < G(\mathcal{Y})$, and (weak) increase in output, $F(T_R^{\star}) \geq F(T^{\star})$, follow from the automation problem in (24). The effect on the labor share follows from its definition in (20).

Worker training The same tools developed for analyzing the automation problem can be used to study the problem of optimal worker training. Formally, the problem of training worker n by choosing new skills $\tilde{x} \in S$ is

$$\max_{\{\tilde{x},\tilde{T}\}} F\left(\tilde{T},\tilde{x}\right) - \Gamma\left(\tilde{x}|x_n, p_n\right) \quad \text{s.t. } \forall_{\ell} D_{\ell} \le p_{\ell},$$
(32)

where the cost of changing skills (Γ) depends on the workers' current skills and mass. I discuss the solution to this problem in Appendix B.5.

The worker training problem is particularly useful when thinking about the introduction of new tasks. Workers are unlikely to be well suited for performing new tasks, increasing mismatch at early stages of adoption. It is then optimal to train them. The introduction of new technologies, like computers and information technologies, will change occupations directly by modifying the tasks carried out by workers. The training process that follows will further change occupations as workers' skills adapt.

4.2 Skill-biased technology

Technical change can also complement workers' skills. This is the case with the introduction of software that complements cognitive over manual skills in the completion of tasks, or heavy machinery, such as cranes, that complements dexterity over brawn. Unlike automation, this type of technical change affects the productivity of workers across tasks without displacing them. But as with automation, technical change is followed by a reassignment of tasks.

Skill-biased technical change can also be directed. It is optimal to favor the skills with the lowest mismatch, concentrating technology on enhancing the skills the workforce already excels at. This contrasts with the way in which automation is directed. Instead of replacing workers at the tasks they are ill-suited for, technology enhances the worker's productivity by increasing the weight of the skills with a better match while reducing the importance of the skills that the workforce lacks.

To make the discussion precise, I impose additional structure on how skill mismatch affects production. Consider two skills, cognitive and manual, and a production technology q as in (1) where the relative importance of skills is governed by a diagonal matrix A =diag $(\alpha, 1 - \alpha)$, with $\alpha \in [0, 1]$. A higher α makes cognitive match more important for production while simultaneously reducing the importance of manual skill match.

The problem is to choose the value of α , taking into account changes in the assignment of tasks to workers. I assume that the cost of changing α is proportional to output; this simplifies the analysis by making the magnitude of the cost comparable to the gains from additional production and captures the idea that changing more productive processes is more expensive than changing less productive ones. The problem is

$$\max_{\{T,\alpha\}} F(T,\alpha) - \Upsilon(\alpha) F(T,\alpha) \qquad \text{s.t. } \forall_n D_n \le p_n.$$
(33)

The optimality condition for α is obtained using the same techniques as before:

$$(M_m - M_c) - \frac{\partial \Upsilon(\alpha)}{\partial \alpha} \ge 0, \tag{34}$$

where M_s is total mismatch in skill s: $M_s \equiv \sum_{n=1}^{N} \int_{\mathcal{Y}_n} (x_{n,s} - y_s)^2 dy$. The first term captures how much production would increase if α increases. The net gain in production is determined by the difference in total mismatch by skill under the current assignment. If there is more mismatch in the manual dimension $(M_m > M_c)$, it is optimal to direct technical change toward cognitive skills by increasing α . In this way, technology reinforces the workforce's bias by giving more weight to skills for which there is a better match. Absent that cost, it is optimal to shift all the weight toward one of the skills, specializing production to depend only on the skill with the lowest mismatch.

Figure 7 shows how the assignment of tasks to workers changes when the weight of cognitive skills increases. The boundaries of occupations shift and become less sensitive to differences in manual skills, discriminating across workers based on differences in their cognitive skills (as $\alpha \rightarrow 1$, the boundaries become vertical). As this happens, worker x_3 becomes less substitutable with others because of the differences in their cognitive skills; recall from Proposition 3 that the elasticity of substitution decreases with the weighted distance between workers' skills. In contrast, workers x_1 and x_2 become more substitutable because they differences in their manual skills, which are now down weighted.



Figure 7: Assignment Example: Increase in the Weight of Cognitive Skills

Note: The figure shows the result of an increase in the weight of cognitive skills in a two-dimensional skill space (cognitive and manual skills). Three types of workers are considered $\{x_1, x_2, x_3\}$ with mass $P = \{0.5, 0.3, 0.2\}$. Tasks are uniformly distributed over the unit square; i.e., $\mathcal{Y} = [0, 1]^2$ and g(y) = 1. The production function q is as in (1) with $A = diag(\alpha, 1 - \alpha)$.

5 Unassigned tasks

I now extend the model to allow tasks to be left unassigned. I consider first the model without automation and contrast the result with those in Section 2, then I introduce automation. Having unassigned tasks also allows automation technology to improve productivity without replacing workers by automating tasks not otherwise assigned. The number of tasks left unassigned depends on the productivity of workers relative to their outside option.

5.1 Optimal assignment with unassigned tasks

I modify the aggregator in equation (4) to only aggregate the output of assigned tasks:

$$F_{\emptyset}(T) = \exp\left(\int_{\mathcal{Y}\setminus\mathcal{Y}_{\emptyset}} \ln q \left(T\left(y\right), y\right) \, dG\right) - 1,\tag{35}$$

where $T: \mathcal{Y} \to \mathcal{X} \cup \{\emptyset\}$ so that tasks can be unassigned and $\mathcal{Y}_{\emptyset} = T^{-1}(\{\emptyset\})$ denotes the set of unassigned tasks. F_{\emptyset} expresses output relative to no assignment $(T(y) = \emptyset$ for all y) so that if all tasks are left unassigned, output is zero. It is immediate that (35) is equivalent to having $q(\emptyset, y) = 1$ in the original formula (4), extending T to take values over \mathcal{X} and the unassigned option. That way, unassigned tasks do not add to the integral, obtaining (35) as a result. Adopting this convention turns out to be useful because it allows me to apply Proposition 1.

The main difference with the results of Section 2 is that the level of the worker's outside option (\underline{w}) affects the assignment. To simplify calculations, I further assume that the outside option is a fraction $\underline{\lambda}$ of total output: $\underline{w}(T) = \underline{\lambda}F(T)$.³³ Then, there exists a vector $\lambda^* \in \mathbb{R}^N_+$ such that min $\lambda_n^* = 0$ and occupations are given by

$$\mathcal{Y}_{n} = \left\{ y \in \mathcal{Y} \,|\, \forall_{\ell} \, \ln q \, (x_{n}, y) - \lambda_{n}^{\star} \geq \ln q \, (x_{\ell}, y) - \lambda_{\ell}^{\star} \quad \wedge \quad \ln q \, (x_{n}, y) - \lambda_{n}^{\star} \geq \underline{\lambda} \right\}.$$
(36)

This condition differs from (11) in the introduction of the second inequality, which compares the output of worker n in the task with the minimum payment they must receive. This ensures that assigning the task is profitable. The unassigned tasks are

$$\mathcal{Y}_{\emptyset} = \{ y \in \mathcal{Y} \, | \, \forall_n \, \ln q \, (x_n, y) - \lambda_n^{\star} < \underline{\lambda} \} \,. \tag{37}$$

The higher the $\underline{\lambda}$, the fewer the tasks assigned for production. A worker's marginal product is given as in equation (15), and their wage is $w_n = \lambda_n^* F(T^*) + \underline{w} = (\lambda_n^* + \underline{\lambda}) F(T^*).$

To fix ideas, consider q as in (1), depending on the quadratic mismatch between the worker's and task's skills. This provides a clear geometrical interpretation for which tasks are left unassigned. Workers will be assigned to a task only if the mismatch is no greater than $a'_x x + a'_y y - \underline{\lambda}$ as this guarantees that the worker produces enough output for it to cover the workers' outside option.³⁴ However, the condition does not imply that the task will be assigned to the worker. This depends on the comparison between workers' productivity as

³³Without this assumption, it is not possible to determine the value of λ independently of the assignment T. The term $\underline{\lambda}$ in (36) must be replaced by $\underline{w}/F(T)$.

³⁴With $a_y = 0$, a task will be assigned only if it lies in an ellipse of radius $\sqrt{a'_x x}$ around the worker's skills. The ellipse's shape depends on the weights in matrix A.



Figure 8: Assignment Example: Unassigned Tasks

in Section 2 (see the first inequality in equation 36).

Which tasks to perform will depend critically on which tasks are more productive given the current technology. This idea is captured by a_y , which determines which tasks generate more output regardless of which worker performs them. A higher cognitive weight in a_y makes cognitive-intensive tasks more likely to be performed.

Figure 8 shows the optimal assignment, with the gray areas representing unassigned tasks. As in Section 4.1, these tasks are located along the boundaries of the task space and the vertices of the assignment. The two panels differ on the weight of task skills in determining the output of a task, as measured by a_y . I assume that $a_y = \overline{a}_y [\alpha_y, 1 - \alpha_y]'$ and vary the relative importance of skills by choosing $\alpha_y \in [0, 1]$. A higher value of α_y makes cognitive-intensive tasks more productive.

Panel 8a presents the assignment under equal skill weights in a_y . Most of the unassigned tasks involve high cognitive skills because performing cognitive-intensive tasks comes at the cost of greater mismatch for workers x_1 and x_2 . In Panel 8b α_y increases, making cognitive-

Note: The figure shows the assignment in a two-dimensional skill space (cognitive and manual skills). Three types of workers are considered $\{x_1, x_2, x_3\}$ with mass $P = \{0.5, 0.3, 0.2\}$. Tasks are uniformly distributed over the unit square; i.e., $\mathcal{Y} = [0, 1]^2$ and g(y) = 1. The production function q is as in (1) with $A = I_2$, $a_x = [0.2, 0.1]'$, and $a_y = \overline{a}_y [\alpha_y, 1 - \alpha_y]'$, with $\overline{a}_y \in \mathbb{R}_+$ and $\alpha_y \in [0, 1]$. The worker's outside option is 0.

intensive tasks more productive and manual-intensive tasks less productive. This makes it worthwhile to reassign workers toward cognitive-intensive tasks. In both cases, only worker x_1 is not fully assigned. x_1 is the least productive worker type.

