The Unification of Labor

Kweku Opoku-Agyemang*

August 2026

Abstract

The division of labor is a foundational organizing principle of economic production. This paper develops a theory in which the *breadth* of a productive unit—the set of tasks it can perform effectively—is endogenous. Specialization emerges when task-specific comparative advantage is sufficiently strong, while broader productive units become optimal when technology increases the productivity of performing multiple tasks within the same unit. We characterize the equilibrium transition between these regimes and show that it depends on the interaction between task-specific productivity, general-purpose capability, and the costs of coordinating production across units. The framework therefore distinguishes technological improvements that raise productivity within existing tasks from those that expand the set of tasks a productive unit can perform effectively. The latter can change not only the productivity of production but its optimal organization: tasks previously allocated across specialized units may become optimally consolidated within broader units. We refer to this process as *generalization* and its limiting organizational configuration as *unification of labor*. Artificial intelligence provides a natural motivation for this mechanism because a sufficiently general-purpose system can perform heterogeneous tasks without the task-specific specialization that characterizes conventional production. The results provide a unified framework in which the division of labor and its partial unification arise as alternative equilibrium responses to the technology of production.

Keywords: Division of labor; specialization; generalization; productive breadth; organization of production; technological change; general-purpose technology; artificial intelligence.
JEL Codes: D23; D24; J24; O33.

*Development Economics X Corporation. Email: kweku@developmenteconomicsx.com. . The conclusions of this research do not necessarily reflect the opinions or official position of Development Economics X Corporation.

1 Introduction

The division of labor is among the foundational ideas of economics. In Smith's (1776) canonical account, specialization raises productivity by allowing productive agents to concentrate on particular activities and exchange with one another. Subsequent theories of comparative advantage, human capital, the firm, and production networks have developed different aspects of the same principle: when productive capabilities differ sufficiently across activities, production gains from allocating different tasks to different agents.

This paper asks what happens when productive capabilities become sufficiently broad that this premise no longer holds.

We develop a theory of production in which the *breadth* of a productive unit—the set of tasks it can perform effectively—is endogenous. When capabilities are narrow and comparative advantage is strong, production is divided across specialized units. When technology increases the set of tasks that a single unit can perform at high productivity, however, the optimal allocation can move in the opposite direction: tasks previously assigned to separate specialists can become optimally assigned to the same unit. We call this process *generalization*, and refer to the resulting organizational configuration as *unification of labor*.

The distinction is between improvements in task productivity and increases in productive breadth. An improvement in productivity within a given task changes the return to performing that task but need not change the organization of production. An increase in productive breadth changes the set of task combinations that a single productive unit can perform effectively, and therefore changes the feasible partition of production itself. The central object of the paper is this endogenous partition.

The mechanism reflects a familiar tradeoff. Dividing production across agents allows each agent to exploit task-specific comparative advantage, but creates coordination costs. Separate productive units must communicate, transmit information, synchronize decisions, and reconcile activities across production stages. A broader productive unit may sacrifice some task-specific productivity, but can economize on these costs by internalizing multiple activities. The optimal organization therefore depends on the relative strength of specialization gains, productive breadth, and coordination costs.

This yields a unified theory of specialization and generalization. When task-specific comparative advantage is sufficiently strong and productive breadth is limited, specialization is optimal. As general-purpose capability increases, the marginal value of additional breadth rises and the optimal task portfolio of each productive unit expands. Beyond a critical level of general-purpose capability, the economy can transition to partial or complete unification. The Smithian division of labor is therefore not imposed as a primitive feature of production; it is one region of the equilibrium of a more general production environment.

Artificial intelligence provides a particularly important motivation for this question. Much of the economics of technological change analyzes technologies that improve productivity in particular tasks or alter the allocation of tasks between labor and capital. General-purpose artificial intelligence can operate on a different margin. The same underlying system can perform a large and potentially expanding set of heterogeneous cognitive tasks. If sufficiently capable artificial agents can perform diverse tasks near the productivity frontier for those tasks, technological progress does not merely make an existing input more productive. It expands the breadth of activities that a single productive unit can perform.

The limiting case illustrates the distinction. Suppose an artificial agent can perform each task required for production at approximately the productivity of the best specialized agent for that task. The standard allocation problem of assigning different tasks to different agents becomes less important because the same productive unit can span the task sequence. Production can then be reorganized around broader productive units even without a technological improvement specific to any individual task. What changes is the optimal *partition of production*.

The model makes this intuition precise. We characterize equilibrium as a function of three forces: heterogeneity in task-specific productivity, the technological breadth of productive units, and the costs of coordinating production across units. The central result is a threshold characterization. When task-specific comparative advantage is sufficiently strong relative to the productivity of generalists, specialization is optimal. As general-purpose capability increases, optimal productive breadth increases. Once the productivity advantage of breadth exceeds the relevant specialization and coordination costs, production becomes partially or fully unified.

The framework also generates implications beyond the allocation of tasks. First, technological progress can increase output while reducing the number of productive units required to execute a given production process. Second, changes in productive breadth can alter the effective boundaries of organizations: activities that previously required coordination across specialized units may become internal to a broader productive unit. Third, the value of specialization itself becomes endogenous to the technological environment. As general-purpose capability expands, some sources of task-specific comparative advantage become less valuable because a single unit can perform a larger set of activities without substantial productivity losses.

These results suggest a broader way to think about technological change. Mechanization, electrification, and information technology changed the productivity of particular inputs and the costs of coordinating production. General-purpose artificial intelligence may also change a more fundamental characteristic of production: the number and variety of tasks that a single productive unit can perform. The economic consequence is therefore potentially organizational as well as technological.

The paper distinguishes two limiting regimes. In a *specialized economy*, production is organized by allocating heterogeneous tasks across units with different comparative advantages. In a *generalized economy*, productive units span a broader collection of tasks, reducing the need to divide production across separate specialists. These regimes need not exhaust the possibilities: the model also permits coexistence and partial generalization. The central endogenous variable is the equilibrium degree of task concentration within productive units.

Our contribution is therefore to make productive breadth a central endogenous object in the theory of production. Existing theories explain the gains from specialization, the coordination costs that limit it, the reallocation of tasks induced by technological change, and the diffusion of general-purpose technologies. We ask how technological change alters the optimal number of tasks assigned to the same productive unit. This additional margin nests specialization and generalization within a common equilibrium framework.

The paper thus provides a theory of when the division of labor is optimal, when broader productive units are optimal, and how technological change can move an economy between the two. The resulting framework places the economics of artificial intelligence within a more general theory of the organization of production: sufficiently general-purpose technology can change not only who performs a task, but the number of tasks that it is efficient for the same productive unit to perform. Although artificial intelligence provides the motivating technological example, the theory does not depend on AI. Any technology that sufficiently expands the set of tasks that a productive unit can perform can generate the same organizational transition.