The role of the outside option An increase in the value of the outside option, \underline{w} , reduces employment by limiting the set of tasks that are profitable to produce. The net effect on wages is nevertheless ambiguous. As the assignment of tasks changes, so does the mismatch of workers in their boundary tasks. Mismatch necessarily decreases for the workers who are not fully assigned. This can cause the wage distribution to become compressed because wages reflect marginal productivity relative to the least productive workers (see equation 17). These effects must be weighted against the increase in the level of wages coming from \underline{w} to determine the net effect of the increase of the outside option.³⁵

5.2 Automation and unassigned tasks

The characterization of automation as a worker-replacing technology changes once tasks can be left unassigned. It is now possible to direct automation toward unassigned tasks, that is, tasks that are not worthwhile for workers to perform (because of low productivity) or tasks for which workers lack the appropriate skills (high mismatch). Moreover, performing additional tasks increases output, potentially raising workers' marginal products and wages.³⁶ The gains in output and worker productivity from automating previously unassigned tasks suggest that automation is more profitable in economies where workers have high outside options, for example, where minimum wage regulations imply high wages or there is a strong safety net in case of non-employment.

³⁵The effects on the wage distribution are similar to those documented by Engbom and Moser (2018).

³⁶These effects are similar in the model to the effects of immigration. The addition of new workers can in principle displace current workers from the tasks they perform; however, the new workers can also perform tasks that were previously left unassigned or tasks for which native workers were not well suited. This includes both low- and high-skill tasks. Consequently, the addition of new workers through immigration can cause an increase in wages via the reduction in mismatch in the occupations of native workers and the increase in productivity by making possible the completion of previously unassigned tasks.



Figure 9: Assignment Example: Unemployment and Automation

Note: The figures show the assignment in a two-dimensional skill space (cognitive and manual skills). Three types of workers are considered $\{x_1, x_2, x_3\}$ with mass $P = \{0.5, 0.3, 0.2\}$. Tasks are uniformly distributed over the unit square; i.e., $\mathcal{Y} = [0, 1]^2$ and g(y) = 1. The production function q is as in (1) with $A = I_2$, $a_x = [0.2, 0.1]'$, and $a_y = \overline{a}_y [\alpha_y, 1 - \alpha_y]'$, with $\overline{a}_y \in \mathbb{R}_+$ and $\alpha_y = 1/2$. The worker's outside option is 0. The automation cost function is $\Omega(r) = r'A_Rr$, with A_R diagonal. The robot's mass is fixed at $p_r = 0.03$. The assignment without the robot is presented in gray.

In general, it is optimal to automate unassigned tasks along the boundaries of occupations. As a result, automation ends up partially displacing workers. To illustrate this, I expand the example in Figure 8 by solving the optimal automation problem. The resulting assignment is presented in Figure 9. The robot is placed so as to automate part of the cognitive-/manual-intensive tasks that were unassigned. The two panels also show how the incentives for automation change with the production technology. If technology favors cognitive-intensive tasks over manual-intensive ones, workers are reassigned away from the latter and into the former (see Figure 8). Consequently, production can be increased by directing automation toward manual-intensive tasks in a way that avoids disrupting the optimal assignment of tasks to workers. In this scenario, technological change makes new tasks available for workers, while automation follows by taking over tasks that are no longer worthwhile for them to perform.³⁷

³⁷This process is similar to the one in Acemoglu and Restrepo (2018b). Changes in technology lead to a reassignment of workers toward more complex (and newer) tasks, while relatively simpler (and older) tasks are automated, displacing workers in the process.

6 Concluding remarks

In this paper, I develop a framework to study occupations, where production occurs by assigning tasks to workers in a multidimensional setting. Occupations arise from the assignment process instead of being taken as a preexisting feature of production, allowing them to change endogenously in response to changes in the economic environment. This flexibility is important when addressing the consequences of worker-replacing technologies like automation or offshoring. These technologies replace workers in some, but not all, of the tasks they perform, transforming occupations.

The model makes precise the role of tasks in defining the marginal product, compensation, and substitutability of workers. All these properties are shaped by how productive workers are at the tasks along the boundaries of their occupations. These are the tasks for which the workers are the least productive and at which they are directly substitutable for other workers. The model also makes it possible to ask about the optimal direction of automation, i.e., which type of tasks should be automated.

References

- Acemoglu, D. and D. Autor (2011). Skills, Tasks and Technologies: Implications for Employment and Earnings. In D. Card and O. Ashenfelter (Eds.), *Handbook of Labor Economics*, Volume 4 of *Handbook of Labor Economics*, Chapter 12, pp. 1043 – 1171. Elsevier.
- Acemoglu, D. and P. Restrepo (2018a). Modeling Automation. AEA Papers and Proceedings 108, 48–53.
- Acemoglu, D. and P. Restrepo (2018b, June). The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment. American Economic Review 108(6), 1488–1542.
- Acemoglu, D. and P. Restrepo (2020). Robots and jobs: Evidence from us labor markets. Journal of Political Economy 128(6), 2188–2244.
- Acemoglu, D. and P. Restrepo (2021, June). Tasks, automation, and the rise in us wage inequality. Working Paper 28920, National Bureau of Economic Research.
- Adenbaum, J. (2021). Endogenous firm structure and worker specialization. https://jacobadenbaum.github.io/files/Adenbaum_JMP.pdf.
- Aghion, P., B. F. Jones, and C. I. Jones (2019, January). Artificial Intelligence and Economic Growth, Book 9, pp. 237–282. University of Chicago Press.
- Atalay, E., P. Phongthiengtham, S. Sotelo, and D. Tannenbaum (2020, April). The evolution of work in the united states. *American Economic Journal: Applied Economics* 12(2), 1–34.
- Aurenhammer, F. (1987, feb). Power Diagrams: Properties, Algorithms and Applications. SIAM Journal on Computing 16(1), 78–96.
- Aurenhammer, F., F. Hoffmann, and B. Aronov (1998, Jan). Minkowski-Type Theorems and Least-Squares Clustering. Algorithmica 20(1), 61–76.
- Autor, D. H. (2013, Sep). The "Task Approach" to Labor Markets: An Overview. Journal for Labour Market Research 46(3), 185–199.
- Autor, D. H., L. F. Katz, and M. S. Kearney (2006, May). The polarization of the u.s. labor market. American Economic Review 96(2), 189–194.
- Autor, D. H., F. Levy, and R. J. Murnane (2003). The Skill Content of Recent Technological Change: An Empirical Exploration. The Quarterly Journal of Economics 118(4), 1279– 1333.
- Baley, I., A. Figueiredo, and R. Ulbricht (2021). Mismatch cycles. https://crei.cat/ wp-content/uploads/2021/09/.
- Baqaee, D. R. and E. Farhi (2019). The macroeconomic impact of microeconomic shocks: Beyond hulten's theorem. *Econometrica* 87(4), 1155–1203.

- Black, S. E. and A. Spitz-Oener (2010). Explaining Women's Success: Technological Change and the Skill Content of Women's Work. *The Review of Economics and Statistics* 92(1), 187–194.
- Blackorby, C. and R. R. Russell (1981). The Morishima Elasticity of Substitution; Symmetry, Constancy, Separability, and its Relationship to the Hicks and Allen Elasticities. *The Review of Economic Studies* 48(1), 147–158.
- Blackorby, C. and R. R. Russell (1989). Will the Real Elasticity of Substitution Please Stand Up? (A Comparison of the Allen/Uzawa and Morishima Elasticities). The American Economic Review 79(4), 882–888.
- Blinder, A. (2009). How Many US Jobs Might be Offshorable? World Economics 10(2), 41-78.
- Blinder, A. S. and A. B. Krueger (2013). Alternative Measures of Offshorability: A Survey Approach. *Journal of Labor Economics* 31(2), S97–S128.
- Carlier, G. (2003). Duality and Existence for a Class of Mass Transportation Problems and Economic Applications, pp. 1–21. Tokyo: Springer Japan.
- Caunedo, J., E. Keller, and Y. Shin (2021, April). Technology and the task content of jobs across the development spectrum. Working Paper 28681, National Bureau of Economic Research.
- Cavounidis, C., V. Dicandia, K. Lang, and R. Malhotra (2021, September). The evolution of skill use within and between jobs. Working Paper 29302, National Bureau of Economic Research.
- Chiappori, P.-A., S. Oreffice, and C. Quintana-Domeque (2012). Fatter attraction: Anthropometric and socioeconomic matching on the marriage market. *Journal of Political Economy* 120(4), 659–695.
- Combemale, C., K. S. Whitefoot, L. Ales, and E. R. H. Fuchs (2021a). How itâs made: A general theory of the labor implications of technological change. https://www.andrew.cmu.edu/user/ales/acfw.pdf.
- Combemale, C., K. S. Whitefoot, L. Ales, and E. R. H. Fuchs (2021b, 08). Not all technological change is equal: how the separability of tasks mediates the effect of technology change on skill demand. *Industrial and Corporate Change dtab026*, 1–27. dtab026.
- de Goes, F., K. Breeden, V. Ostromoukhov, and M. Desbrun (2012, November). Blue Noise Through Optimal Transport. ACM Transactions on Graphics 31(6), 171:1–171:11.
- Deming, D. J. (2017). The Growing Importance of Social Skills in the Labor Market. *The Quarterly Journal of Economics* 132(4), 1593–1640.

- Du, Q., M. Gunzburger, and L. Ju (2010). Advances in Studies and Applications of Centroidal Voronoi Tessellations. Numerical Mathematics: Theory, Methods and Applications 3(2), 119–142.
- Engbom, N. and C. Moser (2018). Earnings Inequality and the Minimum Wage: Evidence from Brazil. Institute Working Paper 7, Opportunity and Inclusive Growth Institute, Federal Reserve Bank of Minneapolis.
- Feenstra, R. C. and J. A. Levinsohn (1995). Estimating Markups and Market Conduct with Multidimensional Product Attributes. The Review of Economic Studies 62(1), 19–52.
- Fogel, J. and B. Modenesi (2021). What is a labor market? classifying workers and jobs using network theory. https://sites.google.com/view/jamiefogel/research.
- Frey, C. B. and M. A. Osborne (2017). The Future of Employment: How Susceptible Are Jobs to Computerisation? *Technological Forecasting and Social Change* 114, 254 280.
- Galichon, A. (2016). Optimal Transport Methods in Economics. Princeton University Press.
- Gathmann, C. and U. Schönberg (2010). How General Is Human Capital? A Task-Based Approach. *Journal of Labor Economics* 28(1), 1–49.
- Gola, P. (2021). Supply and demand in a two-sector matching model. *Journal of Political Economy* 129(3), 940–978.
- Grigsby, J. (2021). Skill heterogeneity and aggregate labor marketdynamics. https://sites.google.com/view/john-grigsby/research.
- Guvenen, F., B. Kuruscu, S. Tanaka, and D. Wiczer (2020, January). Multidimensional skill mismatch. *American Economic Journal: Macroeconomics* 12(1), 210–44.
- Haanwinckel, D. (2021). Supply, demand, institutions, and firms: A theory of labor market sorting and the wage distribution. http://www.econ.ucla.edu/haanwinckel/papers/sdif.pdf.
- Heckman, J. and J. Scheinkman (1987). The Importance of Bundling in a Gorman-Lancaster Model of Earnings. *The Review of Economic Studies* 54 (2), 243–255.
- Heckman, J. J. and G. Sedlacek (1985). Heterogeneity, Aggregation, and Market Wage Functions: An Empirical Model of Self-Selection in the Labor Market. *Journal of Political Economy* 93(6), 1077–1125.
- Hemous, D. and M. Olsen (2021). The Rise of the Machines: Automation, Horizontal Innovation and Income Inequality. American Economic Journal: Macroeconomics, Forthcoming Forthcoming, 1–49.
- Kambourov, G. and I. Manovskii (2009). Occupational Specificity of Human Capital. International Economic Review 50(1), 63–115.

- Kantorovich, L. V. (2006, Mar). On the Translocation of Masses. Journal of Mathematical Sciences 133(4), 1381–1382.
- Kelso, A. S. and V. P. Crawford (1982). Job Matching, Coalition Formation, and Gross Substitutes. *Econometrica* 50(6), 1483–1504.
- Koopmans, T. C. and M. Beckmann (1957). Assignment Problems and the Location of Economic Activities. *Econometrica* 25(1), 53–76.
- Kremer, M. (1993). The O-Ring Theory of Economic Development. The Quarterly Journal of Economics 108(3), 551–575.
- Lindenlaub, I. (2017). Sorting Multidimensional Types: Theory and Application. *The Review* of Economic Studies 84(2), 718–789.
- Lise, J. and F. Postel-Vinay (2020, August). Multidimensional skills, sorting, and human capital accumulation. *American Economic Review* 110(8), 2328–76.
- Lloyd, S. P. (1982). Least Squares Quantization in PCM. IEEE Transactions on Information Theory 28, 129–137.
- Mandelbrot, B. (1962, 02). Paretian Distributions and Income Maximization^{*}. The Quarterly Journal of Economics 76(1), 57–85.
- Martinez, J. and J. Moen-Vorum (2021). Automation potential and diffusion. https://www.josebamartinez.me.
- McKinsey Global Institute (2017). What the Future of Work Will Mean for Jobs, Skills and Wages. Technical report, McKinsey Global Institute.
- Milgrom, P. and I. Segal (2002). Envelope Theorems for Arbitrary Choice Sets. *Econometrica* 70(2), 583–601.
- Poletaev, M. and C. Robinson (2008). Human Capital Specificity: Evidence from the Dictionary of Occupational Titles and Displaced Worker Surveys, 1984-2000. *Journal* of Labor Economics 26(3), 387–420.
- Rosen, S. (1978). Substitution and Division of Labour. *Economica* 45(179), 235–250.
- Roy, A. D. (1951). Some Thoughts on the Distribution of Earnings. Oxford Economic Papers 3(2), 135–146.
- Sattinger, M. (1975). Comparative Advantage and the Distributions of Earnings and Abilities. *Econometrica* 43(3), 455–468.
- Sattinger, M. (1993). Assignment Models of the Distribution of Earnings. Journal of Economic Literature 31(2), 831–880.
- Scarsini, M. (1984). On measures of concordance. Stochastica 8(3), 201–218.

- Sklar, M. (1959). Fonctions de répartition à n dimensions et leurs marges. Technical report, Publications de l'Institut Statistique de l'Université de Paris 8.
- Sobel, J. (1992). How to Count to One Thousand. The Economic Journal 102(410), 1–8.
- Spitz-Oener, A. (2006). Technical Change, Job Tasks, and Rising Educational Demands: Looking outside the Wage Structure. *Journal of Labor Economics* 24 (2), 235–270.
- Stokey, N. L. (2018). Technology, skill, and the wage structure. Journal of Human Capital 12(2), 343–384.
- Tan, J. (2021). Multidimensional heterogeneity and matching in africtional labor market an application to polarization. https://sites.google.com/site/joanneyumintanphd/ research.
- Teulings, C. N. (1995). The wage distribution in a model of the assignment of skills to jobs. Journal of Political Economy 103(2), 280–315.
- Tinbergen, J. (1956). On the Theory of Income Distribution. Weltwirtschaftliches Archiv 77, 155–175.
- Urschel, J. C. (2017). On The Characterization and Uniqueness of Centroidal Voronoi Tessellations. SIAM Journal on Numerical Analysis 55(3), 1525–1547.
- Villani, C. (2009). *Optimal Transport: Old and New*. Grundlehren der mathematischen Wissenschaften. Springer Berlin Heidelberg.
- Webb, M. (2020). The impact of artificial intelligence on the labor market. https://www.michaelwebb.co/webb_ai.pdf.
- Willis, R. J. and S. Rosen (1979). Education and Self-Selection. Journal of Political Economy 87(5, Part 2), S7–S36.
- Xin, S.-Q., B. Lévy, Z. Chen, L. Chu, Y. Yu, C. Tu, and W. Wang (2016, November). Centroidal Power Diagrams with Capacity Constraints: Computation, Applications, and Extension. ACM Trans. Graph. 35(6), 244:1–244:12.
- Yamaguchi, S. (2012). Tasks and Heterogeneous Human Capital. Journal of Labor Economics 30(1), 1–53.

Appendix for Online Publication

Table of Contents

Α	Mathematical Preliminaries	45
В	Proofs	48
	B.1 Existence and uniqueness of optimal task assignment	48
	B.2 Marginal product	53
	B.3 Differentiability of demand	56
	B.4 Directed automation	57
	B.5 Worker training	60
С	Multi-Sector Assignment	61
D	Cosine Similarity Assignments	63

A Mathematical Preliminaries

Definition A.1. [Probability Space] A probability space is a triplet (A, \mathcal{A}, μ) of a set A, a σ -algebra \mathcal{A} on that set and a probability measure $\mu : \mathcal{A} \to [0, 1]$. When the σ -algebra is understood (generally as the Borel σ -algebra) it is omitted.

Definition A.2. [Polish Space] A set A is a polish space if it is separable (allows for a dense countable subset) and metrizable topological space (there exists at least one metric that induces the topology).

Definition A.3. [Coupling] Let (\mathcal{Y}, G) and (\mathcal{X}, P) be two probability spaces. A coupling π of G and P is a joint distribution on $(\mathcal{X} \times \mathcal{Y})$ such that $\int_{\mathcal{X} \times Y} d\pi(x, y) = G(Y)$ for all $Y \in \mathcal{B}(\mathcal{Y})$ and $\int_{X \times \mathcal{Y}} d\pi(x, y) = P(X)$ for all $X \in \mathcal{B}(\mathcal{X})$, where $\mathcal{B}(A)$ denotes the Borel sets of A. So π gives G and P as marginals. Let $\Pi(P, G)$ be the set of all couplings of P and G. When the assignment is given by an assignment function the coupling is deterministic.

Definition A.4. [*h*-transform] Let $h : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}$ be a function. The *h*-transform of a function $f : \mathcal{X} \to \mathbb{R}$ is:

$$f^{h}(y) = \sup_{x \in \mathcal{X}} \left\{ h(x, y) - f(x) \right\}.$$

Definition A.5. [*h*-convex] A function $f : \mathcal{X} \to \mathbb{R}$ is said to be *h*-convex if there exists a function $g : \mathcal{Y} \to \mathbb{R}$ such that:

$$f(x) = \sup_{y \in \mathcal{Y}} \left\{ h(x, y) - g(y) \right\}.$$

Definition A.6. [*h*-subdifferential] The *h*-subdifferential of a function $v : \mathcal{Y} \to \mathbb{R}$ is defined as the set $\partial^h v(y) = \{x \in \mathcal{X} \mid v(y) + v^h(x) = h(x, y)\}.$

The following theorem joins results from optimal transport on the existence of a solution to the Monge-Kantorovich problem and the applicability of Kantorovich's duality to the mass transportation problem:

Theorem A.1. Villani (2009, Thm. 5.10 and Thm 5.30) Let (\mathcal{Y}, G) and (\mathcal{X}, P) be two Polish probability spaces and let $h : \mathcal{X} \times \mathcal{Y} \to \mathbb{R} \cup \{-\infty\}$ be an upper semicontinuous function.

Consider the optimal transport problem:

$$\sup_{\pi\in\Pi(G,P)}\int_{\mathcal{X}\times\mathcal{Y}}h\left(x,y\right)d\pi\left(x,y\right),$$

where function h(x, y) describes the gain (or surplus) of transporting a unit of mass from y to x, and $\Pi(G, P)$ denotes the set of couplings of G and P.

If there exist real valued lower semicontinuous functions $a \in L^{1}(P)$ and $b \in L^{1}(G)$:

$$\forall_{(x,y)\in\mathcal{X}\times\mathcal{Y}} \quad h\left(x,y\right) \leq a\left(x\right) + b\left(y\right),$$

then:

1. There is duality:

$$\begin{split} \sup_{\pi \in \Pi(G,P)} \int_{\mathcal{X} \times \mathcal{Y}} h\left(x,y\right) d\pi\left(x,y\right) &= \inf_{\substack{(\lambda,v) \in L^{1}(P) \times L^{1}(G) \\ \lambda(x) + v(y) \ge h(x,y)}} \int_{\mathcal{X}} \lambda\left(x\right) dP\left(x\right) + \int_{\mathcal{Y}} v\left(y\right) dG\left(y\right) \\ &= \inf_{w \in L^{1}(P)} \int_{\mathcal{X}} \lambda\left(x\right) dP\left(x\right) + \int_{\mathcal{Y}} \lambda^{h}\left(y\right) dG\left(y\right) \\ &= \inf_{v \in L^{1}(G)} \int_{\mathcal{X}} v^{h}\left(x\right) dP\left(x\right) + \int_{\mathcal{Y}} v\left(y\right) dG\left(y\right), \end{split}$$

where $\Pi(G, P)$ is the set of couplings of G and P and f^h denotes the h-transform of function f:

$$f^{h}(y) = \sup_{x \in \mathcal{X}} h(x, y) - f(x).$$

The functions w and v are h-convex since they are the h-transform of one another.