The remainder of the paper proceeds as follows after reviewing the literature. Section II introduces the production environment and defines task breadth, specialization, and generalization. Section III characterizes equilibrium and derives the conditions under which specialization and unification arise. Section IV studies comparative statics and the transition between the two regimes. Section V considers extensions involving heterogeneous agents, endogenous coordination costs, and partial generalization. Section VI discusses implications for firms, labor markets, and technological change. Section VII concludes.

1.1 Related Literature

This paper connects four literatures: the economics of specialization, task-based approaches to technological change, general-purpose technologies, and the organization of production. The common element across these literatures is the allocation of productive activities. The distinction developed here is that the breadth of the productive unit itself is endogenous. Specialization and generalization are therefore alternative equilibrium configurations of the same production problem.

1.2 Specialization and the Organization of Production

The division of labor occupies a foundational position in economic theory. Smith (1776) emphasized that specialization can raise productivity by allowing productive agents to concentrate on narrower activities, facilitating learning, dexterity, and the development of specialized methods. Subsequent theory has formalized the determinants of the extent of specialization.

Becker and Murphy (1992) provide a particularly important benchmark. They show that the gains from specialization must be weighed against the costs of coordinating increasingly specialized productive activities, with the extent of specialization depending on knowledge and coordination costs. (4) Their framework therefore treats the division of labor as an equilibrium outcome rather than a primitive feature of production.

The present paper takes this insight as a starting point but studies a different margin. In Becker and Murphy, the central question is why specialization does not become arbitrarily extensive when narrower specialization raises productivity. Here, the productivity of a productive unit is itself allowed to depend on the breadth of the task set it performs. The economy therefore chooses not only which productive unit performs each task, but also how many tasks optimally belong to the same unit.

This distinction permits specialization and generalization to emerge as alternative equilibria. Specialization is optimal when the productivity gains from concentrating on narrower task sets exceed the costs of maintaining multiple productive units. Generalization becomes optimal when technological progress sufficiently raises the value of productive breadth.

1.3 Tasks, Automation, and Technological Change

A second literature takes tasks rather than occupations or industries as the basic unit of analysis. Autor, Levy, and Murnane (2003) show how computerization changes the composition of work by differentially affecting routine and nonroutine tasks. (3) Subsequent work develops task-based models in which technological change reallocates tasks between labor and capital. Acemoglu and Restrepo (2018, 2019), for example, emphasize the displacement of labor from existing tasks together with the creation of new tasks that can generate a reinstatement effect. (1; 2) Grossman and Rossi-Hansberg (2008) similarly formulate international production in terms of tradeable tasks. (9)

The present paper shares this task-based perspective but introduces a different technological margin. Existing task-based frameworks primarily ask which factor should perform a given task. The question here is instead how many tasks should be performed by the same productive unit. Technological progress can therefore alter not only the allocation of tasks across factors but the optimal partition of the task space itself.

Recent work highlights that automation can strengthen rather than weaken specialization. Gong and Png (2024) show that automation can enable greater specialization by allowing machines to perform particular tasks while reducing the coordination costs associated with dividing those tasks among human workers. (7) Our mechanism is complementary but distinct. We study the opposite organizational margin: technological progress can increase the productivity of assigning a broader collection of tasks to the same productive unit. The relevant technological question is therefore not only which tasks technology performs, but whether technology changes the optimal breadth of the unit performing them.

1.4 General-Purpose Technologies

The literature on general-purpose technologies provides a natural technological counterpart to this framework. Bresnahan and Trajtenberg (1995) characterize general-purpose technologies by their pervasive use, scope for continued improvement, and complementarities with innovation in downstream applications. (5)

The distinction between technological generality and productive breadth is central here. A technology can be general purpose because it is used across many sectors while individual productive units remain highly specialized. Conversely, a sufficiently capable technology may make it efficient for a single productive unit to perform a much broader set of tasks.

Thus,

$$\text{technological generality} \neq \text{productive breadth}. \tag{1}$$

The former concerns the range of applications of a technology; the latter concerns the range of tasks optimally assigned to the same productive unit. The model developed below studies when increases in the former induce increases in the latter.

1.5 Organization, Comparative Advantage, and Productive Breadth

The framework also relates to the literature on comparative advantage and the boundaries of organizations. In conventional models, specialization reflects differences in relative productivity across agents or locations. The present framework does not overturn this logic; instead, it endogenizes the dimensionality over which comparative advantage operates. Technological progress can make the comparative advantage of a sufficiently general productive unit span a larger portion of the task space.

The organization-of-production literature similarly emphasizes the costs of coordinating activities within common organizational structures. Coase (1937), Williamson (1975, 1985), Grossman and Hart (1986), and

Hart and Moore (1990) provide foundational approaches to the determination of organizational boundaries. The productive unit in the present model need not coincide with a legal firm: it may be a worker, team, firm, or other organizational form. What matters is the set of tasks over which the benefits of breadth and the costs of coordination are jointly realized.

This distinction also separates the present framework from theories of human-capital specialization. Those theories emphasize the returns to investing in task-specific capabilities. Here, technological change can alter the relative return to breadth itself. A productive agent may consequently become more valuable not by becoming more specialized, but by acquiring capabilities that span a wider portion of the task space.

1.6 Contribution

The literature therefore provides theories of the productivity gains from specialization, the limits imposed by coordination costs, the reallocation of tasks induced by technological change, and the diffusion of general-purpose technologies. This paper adds a distinct organizational margin:

$$\boxed{\text{How does technological change alter the optimal breadth of a productive unit?}} \tag{2}$$

The distinction is important because a change in technology can alter the structure of production even when the set of tasks itself is unchanged. When the productivity advantage of breadth is sufficiently large relative to the costs of coordination and task-specific specialization, tasks that were previously assigned to separate productive units become optimally consolidated.

The model consequently nests the division of labor and its opposite within a common theory. Specialization and generalization are not different technologies or assumptions about the economy; they are alternative equilibrium allocations generated by the same production environment. In the limiting case of sufficiently powerful general-purpose capability, a productive unit can optimally span a large portion of the task space. We refer to this limiting configuration as *unification of labor*.

The contribution is therefore not a prediction that specialization necessarily disappears. It is a theory of when specialization is optimal, when generalization is optimal, and how technological change can induce a transition between the two.

2 The Production Environment

We begin with a production environment in which output requires the completion of a continuum of tasks. The central object of the model is the set of tasks that an individual agent can perform productively.