- 2. If, furthermore, h is real valued $(h : \mathcal{X} \times \mathcal{Y} \to \mathbb{R})$ and the solution to the Monge-Kantorovich problem is finite $\left(\max_{\pi \in \Pi(G,P)} \int_{\mathcal{X} \times \mathcal{Y}} h(x,y) d\pi(x,y) < \infty\right)$ then there is a measurable h-monotone set $\Gamma \subset \mathcal{X} \times \mathcal{Y}^{38}$ such that for any $\pi \in \Pi(G,P)$ the following statements are equivalent:
 - (a) π is optimal.
 - (b) π is h-cyclically monotone.
 - (c) There is a h-convex function λ such that $\lambda(x) + \lambda^h(y) = h(x, y) \pi$ -almost surely.
 - (d) There exist $\lambda : \mathcal{X} \to \mathbb{R}$ and $v : \mathcal{Y} \to \mathbb{R}$ such that $\lambda(x) + v(y) \ge h(x, y)$ with equality π -almost surely.
 - (e) π is concentrated in Γ .
- 3. If, h is real valued $(h : \mathcal{X} \times \mathcal{Y} \to \mathbb{R})$ and there are functions $c \in L^1(P)$ and $d \in L^1(G)$ such that:

$$\forall_{(x,y)\in\mathcal{X}\times\mathcal{Y}} \quad c(x) + d(y) \le h(x,y).$$

then the dual problem has a solution. There is a function w that attains the infimum.

- 4. (this part from Villani (2009, Thm. 5.30))If:
 - (a) h is real valued $(h : \mathcal{X} \times \mathcal{Y} \to \mathbb{R})$
 - (b) the solution to the Monge-Kantorovich problem is finite:

$$\max_{\pi \in \Pi(G,P)} \int_{\mathcal{X} \times \mathcal{Y}} h(x,y) \, d\pi(x,y) < \infty$$

(c) for any h-convex function $v : \mathcal{Y} \to \mathbb{R} \cup \{-\infty\}$ the subdifferential $\partial^h v(y)$ is single valued G-almost everywhere

 $^{^{38}\}text{If}\;a,\,b$ and h are continuous then Γ is closed.

Then

- (a) there is a unique (in law) optimal coupling π of (G, P).
- (b) the optimal coupling is deterministic: $T: \mathcal{Y} \to \mathcal{X}$.
- (c) the optimal coupling is characterized by the existence of a function h- convex function v such that $T(y) = \partial^h v(y)$.

Finally, Reynolds' transport theorem is used extensively in the text:

Theorem A.2. [Reynolds' Transport Theorem] The rate of change of the integral of a scalar function f within a volume V is equal to the volume integral of the change of f, plus the boundary integral of the rate at which f flows though the boundary ∂V of outward unit normal n:

$$\nabla \int_{V} f(x) \, dV = \int_{V} \nabla f(x) \, dV + \int_{\partial V} f(x) \, (\nabla x \cdot n) \, dA.$$

B Proofs

B.1 Existence and uniqueness of optimal task assignment

Outline of the poof I first relax the problem in (7) to allow for non-deterministic assignments, see Kantorovich (2006) and Koopmans and Beckmann (1957). An assignment is then a joint measure over workers/task pairs: $\pi : \mathcal{X} \times \mathcal{B}(\mathcal{Y}) \to \mathbb{R}_+$, where $\mathcal{B}(\mathcal{Y})$ denotes the Borel sets of \mathcal{Y} , and there are measures P and G so that (\mathcal{X}, P) and (\mathcal{Y}, G) are two Polish spaces. An assignment π is deemed feasible if it is a coupling of measures P and G, see definition A.3 in Appendix A. In terms of the assignment problem, π must guarantee that workers have enough time to perform all the time demanded by their occupations, and each task is completed at most once. Letting $\Pi(P, G)$ be the set of feasible assignments:

$$\pi \in \Pi(P,G) \longleftrightarrow \forall_n \int_{\mathcal{Y}} d\pi(x_n,y) \le p_n \qquad \forall_{Y \in \mathcal{B}(\mathcal{Y})} \sum_{n=1}^N \int_{y \in Y} d\pi(x_n,y) \le G(Y). \quad (B.1)$$

The second condition can be simplified to: $\sum_{n=1}^{N} \pi(x_n, \{y\}) \leq g(y).$

The problem is now to choose a coupling $\pi \in \Pi(P, G)$ to maximize output. I further simplify the problem by applying natural logarithm to the objective function. Doing so reveals the linearity of the problem in the choice variable π . The relaxed optimization problem is:

$$\max_{\pi \in \Pi(P,G)} \sum_{n=1}^{N} \int_{\mathcal{Y}} \ln q\left(x_n, y\right) \, d\pi\left(x_n, y\right). \tag{B.2}$$

Lemma B.1 applies Theorem 5.10 of Villani (2009) to establish duality for the problem:

$$\max_{\pi \in \Pi} \sum_{n=1}^{N} \int_{\mathcal{Y}} \ln q \left(x_n, y \right) \, d\pi \left(x_n, y \right) = \inf_{\substack{(\lambda, \nu) \in \mathbb{R}^N \times L^1(G) \\ \lambda_n + v(y) \ge \ln q(x_n, y)}} \sum_{n=1}^{N} \lambda_n p_n + \int_{\mathcal{Y}} \max_n \left\{ \ln q \left(x_n, y \right) - \lambda_n \right\} dG, \tag{B.3}$$
$$= \inf_{\lambda \in \mathbb{R}^N} \sum_{n=1}^{N} \lambda_n p_n + \int_{\mathcal{Y}} \max_n \left\{ \ln q \left(x_n, y \right) - \lambda_n \right\} dG,$$

where λ and ν are the multipliers (or potentials) of the problem. Lemma B.2 establishes that a solution to the dual problem (λ^*, ν^*) exists. The levels of λ^* and ν^* are only determined up to an additive constant. Both the assignment and the value of the dual problem do not change if λ is increased by a constant κ for all workers and ν is decreased by the same amount for all tasks. I normalize the value of the minimum λ^* to zero. This is convenient when relating the value of λ^* to the marginal product of workers and the wages in the decentralization of the optimal assignment.

The first two conditions on the production function q ensure that the value of the primal problem (B.2) and the dual problem (B.3) are finite. This is the key step in verifying the conditions for Theorem 5.10 of Villani (2009). In particular, the first condition avoids indeterminacies when evaluating the natural logarithm of q for any worker/task pair.

The solution to the dual problem provides a way to construct the optimal assignment T^* . Lemma B.3 applies Theorem 5.30 of Villani (2009) to construct T^* as the sub-differential of v^* . The third condition on the production function q is crucial to establish single-valuedness of the sub-differential of v^* . This gives the formula for the optimal assignment in (8). Galichon (2016, Ch. 5.3) presents an algorithm to solve the dual problem in (B.3).

I now turn to the general proof of the problem.

General setting Consider the set up of Section 1. There are N types of workers $\{x_1, \ldots, x_N\} \equiv \mathcal{X}$. The mass of workers is described by a (discrete) measure P so that $P(x_n) = p_n$. There is a continuum of tasks $y \in \mathcal{Y}$ distributed continuously according to an absolutely continuous measure $G : \mathcal{B}(\mathcal{Y}) \to \mathbb{R}_+$. \mathcal{Y} is assumed compact. The setup in Section 1 further assumes that there is a space of skills \mathcal{S} and that $\mathcal{X}, \mathcal{Y} \subseteq \mathcal{S}$.

Output is produced by completing tasks. A worker of type x_n performing task y produces $q(x_n, y)$. $q : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}$ is a real-valued function. Output for all worker/task pairs is aggregated into a final good:

$$F(\pi) = \begin{cases} \left(\sum_{n=1}^{N} \int \left(q\left(x_{n}, y\right)\right)^{\frac{\sigma-1}{\sigma}} d\pi\left(x_{n}, y\right)\right)^{\frac{\sigma}{\sigma-1}} & \text{if } \sigma > 1\\ \exp\left(\sum_{n=1}^{N} \int_{\mathcal{Y}} \ln q\left(x_{n}, y\right) d\pi\left(x_{n}, y\right)\right) & \text{if } \sigma = 1 \end{cases},$$

where $\pi \in \Pi(P, G)$ is a coupling of P and G (see definition A.3). The coupling π describes the assignment: a mass $\pi(x, y)$ of workers of type x is assigned to task y.

The problem is to maximize output of the final good by choosing an assignment of tasks to workers π . I first transform the objective function so that the problem takes the form of a Monge-Kantorovich problem:

$$\max_{\pi \in \Pi(P,G)} \int h\left(x, y | \sigma\right) d\pi\left(x, y\right),\tag{B.4}$$

where $h(x, y | \sigma) = \begin{cases} (q(x, y))^{\frac{\sigma-1}{\sigma}} & \text{if } \sigma > 1\\ \ln q(x, y) & \text{if } \sigma = 1 \end{cases}$.

The following proposition establishes duality for this problem:

Lemma B.1. If q satisfies the following properties:

- *i.* $\sigma > 1$ or all workers can produce in some task: $\forall_x \exists y \quad q(x,y) > 0$.
- ii. $q(x, \cdot)$ is upper-semicontinuous in y given $x \in \mathcal{X}$.

Then, the following equalities hold:

$$\max_{\pi \in \Pi(P,G)} \int \left(h\left(x, y | \sigma\right)\right)^{\frac{\sigma-1}{\sigma}} d\pi\left(x, y\right) = \inf_{\substack{(\lambda, v) \in \mathbb{R}^{N} \times L^{1}(G)\\\lambda_{n}+v(y) \ge q(x_{n}, y | \sigma)}} \sum_{n=1}^{N} \lambda_{n} p_{n} + \int_{\mathcal{Y}} v\left(y\right) dG\left(y\right)$$
$$= \inf_{\lambda \in \mathbb{R}^{N}} \sum_{n=1}^{N} \lambda_{n} p_{n} + \int_{\mathcal{Y}} \max_{n} \left\{q\left(x_{n}, y | \sigma\right) - \lambda_{n}\right\} dG\left(y\right)$$

Proof. This follows from applying theorem A.1 (Villani, 2009, Thm. 5.10). Note that $\mathcal{Y} \subset \mathbb{R}^n$ and \mathcal{X} is finite so they are both Polish spaces. $h(x, y|\sigma)$ is upper semicontinuous because $f(x) = x^{\frac{\sigma-1}{\sigma}}$ and $f(x) = \ln x$ are continuous and monotone increasing, and q is upper semicontinuous.

It is left to verify that there exist real valued lower semicontinuous functions $a \in L^{1}(P)$ and $b \in L^{1}(G)$:

$$\forall_{(x,y)\in\mathcal{X}\times\mathcal{Y}} \quad h\left(x,y|\sigma\right) \le a\left(x\right) + b\left(y\right)$$

For this let $a(x) = \max_{y \in \mathcal{Y}} \{h(x, y | \sigma)\}$ and b(y) = 0. The max in the definition of a is well defined because h is upper semicontinuous and \mathcal{Y} is compact, furthermore a is finite (either $\sigma > 1$ or, if $\sigma = 1$, h is finite for at least some y guaranteeing a a finite value). Function a is immediately continuous with respect to the discrete topology. The desired equalities follow from part 1 of Theorem A.1.

The dual problem is then to find a value associated with each type of worker $\{\lambda_1, \ldots, \lambda_N\}$. The problem is:

$$\inf_{\lambda \in \mathbb{R}^N} \sum_{n=1}^N \lambda_n p_n + \int_{\mathcal{Y}} v\left(y\right) dG\left(y\right) \quad \text{where: } v\left(y\right) = \max_{n \in \{1,\dots,N\}} \left\{h\left(x, y | \sigma\right) - \lambda_n\right\}.$$
(B.5)

I show that the dual problem has a solution and I use that solution to construct a solution to the Monge-Kantorovich problem in (B.4). Furthermore, the solution will take the form of a deterministic transport and the implied assignment function is the solution to the problem (7) in the main text. Part 3 of Theorem A.1 establishes that solution to the dual problem (B.5) exists.

Lemma B.2. If q satisfies the following properties:

- *i.* $\sigma > 1$ or all workers can produce in all tasks: q(x, y) > 0 for all pairs $(x, y) \in \mathcal{X} \times \mathcal{Y}$.
- ii. $q(x, \cdot)$ is upper-semicontinuous in y given $x \in \mathcal{X}$.

Then there exists $\lambda^{\star} \in \mathbb{R}^N$ such that:

$$\lambda^{\star} \in \underset{\lambda \in \mathbb{R}^{N}}{\operatorname{argmin}} \sum_{n=1}^{N} \lambda_{n} p_{n} + \int_{\mathcal{Y}} \left(\max_{n \in \{1, \dots, N\}} \left\{ h\left(x, y | \sigma\right) - \lambda_{n} \right\} \right) dG\left(y\right).$$

Proof. This follows from applying part 3 of theorem A.1 (Villani, 2009, Thm. 5.10). The function $h(x, y|\sigma)$ is required to be real valued. When $\sigma > 1$ this is verified since q is real valued. When $\sigma = 1$ it is verified under the additional condition that q(x, y) > 0 for all $(x, y) \in \mathcal{X} \times \mathcal{Y}$.

It is left to find functions $c \in L^{1}(P)$ and $d \in L^{1}(G)$ such that:

$$\forall_{(x,y)\in\mathcal{X}\times\mathcal{Y}} \quad c(x) + d(y) \le h(x,y|\sigma).$$

For this let c(x) = 0 and $d(y) = \min_n \{h(x, y|\sigma)\}$. The minimum is well defined because \mathcal{X} is finite.

The final part of Proposition 1 is obtained from applying Theorem 5.30 of Villani (2009), reproduced as part 4 of Theorem A.1. The result is established under the conditions that both $(F(x, y))^{\frac{\sigma-1}{\sigma}}$ and the Monge-Kantorovich problem (B.4) have finite value and the *F*-subdifferential of v is single-valued *G*-almost everywhere.

Lemma B.3. If q is such that:

- *i.* $\sigma > 1$ or all workers can produce in all tasks: q(x, y) > 0 for all pairs $(x, y) \in \mathcal{X} \times \mathcal{Y}$.
- ii. $q(x, \cdot)$ is upper-semicontinuous in y given $x \in \mathcal{X}$.
- iii. q discriminates across workers in almost all tasks: if $q(x_n, y) = q(x_m, y)$ then $x_n = x_m$ G-a.e.

Then there exists $\lambda^* \in \mathbb{R}^N$ that solves the dual problem (B.5). Moreover, let T be defined as:

$$T(y) = \underset{x \in \mathcal{X}}{\operatorname{argmax}} \left\{ h(x, y | \sigma) - \lambda_{n(x)}^{\star} \right\},$$

where T is single-valued G-almost everywhere and it induces a deterministic coupling π^* : $\mathcal{X} \times \mathcal{B}(\mathcal{Y}) \rightarrow \mathbb{R}_+$ that is the unique (in law) solution to the Monge-Kantorovich problem (B.4). π^* is:

$$\pi^{\star}\left(x_{n},Y\right) = \int_{Y \cap T^{-1}\left(x_{n}\right)} dG.$$

Function T is an assignment function and it is the solution to the Monge transportation problem (7).

Proof. The proof follows from applying part 4 of Theorem A.1 (from Villani (2009, Thm. 5.30)). Finiteness of $h(x, y|\sigma)$ is guaranteed if $\sigma > 1$, or if $\sigma = 1$ and q(x, y) > 0 for all pairs $(x, y) \in \mathcal{X} \times \mathcal{Y}$. Finiteness of the value of the Monge-Kantorovich problem is guaranteed since \mathcal{Y} and \mathcal{X} are both compact, and q is upper semicontinuous on y.

It is left to verify that for any *h*-convex function $v : \mathcal{Y} \to \mathbb{R} \cup \{-\infty\}$ the *h*-subdifferential $\partial^h v(y)$ is single valued *G*-almost everywhere. The *h*- subdifferential for a given *y* is given by:

$$\partial^{h} v\left(y\right) = \left\{x \in \mathcal{X} \mid v^{h}\left(x\right) + v\left(y\right) = h\left(x, y \mid \sigma\right)\right\} \quad \text{where} \quad v^{h}\left(x\right) = \sup_{y} \left\{h\left(x, y \mid \sigma\right) - v\left(y\right)\right\}.$$

Because v is h-convex we can instead use its conjugate function $v^{h}(x_{n}) = \lambda_{n}$. Then the h-subdifferential is then equivalently given by:

$$\partial^{h} v(y) = \operatorname*{argmax}_{x \in \mathcal{X}} \left\{ h(x, y | \sigma) - \lambda_{n(x)} \right\}.$$

Becuase $q(\cdot, y)$ is injective in x given y G-a.e., and \mathcal{X} is finite, we get that $\partial^{h} v(y)$ is generically a singleton.

The following lemma establishes the relationship between the multipliers of the transformed problem (B.4) and multipliers of the original problem (B.2).

Lemma B.4. Consider two constrained maximization problems:

$$V(m) = \max_{x} \quad F(x) \qquad s.t. \ h(x) = m; \tag{B.6}$$

$$W(m) = \max_{x} g(F(x)) \qquad s.t. \ h(x) = m, \tag{B.7}$$

where $F : \mathcal{X} \to \mathbb{R}$, $h : \mathcal{X} \to \mathbb{R}^n$, $m \in \mathbb{R}^n$ and $g : \mathbb{R} \to \mathbb{R}$ is strictly monotone. Let $\lambda \in \mathbb{R}^n$ be the multiplier associated with the constraints in (B.6), and $\mu \in \mathbb{R}^n$ the multiplier associated with the constraints in (B.7). Then: $\mu = g'(F(x^*))\lambda$, where x^* is a solution for (B.6) and (B.7).

Proof. Because g is strictly monotone both problems have the same argmax, call it $x^{*}(m)$. The value of each problem is:

$$V(m) = F(x^{\star}(m)) \qquad W(m) = g(F(x^{\star}(m))).$$

By the envelope theorem (Milgrom and Segal, 2002) we know that:

$$\lambda = \frac{\partial V(m)}{\partial m} = \frac{\partial F(x^*)}{\partial x} \frac{\partial x^*(m)}{\partial m} \qquad \mu = \frac{\partial W(m)}{\partial m} = \frac{\partial g(F(x^*))}{\partial F} \frac{\partial F(x^*)}{\partial x} \frac{\partial x^*(m)}{\partial m}.$$

Joining gives the result: $\mu = g'(F(x^*))\lambda$.

B.2 Marginal product

The marginal product of a worker gives the change in output if more workers of that type are used in production. The change in output depends on the tasks that are assigned to additional workers.³⁹ Because of this, it is possible to define the marginal product at a given task, and under some initial assignment. In the main text, I consider the notion of equilibrium marginal products, where the assignment is not taken as given, but it is allowed to react optimally to changes in the supply of workers.

Consider the marginal product of a worker of type x_k at task \overline{y} , given an assignment T. Since task \overline{y} has no mass, output does not change if the task is re-assigned to a worker of type x_k . The marginal product is measured by adding a mass of workers of type x_k and assigning them to a region around task \overline{y} , replacing the workers previously assigned to those tasks. The marginal product at \overline{y} is obtained as the change in output when the mass of added workers tends to zero.

Proposition 4. [Marginal Product] Let T be a deterministic assignment and fix a task $\overline{y} \in \mathcal{Y}_n^{\circ}$. The marginal product of a worker of type x_k at task \overline{y} is:

$$MP(x_k, \overline{y}|T) = F(T) \left(\ln q(x_k, \overline{y}) - \ln q(x_n, \overline{y}) \right),$$

where $F(T) = \exp\left(\int \ln q\left(T\left(y\right), y\right) dG\right)$ and $T(\overline{y}) = x_n$.