This allows us to distinguish two dimensions of productive capability that are often conflated: productivity conditional on performing a task, and the breadth of tasks that an agent can perform. The first determines comparative advantage; the second determines the scope for specialization.

2.1 Tasks and Agents

There is a continuum of tasks indexed by

$$j \in [0, 1].$$

A unit of the final good requires one unit of each task. Let q_j denote the quantity of task j produced. Final output is given by the Leontief technology

$$Y = \min_{j \in [0,1]} q_j.$$

The Leontief specification is not essential for the main mechanism, but it makes the organizational margin transparent: production requires the coordinated completion of a set of heterogeneous activities.

There is a unit mass of potentially productive agents indexed by i. An agent may allocate effort across the tasks in which it is capable of producing. Let

$$a_{ij} \geq 0$$

denote agent i's productivity in task j, measured as task output per unit of effort. If agent i devotes effort e_{ij} to task j, its contribution to that task is

$$q_{ij} = a_{ij} e_{ij}.$$

The total effort of an agent is bounded by

$$\int_0^1 e_{ij} \, dj \leq 1.$$

The matrix of productivities $\{a_{ij}\}$ constitutes the primitive distribution of productive capabilities. We separate two features of this distribution. The first is *task-specific productivity*: conditional on performing task j, how efficiently does agent i perform it? The second is *task breadth*: over how many tasks can agent i attain a sufficiently high level of productivity?

This distinction is central to the analysis. A technology that raises a_{ij} for a fixed j makes an agent more productive at an existing task. A technology that raises the number of j for which a_{ij} is high expands the agent's productive breadth. The former strengthens production within a given division of labor; the latter can change the optimal division of labor itself.

2.2 Capability Sets and Task Breadth

For each agent i, define its capability set as

$$\mathcal{C}_i(\bar{a}) = \{j \in [0,1] : a_{ij} \geq \bar{a}\},$$

where $\bar{a} > 0$ is a productivity threshold. The associated task breadth is

$$B_i(\bar{a}) = \mu\left(\mathcal{C}_i(\bar{a})\right),$$

where $\mu(\cdot)$ denotes Lebesgue measure.

Task breadth therefore measures the fraction of the production process over which an agent can operate at or above the benchmark productivity $\bar{a}$. The benchmark can be interpreted as the productivity required for an agent to be considered a viable producer of a task. Results below are robust to alternative definitions based on productivity ranks rather than a fixed threshold.

For expositional purposes, we first consider a symmetric benchmark in which agents have a common potential for general-purpose production but differ in the tasks at which they possess comparative advantage. Specifically, productivity can be written as

$$a_{ij} = \bar{a}_i g_{ij},$$

where $\bar{a}_i$ captures the general productive capability of agent i and g_{ij} captures task-specific comparative advantage.

The parameter $\bar{a}_i$ is therefore distinct from the distribution of comparative advantage. An increase in $\bar{a}_i$ makes an agent more productive across a broad set of tasks, whereas a change in g_{ij} can alter the relative productivity of the agent across tasks without necessarily changing its task breadth.

This distinction permits us to define a general-purpose capability parameter, denoted by G. A higher G expands the set of tasks over which an agent can operate at high productivity. Formally, we write

$$a_{ij}(G) = \bar{a}\, g_{ij}(G),$$

with

$$\frac{\partial B_i(G)}{\partial G} > 0.$$

The parameter G is thus a reduced-form measure of the technological generality of an agent.

The interpretation of G is deliberately broader than artificial intelligence. Any technology that allows a single productive unit to perform a wider range of activities can be represented in this way. Artificial intelligence is one particularly important example because its potential capabilities span tasks that historically have been distributed across distinct occupations.

2.3 Specialists and Generalists

We next formalize the distinction between specialization and generalization.

An allocation of tasks is represented by a measurable correspondence

$$S_i \subseteq [0, 1],$$

where S_i is the set of tasks assigned to agent i. Feasibility requires

$$S_i \subseteq \mathcal{C}_i$$

for every productive agent.

Define the breadth of agent i under allocation S as

$$b_i = \mu(S_i).$$

An agent is a *specialist* if its assigned task set is sufficiently narrow relative to the production process, while an agent is a *generalist* if it performs a sufficiently broad set of tasks. More fundamentally, rather than imposing a discrete distinction between specialists and generalists, we characterize an allocation by its distribution of task breadth,

$$F_B(b).$$

The degree of specialization in the economy can then be summarized by the dispersion of task assignments across agents, while the degree of generalization can be summarized by the average breadth of individual agents,

$$\bar{B} = \int b_i \, di.$$

Because every task must ultimately be assigned to at least one producer, an increase in the breadth of individual agents need not imply a proportional increase in the total number of tasks performed. The key organizational margin is therefore not the total amount of productive activity, but the number of distinct activities internalized by a given productive unit.

This observation motivates our terminology. We use *division of labor* to describe an allocation in which production is distributed across agents with relatively narrow task portfolios. We use *unification of labor* to describe an allocation in which a relatively broad set of production tasks is performed within the same agent. The associated technological forces are referred to as *specialization* and *generalization*, respectively.

The terminology is intended to emphasize that these are opposite margins of the same underlying object. Specialization reduces the breadth of an individual's productive portfolio in order to exploit differences in comparative advantage. Generalization expands that portfolio as technology makes it increasingly efficient for one agent to perform heterogeneous activities.

2.4 Productivity and Coordination

Dividing production across agents can increase task-level productivity because agents can concentrate on activities for which they have comparative advantage. It can also create costs because production must be coordinated across those agents.

Let $N(S)$ denote the number of distinct productive agents involved in an allocation S. We assume that coordination costs are increasing in the number of productive relationships that must be managed. A convenient reduced-form specification is

$$C(N) = \kappa\phi(N),$$

where

$$\kappa \geq 0, \qquad \phi'(N) > 0, \qquad \phi''(N) \geq 0.$$

The parameter κ captures the technological and institutional environment governing coordination. It encompasses, among other things, the costs of communicating information, transferring intermediate outputs, monitoring performance, contracting, and synchronizing decisions. The convexity of ϕ allows coordination costs to rise disproportionately as production becomes distributed across many agents.

This formulation deliberately abstracts from the particular source of coordination costs. What matters for our argument is that dividing production among more agents is not costless. The classical gains from specialization must therefore be weighed against the organizational costs created by specialization.