When task \overline{y} is re-assigned from x_n to x_k output changes by $\ln q(x_k, \overline{y}) - \ln q(x_n, \overline{y})$. The marginal product takes into account the opportunity cost of assigning task \overline{y} to x_k , which comes from the capacity constraint of tasks. The derivative of output takes into account the scale of production at the current assignment. Task \overline{y} is required to be in the interior of \mathcal{Y}_n for technical reasons. If $\overline{y} \in \mathcal{Y}_n \cap \mathcal{Y}_m$ it becomes necessary to specify the region around \overline{y} to which x_k will be assigned.

The proof of the result is complicated because the task \overline{y} has dimension zero in the space of tasks, which has dimension $d \geq 1$. Before showing the general proof for the result, I consider the single-dimensional case where the argument is simpler. I further assume that $y \sim U([0, 1])$. When d = 1 the production function can be written as:

$$F(T) = \exp\left(\int_{0}^{1} \ln q\left(T\left(y\right), y\right) dy\right).$$

Fix a task $\overline{y} \in (0,1)$ and consider adding a mass ϵ of workers of type x_k . Workers are

 $^{^{39}}$ Unlike traditional production functions, the amount of an input used by the firm in production and what that input is used for are not the same.

assigned to the set $C_{\overline{y},\epsilon} = \left\{ y \mid |y - \overline{y}| < \frac{\epsilon}{2} \right\} = \left[\overline{y} - \frac{\epsilon}{2}, \overline{y} + \frac{\epsilon}{2} \right]$. The new assignment is:

$$T_{\epsilon}(y) = \begin{cases} T(y) & \text{if } y \notin C_{\overline{y},\epsilon} \\ 0 & \text{if } y \in C_{\overline{y},\epsilon} \land x \neq x_k \\ 1 & \text{if } y \in C_{\overline{y},\epsilon} \land x = x_k \end{cases}$$

The change in output is:

$$F(T_{\epsilon}) - F(T) = F(T) \left(\exp \left(\int_{\overline{y} - \frac{\epsilon}{2}}^{\overline{y} + \frac{\epsilon}{2}} \left(\ln q(x_k, y) - \ln q(T(y), y) \right) dy \right) - 1 \right)$$

The marginal product is:

$$\mathrm{MP}\left(x_{k}, \overline{y}|T\right) = \left.\frac{\partial F\left(R_{\epsilon}\right)}{\partial \epsilon}\right|_{\epsilon=0} = \lim_{\epsilon \to 0} \frac{F\left(R_{\epsilon}\right) - F\left(T\right)}{\epsilon},$$

replacing and applying L'HÃŽpital's rule:

$$\operatorname{MP}\left(x_{k}, \overline{y}|T\right) = F\left(R\right) \left. \frac{\partial \exp\left(\int_{\overline{y}-\frac{\epsilon}{2}}^{\overline{y}+\frac{\epsilon}{2}} \left(\ln q\left(x_{k}, y\right) - \ln q\left(T\left(y\right), y\right)\right) dy\right)}{\partial \epsilon} \right|_{\epsilon=0}$$

The derivative follows from Leibniz's rule. Generically $\overline{y}\in\mathcal{Y}_n^{\mathrm{o}}$ and:

$$MP(x_k, \overline{y}|T) = F(T) \left[\frac{1}{2} \left(\ln q \left(x_k, \overline{y} + \frac{\epsilon}{2} \right) - \ln q \left(T(y), \overline{y} + \frac{\epsilon}{2} \right) \right) + \frac{1}{2} \left(\ln q \left(x_k, \overline{y} - \frac{\epsilon}{2} \right) - \ln q \left(T(y), \overline{y} - \frac{\epsilon}{2} \right) \right) \right]_{\epsilon=0}$$
$$= F(T) \left(\ln q \left(x_k, \overline{y} \right) - \ln q \left(x_n, \overline{y} \right) \right).$$

If $\overline{y} \in \mathcal{Y}_n \cap \mathcal{Y}_m$ the marginal product takes into account that x_k replaces different types of workers around \overline{y} :

$$MP(x_k, \overline{y}|T) = F(T)\left(\ln q(x_k, \overline{y}) - \frac{\ln q(x_n, \overline{y}) + \ln q(x_m, \overline{y})}{2}\right).$$

In multiple dimensions, the treatment of the boundary cases becomes intractable, except in very specific cases for which similar expressions are obtained.

I now provide the general proof of the result.

Proof. Recall that the space of skills is of dimension d. Changing the assignment of tasks to workers in any region of dimension less than d will have no impact on output. To compute the effect on output

of the added workers it is necessary to proceed one dimension at a time. Consider a region formed as a hypercube around \overline{y} , with sides of length ϵ_i , denote this region by $C_{\overline{y},\epsilon} = \{y | \forall_i | y_i - \overline{y}_i | \leq \frac{\epsilon_i}{2}\}$. Note that as all $\epsilon_i \to 0$ the region $C_{\overline{y},\epsilon} \to \{\overline{y}\}$. The assignment is modified as in the single-dimensional example:

$$T_{\epsilon}(y) = \begin{cases} T(y) & \text{if } y \notin C_{\overline{y},\epsilon} \\ 0 & \text{if } y \in C_{\overline{y},\epsilon} \land x \neq x_k \\ 1 & \text{if } y \in C_{\overline{y},\epsilon} \land x = x_k \end{cases}$$

The difference in production between the two assignments is:

$$F(T_{\epsilon}) - F(T) = F(T) \left(\exp \left(\int_{\overline{y}_{1} - \frac{\epsilon_{1}}{2}}^{\overline{y}_{1} + \frac{\epsilon_{1}}{2}} \int_{\overline{y}_{d} - \frac{\epsilon_{d}}{2}}^{\overline{y}_{d} + \frac{\epsilon_{d}}{2}} (\ln q(x_{k}, y) - \ln q(T(y), y)) dy \right) - 1 \right)$$

I proceed by computing the change in output when the region $C_{\overline{y},\epsilon}$ changes. The change has to be computed one dimension at a time. If all dimensions are changed simultaneously the change in Fgoes to zero (this can be verified directly using Reynold's transport theorem- Theorem A.2). The change in output when $C_{\overline{y},\epsilon}$ changes in the d^{th} dimension is:

$$\frac{\partial F\left(T_{\epsilon}\right)}{\partial \epsilon_{d}} = F\left(T\right) \begin{pmatrix} \overline{y}_{1} + \frac{\epsilon_{1}}{2} & \overline{y}_{d-1} + \frac{\epsilon_{d}}{2} \\ \int & \int \\ \overline{y}_{1} - \frac{\epsilon_{1}}{2} & \cdots & \int \\ \overline{y}_{d-1} - \frac{\epsilon_{d}}{2} \end{pmatrix} \frac{1}{2} \left(\ln q\left(x_{k}, \overline{y}^{+}\left(d,\epsilon\right)\right) - \ln q\left(T\left(\overline{y}^{+}\left(d,\epsilon\right)\right), \overline{y}^{+}\left(d,\epsilon\right)\right)\right) \\ + \frac{1}{2} \left(\ln q\left(x_{k}, \overline{y}^{-}\left(d,\epsilon\right)\right) - \ln q\left(T\left(\overline{y}^{-}\left(d,\epsilon\right)\right), \overline{y}^{-}\left(d,\epsilon\right)\right)\right) dy_{1} \dots dy_{d-1}\right),$$

where $\overline{y}^+(d,\epsilon) \equiv (\overline{y}_1,\ldots,\overline{y}_{d-1},\overline{y}_d+\epsilon/2)^T$ and $\overline{y}^-(d,\epsilon) \equiv (\overline{y}_1,\ldots,\overline{y}_{d-1},\overline{y}_d-\epsilon/2)^T$. Applying the same procedure iteratively we obtain the change in output as x_k is assigned to

Applying the same procedure iteratively we obtain the change in output as x_k is assigned to tasks around \overline{y} in all directions:

$$\mathrm{MP}\left(x_{k},\overline{y}|T\right) = \left.\frac{\partial^{d}F\left(T_{\epsilon}\right)}{\partial\epsilon_{1}\cdots\partial\epsilon_{d}}\right|_{\epsilon=0} = F\left(T\right)\left(\ln q\left(x_{k},\overline{y}\right) - \ln q\left(x_{n},\overline{y}\right)\right).$$

B.3 Differentiability of demand

Lemma B.5. Let $\lambda \in \mathbb{R}^N$ be a vector of multipliers. If q satisfies Assumption (2) then D_n is continuously differentiable with respect to λ and:

$$\frac{\partial D_n}{\partial \lambda_m} = \frac{\operatorname{area}\left(\mathcal{Y}_n\left(w\right) \cap \mathcal{Y}_m\left(w\right)\right)}{2\sqrt{\left(x_n - x_m\right)'} A' A\left(x_n - x_m\right)} \ge 0.$$

Proof. The proof follows from an application of Reynolds' Transport Theorem (Theorem A.2). In order to apply Reynolds' theorem recall that $D_m = \int_{\mathcal{Y}_m} \rho(y) \, dy$, where ρ is the density of tasks in the space. In our case $\rho(y) = 1$. So the volume is \mathcal{Y}_m and the function is the density of tasks.

The second term in the theorem measures the rate at which the density flows in and out of the volume. The density flows out and into other workers as tasks are reassigned. Consider the flow into of \mathcal{Y}_m and out of \mathcal{Y}_k . The flow is in the direction $\frac{A(x_k-x_m)}{\sqrt{(x_k-x_m)'A'A(x_k-x_m)}}$ and through the shared boundary of \mathcal{Y}_m and \mathcal{Y}_k , given by $\mathcal{Y}_m \cap \mathcal{Y}_k$. Note that when prices change the hyperplanes that define the boundaries of the demand sets move in parallel. Applying the theorem:

$$\frac{\partial D_m}{\partial w_n} = \int_{\mathcal{Y}_m} \frac{\partial \rho\left(y\right)}{\partial w_n} dy + \sum_{k \neq m} \int_{\mathcal{Y}_m \cap \mathcal{Y}_k} \rho\left(y\right) \left(\frac{\frac{\partial y \cdot \frac{A(x_k - x_m)}{\sqrt{(x_k - x_m)'A'A(x_k - x_m)}}}{\partial w_n}\right) dy$$

Note that for all $y \in \mathcal{Y}_m \cap \mathcal{Y}_k$ lie in a plane perpendicular to $A(x_k - x_m)$. Then they can be always expressed as $y = y_\lambda + a\vec{v}$ where $a \in \mathbb{R}$, \vec{v} is a vector perpendicular to $A(x_k - x_m)$ and $y_\lambda = (1 - \lambda) x_k + \lambda x_m$ is such that $y_\lambda \in \mathcal{Y}_m \cap \mathcal{Y}_k$. Then the change $y \in \mathcal{Y}_m \cap \mathcal{Y}_k$ is equal to the change in y_λ .

$$\frac{\partial D_m}{\partial w_n} = \sum_{k \neq m} \int_{\mathcal{Y}_m \cap \mathcal{Y}_k} \rho\left(y\right) \left(\frac{\frac{\partial y_\lambda \cdot \frac{A(x_k - x_m)}{\sqrt{(x_k - x_m)'A'A(x_k - x_m)}}}{\partial w_n}\right) dy$$

The value of λ is obtained from the equation for the hyperplane that defines $\mathcal{Y}_m \cap \mathcal{Y}_k$:

$$\lambda = \frac{(x_m - x_k)' A (x_m - x_k) + w_m - w_k}{2 (x_m - x_k)' A (x_m - x_k)},$$

so:

$$\frac{\partial y_{\lambda} \cdot \frac{x_k - x_m}{\sqrt{(x_k - x_m)'(x_k - x_m)}}}{\partial w_n} = \begin{cases} \frac{1}{2\sqrt{(x_n - x_m)'A'A(x_n - x_m)}} & \text{if } k = n\\ 0 & \text{otw} \end{cases}$$

Replacing:

$$\frac{\partial D_m}{\partial w_n} = \frac{\int_{\mathcal{Y}_m \cap \mathcal{Y}_n} \rho\left(y\right) dy}{2\sqrt{(x_n - x_m)'(x_n - x_m)}} = \frac{\operatorname{area}\left(\mathcal{Y}_n\left(w\right) \cap \mathcal{Y}_m\left(w\right)\right)}{2\sqrt{(x_n - x_m)'(x_n - x_m)}},$$

which completes the proof.

B.4 Directed automation

Proposition 5. Consider the automation problem in 24 and let $\mu \in \mathbb{R}^{N+1}$ characterize an assignment according to 27. If q is differentiable then the first order conditions of the problem are:

$$F_{R}(\mu, r) \int_{\mathcal{Y}_{R}} \frac{\partial \ln q(r, y)}{\partial r} dy - \frac{\partial \Omega(r, p_{r})}{\partial r} = 0; \qquad [r]$$

$$F_R(\mu, r) \mu_R - \frac{\partial \Omega(r, p_r)}{\partial p_r} = 0. \qquad [p_r]$$

Proof. After replacing T_R for μ in the problem, and abusing notation, the corresponding Lagrangian is:

$$\max_{\{r, p_r, \mu, \Lambda\}} \mathcal{L} = F_R(\mu, r) - \Omega(r, p_r) + \sum_{n=1}^N \Lambda_n (p_n - D_n) + \Lambda_R (p_r - D_R).$$
(B.8)

The multipliers of the workers/robot capacity constraints are given by the vector $\Lambda \in \mathbb{R}^{N+1}$.

The first order condition of interest is with respect to the skills of the robot:

$$\frac{\partial \mathcal{L}}{\partial r} = \frac{\partial F_R(\mu, r)}{\partial r} - \frac{\partial \Omega(r, p_r)}{\partial r} - \sum_{n=1}^N \Lambda_n \frac{\partial D_n}{\partial r} - \Lambda_R \frac{\partial D_R}{\partial r}.$$
(B.9)

Following de Goes et al. (2012) and using the result in Lemma B.4 the first order condition becomes:

$$\frac{\partial \mathcal{L}}{\partial r} = \frac{\partial F_R(\mu, r)}{\partial r} - \frac{\partial \Omega(r, p_r)}{\partial r} - F_R(\mu, r) \left(\sum_{n=1}^N \mu_n \frac{\partial D_n}{\partial r} - \mu_R \frac{\partial D_R}{\partial r} \right).$$
(B.10)

I proceed by computing separately the first term of the first order condition:

$$\frac{\partial F_R(\mu, r)}{\partial r} = F_R(\mu, r) \left(\sum_{n=1}^N \frac{\partial \int_{\mathcal{Y}_n} \ln q(x_n, y) \, dy}{\partial r} + \frac{\partial \int_{\mathcal{Y}_R} \ln q(r, y) \, dy}{\partial r} \right)$$

Each of the derivatives follows from Reynold's theorem.

$$\begin{aligned} \frac{\partial \int_{\mathcal{Y}_n} \ln q\left(x_n, y\right) dy}{\partial r} &= \int_{\mathcal{Y}_n} \frac{\partial \ln q\left(x_n, y\right)}{\partial r} dy + \int_{\mathcal{Y}_n \cap \mathcal{Y}_R} \ln q\left(x_n, y\right) \frac{\partial y \cdot c_{nr}}{\partial r} dy \\ &= \int_{\mathcal{Y}_n \cap \mathcal{Y}_R} \ln q\left(x_n, y\right) \frac{\partial y \cdot c_{nr}}{\partial r} dy, \end{aligned}$$

where $c_{nr} = \frac{2A(x_n-r)}{\sqrt{(x_n-r)'A(x_n-r)'}}$ is the normal vector to the direction in which the boundary is moving.

In a similar way:

$$\frac{\partial \int_{\mathcal{Y}_R} \ln q\left(r, y\right) dy}{\partial r} = \int_{\mathcal{Y}_R} \frac{\partial \ln q\left(r, y\right)}{\partial r} dy + \sum_{n=1}^N \int_{\mathcal{Y}_n \cap \mathcal{Y}_R} \ln q\left(x_n, y\right) \frac{\partial y \cdot c_{rn}}{\partial r} dy,$$

where $c_{rn} = -c_{nr}$. Joining and reorganizing we get:

$$\frac{1}{F_R(\mu,r)}\frac{\partial F_R(\mu,r)}{\partial r} = \int_{\mathcal{Y}_R} \frac{\partial \ln q(r,y)}{\partial r} dy + \sum_{n=1}^N \int_{\mathcal{Y}_n \cap \mathcal{Y}_R} \left(\ln q(x_n,y) - \ln q(r,y)\right) \frac{\partial y \cdot c_{nr}}{\partial r} dy.$$

Note now that by the definition of the boundary $\ln q(x_n, y) - \ln q(r, y) = \mu_n - \mu_r$ for all $y \in \mathcal{Y}_n \cap \mathcal{Y}_R$. Then:

$$\frac{1}{F_R(\mu,r)}\frac{\partial F_R(\mu,r)}{\partial r} = \int_{\mathcal{Y}_R} \frac{\partial \ln q(r,y)}{\partial r} dy + \sum_{n=1}^N (\mu_n - \mu_r) \int_{\mathcal{Y}_n \cap \mathcal{Y}_R} \frac{\partial y \cdot c_{nr}}{\partial r} dy$$

Finally note that $\frac{\partial D_n}{\partial r} = \int_{\mathcal{Y}_n \cap \mathcal{Y}_R} \frac{\partial y \cdot c_{nr}}{\partial r} dy$, which follows from applying Reynold's Theorem (again) to D_n .

$$\frac{1}{F_R(\mu,r)}\frac{\partial F_R(\mu,r)}{\partial r} = \int_{\mathcal{Y}_R} \frac{\partial \ln q(r,y)}{\partial r} dy + \sum_{n=1}^N \left(\mu_n - \mu_r\right) \frac{\partial D_n}{\partial r}.$$

When the location of the robot (r) is changed, there is a change in output due to the change in mismatch inside the region previously assigned to the robot (\mathcal{Y}_R) , that is given by the first term. There is also a change in the demand for workers, only workers who are neighbors of the robot are affected. When their demand is affected, the demand of the robot changes in the opposite direction. The demand for worker *n* changes by $\frac{\partial D_n}{\partial r}$, that is valued by the planner at $\lambda_n - \lambda_r$. Recall that λ_n is the shadow price of the supply of a worker.

It is left to spell out the first term:

$$\int_{\mathcal{Y}_R} \frac{\partial \ln q(r, y)}{\partial r} dy = \int_{\mathcal{Y}_R} \frac{\partial a'_x r - (r - y)' A(r - y)}{\partial r} dy$$
$$= \int_{\mathcal{Y}_R} (a_x - 2Ar + 2Ay) dy$$
$$= 2D_R \left(\frac{a_x}{2} - A(r - b_R)\right),$$

where $b_R = \frac{\int_{\mathcal{Y}_R} y dy}{D_R}$ is the barycenter (centroid, average or center of mass) of the tasks assigned to r.

It is now possible to obtain the first order condition of the problem with respect to the location of the robot. Note that since the total demand is constant it holds that:

$$\frac{\partial D_R}{\partial r} = -\sum_{n=1}^N \frac{\partial D_n}{\partial r}.$$

Then:

$$\frac{\partial \mathcal{L}}{\partial r} = \frac{\partial F_R(\mu, r)}{\partial r} - \frac{\partial \Omega(r, p_r)}{\partial r} - F_R(\mu, r) \left(\sum_{n=1}^N (\mu_n - \mu_R) \frac{\partial D_n}{\partial r} \right).$$
(B.11)

Replacing for $\frac{\partial F_R(\mu, r)}{\partial r}$ we get:

$$\frac{\partial \mathcal{L}}{\partial r} = 2F_R(\mu, r) D_R\left(\frac{a_x}{2} - A(r - b_R)\right) - \frac{\partial \Omega(r, p_r)}{\partial r}.$$
(B.12)

The first order condition does not include the effect of r on the demand for workers since the gains cancel with the reductions/increases of slack in the feasibility constraints.