An allocation S generates gross output

$$Y^G(S)$$

and incurs coordination cost $C(N(S))$. Net output is

$$Y(S) = Y^G(S) - C(N(S)).$$

The equilibrium organization of production solves

$$S^* \in \arg\max_S \left\{ Y^G(S) - C(N(S)) \right\}.$$

This formulation makes clear why generalization can arise even when specialized agents retain comparative advantages. A generalist need not be the most productive agent at every individual task. It is sufficient that the productivity loss from assigning several tasks to one agent is smaller than the coordination cost saved by doing so.

2.5 The Specialization–Generalization Tradeoff

To make this tradeoff transparent, consider a production process containing K tasks. Suppose a specialized allocation assigns each task to the agent with the highest productivity in that task. Let the resulting average task productivity be A_S.

Now consider an agent capable of performing all K tasks at productivity $A_G(G)$. Under a fully unified allocation, gross production is determined by $A_G(G)$, while the number of agents requiring coordination falls relative to the specialized allocation.

The difference in net production between the two organizational forms can be written schematically as

$$\Delta(G) = \underbrace{A_G(G) - A_S}_{\text{productivity cost or gain}} + \underbrace{C(N_S) - C(N_G)}_{\text{coordination-cost saving}},$$

where N_S and N_G denote the number of agents involved under specialization and unification, respectively.

A fully generalized allocation is preferred whenever

$$A_G(G) - A_S > C(N_G) - C(N_S).$$

This condition captures the central mechanism of the paper. Generalization does not require the disappearance of comparative advantage. Nor does it require a generalist to dominate every specialist in every task. What matters is whether the productivity sacrifice from assigning a broader portfolio to an agent is smaller than the organizational costs avoided by reducing the number of agents across which production is divided.

As G increases, $A_G(G)$ can rise and the feasible task breadth of a generalist can expand. The right-hand side of the condition is determined by the organizational structure of production and need not change with G. Consequently, sufficiently large increases in general-purpose capability can cause the inequality to reverse.

This gives rise to the possibility of a technological threshold,

$$G^*,$$

such that

$$\begin{cases} \text{specialization is optimal,} & G < G^*, \\ \text{generalization is optimal,} & G > G^*. \end{cases}$$

The existence and uniqueness of such a threshold, as well as its dependence on the distribution of comparative advantage and the structure of coordination costs, are established formally below.

2.6 A Continuum Formulation

The discrete formulation is useful for intuition, but the continuum formulation clarifies the distinction between task productivity and task breadth. Let the productivity of a generalist agent be represented by

$$a_G(j; G) = a_S(j)\, \rho(j; G),$$

where $a_S(j)$ is the productivity attainable by the best specialized agent in task j and

$$0 \leq \rho(j; G) \leq 1.$$

The function $\rho(j; G)$ measures how close a generalist comes to the task-specific frontier. General-purpose capability expands as

$$\frac{\partial \rho(j; G)}{\partial G} \geq 0.$$

Define the generalist's effective task set as

$$\mathcal{G}(G) = \{j : \rho(j; G) \geq \bar{\rho}\},$$

with breadth

$$B_G(G) = \mu(\mathcal{G}(G)).$$

A purely task-specific technology increases $a_S(j)$ while leaving $B_G(G)$ unchanged. A general-purpose technology instead increases $B_G(G)$ and potentially raises productivity across many heterogeneous tasks simultaneously. The distinction is therefore not between "old" and "new" technology, but between technological progress along two different dimensions of the production function.

This distinction is essential for interpreting the model's implications for artificial intelligence. If AI merely raises productivity within existing occupations, its primary effect can be represented by an increase in $a_S(j)$. If AI also makes a single productive system capable of performing tasks that were previously outside its capability set, its effect is represented by an increase in $B_G(G)$. The latter can alter the equilibrium organization of production even when its effect on productivity within any individual task is modest.

2.7 Equilibrium Objects

The analysis that follows focuses on three equilibrium objects.

First, we characterize the *task allocation*

$$S^* = \{S_i^*\}_i,$$

which determines which agents perform which activities.

Second, we characterize the *task breadth distribution*

$$F_B^*(b),$$

which describes how broadly productive activities are distributed within individual agents.

Third, we characterize the *organizational concentration* of production. Let

$$\Omega = \int \omega_i \, di,$$

where ω_i denotes the share of production attributable to agent i. A high value of concentration corresponds to production being performed by relatively few agents, while a low value corresponds to production being distributed across many agents.

These objects should be distinguished. A more concentrated production structure need not imply less competition, just as a broader task portfolio need not imply higher productivity. The model is concerned with the endogenous organization of tasks, not with the market structure of the final-good industry.

The principal comparative-static parameter is the general-purpose capability G. We therefore study how

$$(S^*, F_B^*, \Omega^*)$$

responds to changes in G.

The central theoretical result will show that increasing G can produce a qualitative change in the equilibrium allocation rather than a smooth increase in output alone. In the specialized regime, increases in productivity primarily strengthen the existing division of labor. Once general-purpose capability becomes sufficiently high, however, further increases in G increase the optimal breadth of individual agents and reduce the gains from dividing production across them. Technological progress can therefore change the organizational principle of production itself.

2.8 Discussion

Several features of the environment are worth emphasizing.

First, specialization is not assumed. It emerges because heterogeneous task productivity creates gains from assigning activities to agents with comparative advantage. This allows the classical division of labor to appear as an equilibrium outcome rather than as a primitive feature of the model.

Second, generalization is not defined by superior productivity at every task. It is defined by the expansion of the set of tasks that an agent can perform sufficiently well. This makes generalization a statement about the *dimensionality* of productive capability rather than its level.

Third, unification of labor is not synonymous with autarky. A generalized economy can continue to contain multiple agents, markets, and exchange. What changes is the breadth of activities performed within each productive unit. In particular, the theory allows for interior allocations in which some activities remain specialized while others are unified.

Finally, the model does not presume that technological progress necessarily leads to generalization. The result depends on the interaction between general-purpose capability, comparative advantage, and coordination costs. If specialization generates sufficiently large task-specific productivity gains, or if coordination costs are negligible, the division of labor can remain optimal even when general-purpose capability improves. The generalized regime arises only when the expansion of task breadth becomes sufficiently valuable relative to the productivity gains and organizational benefits of specialization.

The next section derives the equilibrium allocation and establishes the conditions under which specialization, partial generalization, and full unification of labor arise.

3 Equilibrium Organization of Production

This section characterizes the equilibrium organization of production. The central result is that specialization and generalization arise as different solutions to the same allocation problem. When comparative advantage is sufficiently strong relative to the costs of coordinating production across agents, the efficient allocation assigns narrow task portfolios to different agents. As general-purpose capability increases, the optimal breadth of individual agents increases continuously and, under appropriate conditions, production passes through a partially generalized regime before becoming fully unified.