This is a necessary condition for an optimum. It does not fully characterize the solution. In fact, there can be, in general, multiple solutions to the problem. The first order condition is also silent about the location of the region assigned to r. Instead, it prescribes the relationship between the region's centroid and the location of r. It is convenient to see what happens when $a_x = 0$ and $\frac{\partial \Omega(r, p_r)}{\partial r} = 0$. Then the necessary condition reduces to make r equal to the barycenter of its region. The first order condition with respect to p_r is:

$$\frac{\partial F}{\partial p_{r}} = F_{R}\left(\mu, r\right) \mu_{R} - \frac{\partial \Omega\left(r, p_{r}\right)}{\partial p_{r}}.$$

The first order condition with respect to μ requires more work, but it follows from applying again Reynolds' Transport Theorem.

$$\frac{\partial F}{\partial \mu_n} = p_n - D_n.$$

B.5 Worker training

I introduce the problem of optimal worker training in a similar way to the automation problem described above. The objective in both cases is to reduce the mismatch between tasks and workers, now by modifying workers' skills. Crucially, as the skills of the worker change the assignment also changes, altering the tasks in the workers' occupation.

Formally, the problem of training worker n by choosing new skills $\tilde{x} \in \mathcal{S}$ is:

$$\max_{\{\tilde{x},\tilde{T}\}} F\left(\tilde{T},\tilde{x}\right) - \Gamma\left(\tilde{x}|x_n, p_n\right) \qquad \text{s.t. } \forall_{\ell} D_{\ell} \le p_{\ell}, \tag{B.13}$$

where the cost of changing skills (Γ) depends on the workers' current skills and mass. The first order condition of the problem is:

$$F\left(\tilde{T}^{\star}\right)\int_{\mathcal{Y}_{n}}\frac{\partial\ln q\left(\tilde{x},y\right)}{\partial\tilde{x}}dG\left(y\right)-\frac{\partial\Gamma\left(\tilde{x}|x_{n},p_{n}\right)}{\partial\tilde{x}}=0_{d\times1}.$$
(B.14)

The interpretation is the same as in the automation problem. The objective is to minimize skill mismatch across the tasks in the worker's occupation given the cost of changing the workers' skills. If q is given by (1) this is achieved by setting \tilde{x} to the centroid of the occupation, and adjusting for the weight of skills in production (a_x) and the marginal cost of changing the worker's skills. Even if acquiring skills was costless, it is not always optimal to increase the worker's skills, doing so can generate its own costs as mismatch increases with respect to the boundary tasks of the worker's occupation. The problem is further complicated by the ambiguous effects on total output, since training one worker can induce higher mismatch for other workers, as the assignment changes. Because of this, condition (B.14) is only necessary, and not sufficient, for characterizing the optimal worker training. However, Lloyd's algorithm still applies.

C Multi-Sector Assignment

There is a tight link between the model presented in Section 1 and a multi-sector assignment problem. Consider an alternative interpretation of the setup where the roles of tasks and workers are reversed, so that there is a continuum of workers and finitely many tasks (or sectors). Under this interpretation it is convenient to think of each sector as having p_n identical firms, each can hire a single worker. The problem is for each sector to select which workers to hire, or for each worker to select a sector as in the traditional Roy (1951) model.⁴⁰ A worker $y \in \mathcal{Y}$ in sector $x_n \in \mathcal{X}$ produces an output of $q(x_n, y)$.

Each sector differs in how it ranks potential workers. Differences in the skills of workers make them better suited for certain sectors over others. The rank of workers in each sector reflects the productivity of the worker as captured by q:

$$r_n(y) \equiv \Pr\left(q\left(x_n, Y\right) \le q\left(x_n, y\right)\right),\tag{C.1}$$

where the probability is taken with respect to the distribution of workers G. The functions $r_n : \mathcal{Y} \to [0, 1]$ provide a single -dimensional value (the rank) for each worker in each sector. The rank is 'sector' specific as different sectors value workers differently. The conditions in Proposition 1 ensure that r_n is continuous.

We can further make use of it by defining a surplus function:

$$Q_n(r_n(y)) \equiv q(x_n, y). \tag{C.2}$$

The surplus function $Q_n : [0, 1] \to \mathbb{R}_+$ gives the output of a worker with rank r in sector n. It follows from the definition of the rank in (C.1) that Q_n is increasing. The definition of the rank functions r_n and the surplus function Q follows Gola (2021).

This alternative setup can be solved directly in terms of the multi-dimensional distribution of workers and sectors and the production function q as in Section 2 or in terms of the sector-specific (single-dimensional) rankings and the surplus function Q_n . The second approach is used by Gola (2021), who exploits the fact that the joint distribution of rankings across sectors is a Copula (Sklar, 1959). This is the joint distribution over tuples $(r_1(y), \ldots, r_N(y))$ for $y \in \mathcal{Y}$. In both cases the solution is a tuple of profits $(\lambda_1, \ldots, \lambda_N)$ such that firms in sector n have a profit of λ_n . These profits are supported by wage functions $\nu_n(r) \equiv Q_n(r) - \lambda_n$ that give the wage of a worker in sector n with rank r. These prices ensure that markets clear when workers choose sectors by maximizing their wage.

While the setup of Section 1 makes it possible to work directly with the primitives of the model without relying on differentiability, and to tackle the bundling of tasks and changes in occupations, the setup in this appendix provides additional intuition on the role of skill heterogeneity across workers. As in Gola (2021), the concordance of the joint distribution of rankings is informative about the effective supply of skills in the economy.⁴¹ If concordance is high, the same workers are ranked high/low across sectors and the effective supply of skills

⁴⁰This is equivalent to an economy where we maintain the interpretation of workers and firms as in Section 1 but we solve the assignment from the perspective of workers that choose which tasks to perform or equivalently rent machines (tasks) from firms (capitalists).

⁴¹A copula, C_1 , is more concordant than C_2 , if $C_1(\vec{r}) \ge C_2(\vec{r})$ for all $\vec{r} \in [0, 1]^N$ (Scarsini, 1984).

is low because various sectors (x_n) are after the same (type of) workers (y). See Gola (2021) for the implications of changes in concordance on sectoral wages.

In terms of the model used in the main text, a high concordance means that workers (x_n) are good/bad at performing the same tasks. An increase in concordance (either because of a change in q or a change in the distribution or workers or tasks) makes it so that workers are more substitutable as they are more alike in the performance of tasks. Wage differentials would decrease because of the same reason. To be precise, an increase in the concordance of the joint distribution of tasks ranks would be reflected in a reduction of the productivity differentials of workers at boundary tasks.

Higher concordance also reduces total output because it reflects the fact that workers are good at performing the same set of tasks. Because these tasks are in fixed demand (G), most workers end up being assigned to tasks in which they are not productive. The flip-side of this is a change in technology or worker training that reduces concordance. Such a change reflects workers becoming more specialized in performing different types of tasks, with the workers that are good at some tasks being bad in others. A decrease in concordance increases output by allowing for more specialization in production.

D Cosine Similarity Assignments

I now consider the optimal assignment of tasks to workers when the mismatch is measured by the cosine similarity between the skill vectors of workers and tasks respectively. The cosine similarity between two vectors is given by the cosine of the angle between the vectors:

$$\cos(\theta_{xy}) = \frac{x'y}{\|x\| \|y\|},$$
 (D.1)

where $||z|| = \sqrt{\sum_{i=1}^{d} z_i^2}$ is the (Euclidean-)norm of a vector. The cosine similarity varies between 0 and 1. The main advantage of the cosine similarity is that if effectively reduces the dimensionality of the problem from d to a single dimension. This is because only the angle between two vectors, and not their magnitudes, matters for this notion of mismatch.

The set up of the model is just as in Section 1 with just two modifications. First, the set of tasks can be restricted (without loss) to tasks on a circle in the positive quadrant (or the surface of a sphere in the positive orthant for higher dimensions). Second, the task output function is:

$$q(x,y) = \exp\left(f^x(x) + \cos\left(\theta_{xy}\right)\right),\tag{D.2}$$

where $f^x(\cdot)$ captures the role of workers' skills in production. Figure D.1a presents the setup. Each worker x_1 , x_2 and x_3 is paired with an angle θ_1^x , θ_2^x , θ_3^x . The cosine similarity between a worker and a task in is the cosine of the difference between their angles.



Figure D.1: Assignment Example - Cosine Similarity Loss

Note: The left panel depicts the setup of the model when mismatch is measured according to the cosine similarity between workers and tasks as in (D.2). Because the cosine similarity is independent of magnitudes only the tasks in along gray circumference are considered. There are three types of workers $\{x_1, x_2, x_3\}$ with mass $P = \{0.4, 0.3, 0.3\}$ at angles $\{\theta_1^x, \theta_2^x, \theta_3^x\}$. The right panel shows the optimal assignment. The assignment is characterized by two angles, θ_{12}^y and θ_{13}^y , which determine the boundaries of between \mathcal{Y}_1 and \mathcal{Y}_2 , and \mathcal{Y}_1 and \mathcal{Y}_3 , respectively. These angles partition the space as shown by the two rays that intersect the circumference.

Figure D.1b shows the optimal assignment. In this setup, the assignment consist on two cutoff angles $(\theta_{12}^y \text{ and } \theta_{12}^y)$ which characterize rays partitioning the space of skills. The solution is the same regardless of whether tasks are collapsed onto the circle or they are distributed on the whole plane. The magnitude of the tasks' vectors does not change the cosine distance from the task to the worker and therefore does not affect mismatch. This is undoubtedly a strong assumption, it implies that the magnitude of the skill vectors only affects production by determining an absolute productive of workers which is independent of the assignment. This productivity is captured by the value of $f^x(x)$.

The assignment in Figure D.1b also reveals the logic behind single-dimensional assignments. The productivity of workers as measured by $f^x(x_1)$, $f^x(x_2)$, and $f^x(x_3)$ establishes an ordering of workers prior to the assignment. It is optimal to assign tasks to the most productive worker first to ensure that all of its time is used in production. In the example this is worker x_3 . In fact, the value of θ_{13}^y in the assignment is such that a fraction p_3 of the tasks in the circle are assigned to worker x_3 . Following this process, it is optimal to assign tasks to worker x_2 next. This gives angle θ_{12}^y such that the tasks between it and $\pi/2$ are assigned to worker x_2 . The remaining tasks are assigned to worker x_1 , who is the least productive.