The analysis proceeds in three steps. We first characterize the efficient allocation conditional on the number of productive agents. We then characterize the optimal number of agents and establish the conditions under which specialization or generalization is preferred. Finally, we show that sufficiently large increases in general-purpose capability generate a structural transition in the organization of production.

3.1 The Planner's Problem

We begin with the social planner's problem. Since the final good is produced through a collection of complementary tasks, the planner chooses both the allocation of tasks across agents and the amount of effort devoted to each task.

For an allocation $S = \{S_i\}_i$, let $b_i = \mu(S_i)$ denote the breadth of agent i. Conditional on S, the productivity of agent i in task j is a_{ij}. The planner chooses effort allocations $\{e_{ij}\}$ to maximize net output:

$$\max_{\{S_i, e_{ij}\}} \left\{ \min_{j \in [0,1]} \int_i a_{ij} e_{ij} \, di - C(N(S)) \right\}, \tag{3}$$

subject to

$$\int_{S_i} e_{ij} \, dj \leq 1 \qquad \forall i. \tag{4}$$

The planner's problem contains two distinct margins. Conditional on an organizational structure, effort must be allocated toward the agents with the highest productivity in each task. But the planner must also determine whether those agents should be the same across tasks or whether production should be distributed across a larger set of agents.

The second margin is the object of interest. It determines the equilibrium degree of division of labor.

For much of the analysis, it is useful to impose a symmetric environment in which agents have identical general-purpose capability but heterogeneous comparative advantages. We therefore write productivity as

$$a_{ij}(G) = A(G) g_{ij}, \tag{5}$$

where $A(G)$ captures general-purpose capability and g_{ij} captures task-specific comparative advantage. The normalization is chosen such that

$$\int_0^1 g_{ij} \, dj = 1. \tag{6}$$

An increase in $A(G)$ therefore raises the productivity of an agent across tasks, while changes in the distribution of g_{ij} determine the gains from assigning different tasks to different agents.

The key distinction is that an increase in $A(G)$ need not imply that the agent becomes more general. Generalization requires that productivity remain sufficiently high over a larger set of tasks. We therefore allow the productivity frontier to depend on G through both its level and its breadth.

3.2 Task Assignment Conditional on Organizational Structure

Consider first an economy in which N agents are active in production. Suppose that the planner assigns each task to the agent with the highest productivity among the N active agents. Let

$$M_N(j;G) = \max_{i \leq N} a_{ij}(G) \tag{7}$$

denote the productivity frontier available for task j.

The average productivity of the task frontier is

$$P_N(G) = \int_0^1 M_N(j;G)\, dj. \tag{8}$$

Under the complementary production technology, aggregate gross productivity is governed by the least productive required task. To avoid obscuring the organizational mechanism with the details of task balancing, we normalize task requirements such that the planner can implement the frontier allocation without loss. Gross output conditional on N is then represented by

$$Y_N^G(G) = P_N(G). \tag{9}$$

The net value of an N-agent organization is

$$V_N(G) = P_N(G) - C(N). \tag{10}$$

The equilibrium number of productive agents is therefore

$$N^*(G) \in \arg\max_{N \geq 1} \left\{ P_N(G) - C(N) \right\}. \tag{11}$$

This representation makes the basic tradeoff transparent. Increasing N expands the set of comparative advantages available to the economy and therefore weakly increases the task-level productivity frontier. But increasing N also raises coordination costs.

The classical division of labor corresponds to a region in which the marginal gain from expanding the task-specific productivity frontier exceeds the marginal coordination cost. Generalization corresponds to a region in which the productivity frontier can be approximated by a smaller number of broad-capability agents.

3.3 The Value of Specialization

Let $P_1(G)$ denote the productivity frontier when a single agent performs the entire production process. Let

$$\Delta_N(G) = P_N(G) - P_1(G) \tag{12}$$

denote the gross productivity gain from making N agents available rather than relying on a single generalist.

We refer to $\Delta_N(G)$ as the *specialization gain*. It captures the value generated by access to heterogeneous comparative advantages.

The corresponding net gain from dividing production among N agents is

$$\Gamma_N(G) = \Delta_N(G) - [C(N) - C(1)]\,. \tag{13}$$

An N-agent organization is preferred to complete unification if and only if

$$\Gamma_N(G) \geq 0. \tag{14}$$

This condition is useful because it separates the two forces governing the organization of production. The first term is technological: heterogeneous agents provide comparative advantages that a single agent cannot fully exploit. The second is organizational: exploiting those comparative advantages requires production to be divided across agents.

The standard theory of specialization focuses primarily on the first term. Our framework emphasizes that its importance is endogenous to general-purpose capability.

3.4 General-Purpose Capability and the Erosion of Specialization Gains

Suppose that an increase in G raises the productivity of an agent across tasks and expands the set of tasks for which the agent can operate near the task-specific frontier. We impose the following condition.

Assumption 1 (Generalization)**.** *For any task j,*

$$\frac{\partial a_{ij}(G)}{\partial G} \geq 0,$$

and for any pair of tasks j, k for which agent i has sufficiently high productivity,

$$\frac{\partial}{\partial G} \left| a_{ij}(G) - a_{ik}(G) \right| \leq 0.$$

The second condition captures the idea that general-purpose technological progress compresses the dispersion of an agent's productivity across tasks. The agent becomes not merely more productive, but more uniformly capable.

This has a direct implication for the value of comparative advantage. If a generalist becomes increasingly productive in tasks outside its original area of comparative advantage, the productivity gap between the best specialized agent and the generalist narrows.

Formally, suppose that

$$\frac{\partial \Delta_N(G)}{\partial G} < 0 \qquad \text{for } N > 1. \tag{15}$$

The specialization gain is then decreasing in general-purpose capability.

This condition is not necessary for all forms of generalization, but it identifies the economically important case. General-purpose technological progress can simultaneously raise the productivity of specialized and generalized production while reducing the *relative* advantage of dividing production among agents.

This distinction is central. It is entirely possible for

$$\frac{\partial P_N(G)}{\partial G} > 0$$

for every N while also having

$$\frac{\partial \Gamma_N(G)}{\partial G} < 0.$$

Thus, general-purpose technological progress can raise productivity under every organizational structure while nevertheless making specialization less attractive.

3.5 The Optimal Number of Agents

The planner chooses N according to equation (11). Define the marginal specialization gain from adding the Nth agent as

$$MS_N(G) = P_N(G) - P_{N-1}(G), \tag{16}$$

and the marginal coordination cost as

$$MC_N = C(N) - C(N-1). \tag{17}$$

The Nth agent is active whenever

$$MS_N(G) \geq MC_N. \tag{18}$$

This is the fundamental equilibrium condition.

Under conventional specialization, the marginal specialization gain is large because a new agent brings a distinct comparative advantage that cannot easily be replicated by existing agents. As general-purpose capability increases, however, existing agents become capable of performing a broader range of tasks. The incremental value of adding another specialized agent falls.

The resulting comparative static is immediate.

[Declining Optimal Division of Labor] Suppose that $MS_N(G)$ is strictly decreasing in G for every $N > 1$ and that MC_N is weakly increasing in N. Then the optimal number of productive agents $N^*(G)$ is weakly decreasing in G.

Proof. For any $N > 1$, the condition for the Nth agent to be active is

$$MS_N(G) \geq MC_N.$$

By assumption, $MS_N(G)$ decreases in G, while MC_N does not. Hence an increase in G cannot cause an agent that was previously inactive to become necessary solely through the specialization margin. It can only preserve or eliminate existing agents from the optimal allocation. Therefore $N^*(G)$ is weakly decreasing in G. □

The proposition provides the first formal result of the paper. It does not say that general-purpose technology necessarily reduces employment or the number of agents in the economy. Rather, it says that when the relevant technological change reduces the marginal value of comparative advantage, the efficient organization of a fixed production process becomes less divided across agents.

3.6 The Specialization Equilibrium

We can now define the classical regime.

[Specialization Equilibrium] An equilibrium is a *specialization equilibrium* if

$$N^*(G) > 1$$

and the equilibrium task allocation assigns different subsets of tasks to agents such that

$$\mu(S_i^*) < 1$$

for the active agents.

The equilibrium is characterized by a positive division-of-labor margin. Different agents perform different tasks because their comparative advantages are sufficiently valuable to offset the costs of coordination.

A sufficient condition is

$$MS_N(G) > MC_N \qquad \text{for some } N > 1. \tag{19}$$

If this condition holds for all $N \leq \bar{N}$, the equilibrium involves at least $\bar{N}$ productive agents.

The important point is that specialization is not imposed by the production function. It is generated by the interaction between heterogeneous productivity and coordination costs.

3.7 The Unification Equilibrium

At the opposite extreme, consider an allocation in which a single general-purpose agent performs the entire production process.

[Unification Equilibrium] An equilibrium is a *unification equilibrium* if

$$N^*(G) = 1$$

and the active agent performs the full task set,

$$S_1^* = [0, 1].$$

The condition for complete unification is

$$P_1(G) - C(1) \geq P_N(G) - C(N) \qquad \forall N > 1. \tag{20}$$

Equivalently,

$$\Delta_N(G) \leq C(N) - C(1) \qquad \forall N > 1. \tag{21}$$

The interpretation is particularly revealing. Complete unification does not require the generalist to be the most productive possible agent at every task. It requires only that the aggregate gain from exploiting specialized comparative advantages be smaller than the coordination cost required to exploit them.

Thus, unification can emerge even in an economy in which comparative advantage continues to exist.

3.8 A Threshold for the Organizational Regime

The previous results allow us to define the critical level of general-purpose capability.

Let

$$G^* = \inf \left\{ G : \Delta_N(G) \leq C(N) - C(1) \quad \forall N > 1 \right\}. \tag{22}$$

Provided the set is nonempty, G^* is the lowest level of general-purpose capability at which complete unification is optimal.

[Specialization-to-Generalization Transition] Suppose that:

1. $P_N(G)$ is increasing in G for every N;
2. $\Delta_N(G)$ is strictly decreasing in G for every $N > 1$;
3. $C(N)$ is increasing and convex in N;
4. $\lim_{G \to \infty} \Delta_N(G) = 0$ for every finite N.

Then there exists a finite threshold G^* such that:

1. for $G < G^*$, at least one specialized allocation is strictly optimal;
2. at $G = G^*$, the specialized and unified allocations can coexist;
3. for $G > G^*$, the fully unified allocation is uniquely optimal.

Proof. By assumption 2, $\Delta_N(G)$ is strictly decreasing in G. By assumption 4, for every $N > 1$ there exists a sufficiently large G such that

$$\Delta_N(G) < C(N) - C(1).$$

Since $C(N)$ is increasing, the condition holds simultaneously for all $N > 1$ at sufficiently large G. Hence the unified allocation is optimal for sufficiently large G.

For sufficiently low G, suppose that there exists an $N > 1$ such that

$$\Delta_N(G) > C(N) - C(1).$$

Then specialization strictly dominates unification. Continuity of $\Delta_N(G)$ implies the existence of a threshold at which equality holds. Strict monotonicity implies that the inequality reverses only once. The result follows. □

The theorem formalizes the central claim of the paper. The division of labor and the unification of labor are not competing assumptions about the economy. They are equilibrium outcomes generated by the same production technology under different levels of general-purpose capability.

3.9 Partial Generalization

Complete specialization and complete unification are useful limiting cases, but the most interesting equilibrium may lie between them. To characterize this possibility, suppose that the task set can be partitioned into two regions,

$$[0, 1] = \mathcal{S}(G) \cup \mathcal{U}(G),$$

where $\mathcal{S}(G)$ denotes tasks assigned according to comparative advantage across specialized agents and $\mathcal{U}(G)$ denotes tasks internalized by a generalist.

Define the generalized share of production as

$$\theta(G) = \mu(\mathcal{U}(G)). \tag{23}$$

We refer to $\theta(G)$ as the *degree of generalization.* Thus,

$$\theta(G) = 0$$

corresponds to complete specialization, while

$$\theta(G) = 1$$

corresponds to complete unification.

The equilibrium value of θ solves

$$\theta^*(G) \in \arg\max_{\theta \in [0,1]} \{Y(\theta, G)\}. \tag{24}$$

Suppose the marginal productivity advantage of specialization over generalization for task j is

$$d(j, G) = a_S(j) - a_G(j; G), \tag{25}$$

and let $\lambda(G)$ denote the marginal coordination cost generated by assigning an additional task to a separate specialist. A task is assigned to a specialist whenever

$$d(j, G) > \lambda(G), \tag{26}$$

and to the generalist whenever

$$d(j, G) \leq \lambda(G). \tag{27}$$

If general-purpose capability reduces $d(j, G)$ for every task, the set of tasks satisfying the latter inequality expands monotonically with G.

[Monotone Generalization] Suppose that

$$\frac{\partial d(j, G)}{\partial G} < 0$$

for almost every task j. Then the equilibrium degree of generalization $\theta^*(G)$ is weakly increasing in G.

Proof. A task assigned to the generalist satisfies

$$d(j, G) \leq \lambda(G).$$

Because $d(j, G)$ is decreasing in G, an increase in G weakly expands the set of tasks satisfying this condition. Therefore $\mu(\mathcal{U}(G))$ weakly increases, implying that $\theta^*(G)$ is weakly increasing. $\square$

This result is important because it shows that the transition need not be discontinuous. General-purpose technological progress can first induce the internalization of a subset of tasks, leaving the remaining activities specialized. As G rises further, the generalized region expands. Full unification is the limiting case.

The resulting organizational transition can therefore be represented as

$$\theta^*(G) = \begin{cases} 0, & G < G_0, \\ \in (0,1), & G_0 \leq G < G_1, \\ 1, & G \geq G_1. \end{cases} \tag{28}$$

The intermediate region is economically important. It corresponds to an economy in which AI or another general-purpose technology does not eliminate specialization but progressively changes its scope. Some tasks remain valuable precisely because they require highly specialized capabilities, while other tasks become efficiently internalized by general-purpose agents.

3.10 When Does Generalization Fail to Occur?

The theory does not imply that all sufficiently capable technologies eliminate specialization. The transition depends on the structure of comparative advantage.

Suppose, for example, that task-specific productivity satisfies

$$\sup_i a_{ij}(G) - \sup_{i \in \mathcal{G}} a_{ij}(G) \geq \delta > 0 \tag{29}$$

for a non-negligible set of tasks even as G becomes large. These tasks remain sufficiently difficult for generalists relative to specialists. In this case,

$$\liminf_{G \to \infty} \Delta_N(G) > 0,$$

and complete unification need not emerge.

This observation provides an important qualification. Generalization is not equivalent to the disappearance of expertise. An economy can contain highly capable general-purpose agents alongside specialists whose capabilities remain sufficiently differentiated. The theory predicts unification only when general-purpose capability becomes sufficiently close to the task-specific frontier.

This distinction also clarifies the meaning of the AGI benchmark. The relevant economic condition is not that an artificial system be "intelligent" in an abstract sense. What matters is whether the system can perform a sufficiently broad set of economically relevant tasks at sufficiently high productivity.

3.11 Generalization and the Number of Agents

The preceding analysis also establishes a relationship between task breadth and the number of productive agents.

Let $\bar{B}^*(G)$ denote equilibrium average task breadth. Under the monotonicity conditions above,

$$\frac{d\bar{B}^*(G)}{dG} \geq 0, \tag{30}$$

while

$$\frac{dN^*(G)}{dG} \leq 0. \tag{31}$$

Thus general-purpose technological progress can generate a simultaneous increase in the breadth of individual productive portfolios and a decline in the number of agents required to organize a given production process.

These comparative statics should not be interpreted as a prediction that the total number of workers in an economy must fall. The model holds the production process fixed. In a general equilibrium setting, lower organizational costs and higher productivity can generate new goods, services, and tasks, potentially increasing the demand for productive agents elsewhere in the economy. The result is instead about the *organization of a given production process*: as agents become more general, fewer distinct agents need be coordinated to perform its constituent activities.

3.12 A Reinterpretation of Comparative Advantage

The model also yields a useful reinterpretation of comparative advantage.

In the standard specialization framework, comparative advantage determines which agent should perform which task. In the generalized framework, comparative advantage determines which tasks remain worth separating across agents.

This is a subtle but important change in the economic role of heterogeneity. Let

$$R_j(G) = \max_i a_{ij}(G) - a_G(j;G) \tag{32}$$

denote the residual comparative advantage of a specialist over a generalist in task j. The task remains specialized only if

$$R_j(G) > \lambda(G). \tag{33}$$

As G rises, the relevant question is therefore no longer whether comparative advantage exists. It almost always will. The relevant question is whether comparative advantage is *large enough to justify organizational separation.*

This distinction provides a useful way to think about the economic significance of AGI. General-purpose technology need not eliminate heterogeneity in productive ability. It can instead reduce the economic value of that heterogeneity by making the organizational cost of exploiting it larger relative to the remaining productivity differences.

3.13 Main Result

The preceding analysis can be summarized in the following result.

[Unification of Labor] Consider an economy in which:

1. production requires a collection of heterogeneous tasks;
2. agents differ in task-specific productivity;
3. coordination costs increase with the number of agents across whom production is divided;
4. general-purpose technological progress expands the set of tasks an individual agent can perform at high productivity; and
5. general-purpose capability weakly reduces the residual productivity advantage of specialists over generalists.

Then an increase in general-purpose capability weakly increases equilibrium task breadth and weakly decreases the optimal degree of division of labor. If general-purpose capability becomes sufficiently high that the residual gains from comparative advantage are smaller than the coordination costs required to exploit them, the equilibrium converges to unification of labor.

The theorem provides the central organizing principle for the remainder of the paper:

$$\boxed{\text{General-purpose capability} \uparrow \quad \Longrightarrow \quad \text{task breadth} \uparrow \quad \Longrightarrow \quad \text{division of labor} \downarrow}. \tag{34}$$

The result does not overturn the classical theory of specialization. It embeds it. Specialization remains optimal whenever heterogeneous comparative advantage is sufficiently valuable relative to the costs of coordination. What changes in the presence of sufficiently general productive capabilities is the technological boundary of that result.

In the classical regime, the question is which agent should perform each task. In the generalized regime, the question becomes which tasks should be performed by the same agent.

This change in the allocation problem is the central economic consequence of generalization.

4 Comparative Statics and the Transition to Generalization

The preceding section established the basic equilibrium tradeoff between specialization and generalization. This section studies how that tradeoff responds to the fundamental parameters of the production environment. Three comparative statics are central. First, an increase in general-purpose capability increases the set of tasks that an individual agent can perform effectively and thereby reduces the residual value of specialization. Second, greater heterogeneity in task-specific comparative advantage strengthens the case for division of labor and delays the transition to generalization. Third, higher coordination costs accelerate the transition toward unified production.

These forces generate a phase diagram in which the equilibrium organization of production moves from specialization to partial generalization and, under sufficiently strong general-purpose capability, to complete unification. The transition is not simply a consequence of higher productivity. It arises because general-purpose technological progress changes the relative productivity of alternative organizational forms.

4.1 General-Purpose Capability

We begin with the central comparative static. Let G denote the economy's level of general-purpose capability. For a task j, define the productivity frontier available to a generalist as

$$a_G(j;G) = a_S(j)\rho(j;G), \tag{35}$$

where $a_S(j)$ is the productivity of the best specialized producer of task j and

$$0 \leq \rho(j;G) \leq 1. \tag{36}$$

The function $\rho(j;G)$ measures the extent to which a general-purpose agent can approach the task-specific frontier. Generalization corresponds to an increase in $\rho(j;G)$ over a broad set of tasks.

For each task, define the residual specialization advantage

$$r(j, G) = a_S(j) - a_G(j; G) = a_S(j)[1 - \rho(j; G)]. \tag{37}$$

The residual advantage is the productivity gain from assigning task j to a specialist rather than to a generalist.

The key assumption underlying the transition is

$$\frac{\partial \rho(j; G)}{\partial G} > 0 \qquad \text{for almost every } j, \tag{38}$$

which implies

$$\frac{\partial r(j, G)}{\partial G} < 0. \tag{39}$$

Thus, general-purpose technological progress need not eliminate comparative advantage. Instead, it reduces the productivity advantage that remains after the generalist becomes more capable.

Let $\lambda(G)$ denote the marginal organizational cost of assigning a task to an additional specialized agent. A task is optimally assigned to a specialist whenever

$$r(j, G) > \lambda(G), \tag{40}$$

and to the generalist whenever

$$r(j, G) \leq \lambda(G). \tag{41}$$

Consequently, the equilibrium generalized task set is

$$\mathcal{U}(G) = \{j : r(j, G) \leq \lambda(G)\}, \tag{42}$$

and its measure,

$$\theta(G) = \mu(\mathcal{U}(G)), \tag{43}$$

is the degree of generalization.

[General-Purpose Capability] Suppose that $r(j, G)$ is strictly decreasing in G for almost every task and that $\lambda(G)$ is weakly increasing in G. Then the equilibrium degree of generalization satisfies

$$\frac{d\theta(G)}{dG} \geq 0. \tag{44}$$

If the set of tasks for which $r(j, G) = \lambda(G)$ has measure zero, the inequality is strict whenever an increase in G changes the set of tasks assigned to the generalist.

Proof. An increase in G weakly lowers $r(j, G)$ for every task. Since $\lambda(G)$ is weakly increasing, the set satisfying

$$r(j, G) \leq \lambda(G)$$

weakly expands. Therefore $\mu(\mathcal{U}(G))$ weakly increases. If the boundary set has measure zero and the increase in G causes a positive-measure set of tasks to cross the assignment boundary, the increase in $\theta(G)$ is strict. □

The proposition establishes a central comparative static: greater general-purpose capability increases the breadth of production internalized by individual agents. Importantly, this result concerns the organization of production rather than output alone. Output can increase under every organizational structure while the optimal allocation simultaneously becomes less specialized.

4.2 The Intensive Margin of Generalization

The preceding result describes the extensive margin—whether a task is assigned to a generalist or a specialist. We next consider the intensive margin: how much broader an individual agent's task portfolio becomes.

Suppose that the productivity of a generalist over tasks is described by a smooth capability function

$$a_G(j; G) = A(G) \exp[-\sigma(G)d(j)], \tag{45}$$

where $d(j)$ measures the distance of task j from the agent's core capability, $A(G)$ is general productivity, and $\sigma(G)$ captures the rate at which productivity declines as the agent moves across tasks.

General-purpose technological progress can operate through either channel. It can increase $A(G)$, raising productivity throughout the task space, or reduce $\sigma(G)$, flattening the productivity gradient across tasks.

The latter is particularly important for generalization. A decline in $\sigma(G)$ means that moving into a new task entails a smaller productivity penalty. Thus, the technological change relevant for unification is not necessarily an increase in the productivity frontier at its peak. It can instead be a flattening of the frontier.

Define an agent's effective task breadth as

$$B(G) = \mu \left\{ j : a_G(j; G) \geq \bar{a} \right\}. \tag{46}$$

Under the capability function above,

$$B(G) = \mu \left\{ j : d(j) \leq \frac{1}{\sigma(G)} \log \left(\frac{A(G)}{\bar{a}} \right) \right\}. \tag{47}$$

Hence

$$\frac{dB(G)}{dG} > 0 \tag{48}$$

whenever general-purpose progress either raises $A(G)$ or lowers $\sigma(G)$ sufficiently.

This provides a useful interpretation of generalization. The relevant technological innovation is one that expands the feasible set of productive activities within a single agent's capability frontier.

4.3 Comparative Advantage

General-purpose capability is not the only determinant of the equilibrium organization of production. The strength of comparative advantage also matters.

Let the productivity of specialist i in task j be

$$a_{ij} = \bar{a}_j \exp(\varepsilon_{ij}), \tag{49}$$

where ε_{ij} captures agent-task-specific comparative advantage. Let σ_ε^2 denote the dispersion of these comparative advantages.

A higher σ_ε^2 increases the gains from assigning each task to the agent with the highest task-specific productivity. In other words, it increases the value of specialization.

Let

$$\Delta(G, \sigma_\varepsilon) \tag{50}$$

denote the aggregate productivity gain from exploiting comparative advantage relative to a unified allocation.

Under standard regularity conditions,

$$\frac{\partial \Delta}{\partial \sigma_\varepsilon} > 0. \tag{51}$$

The condition for generalization can therefore be written as

$$\Delta(G, \sigma_\varepsilon) \leq C(N) - C(1). \tag{52}$$

Since Δ is increasing in comparative-advantage heterogeneity and decreasing in general-purpose capability, the critical level of general-purpose capability satisfies

$$G^* = G^*(\sigma_\varepsilon, \kappa). \tag{53}$$

The comparative statics are

$$\frac{\partial G^*}{\partial \sigma_\varepsilon} > 0 \tag{54}$$

and

$$\frac{\partial G^*}{\partial \kappa} < 0. \tag{55}$$

Greater heterogeneity in comparative advantage delays generalization, whereas greater coordination costs accelerate it.

[Comparative Advantage and the Generalization Threshold] Suppose that the aggregate specialization gain $\Delta(G, \sigma_\varepsilon)$ is decreasing in G and increasing in σ_ε. Then the critical level of general-purpose capability required for complete unification satisfies

$$\frac{\partial G^*}{\partial \sigma_\varepsilon} > 0. \tag{56}$$

Proof. At the threshold,

$$\Delta(G^*, \sigma_\varepsilon) = C(N) - C(1).$$

Differentiating implicitly with respect to σ_ε yields

$$\frac{\partial G^*}{\partial \sigma_\varepsilon} = -\frac{\partial \Delta / \partial \sigma_\varepsilon}{\partial \Delta / \partial G}.$$

The numerator is positive and the denominator is negative. Hence

$$\frac{\partial G^*}{\partial \sigma_\varepsilon} > 0.$$

□

The result gives a precise meaning to the claim that generalization does not imply the disappearance of specialization. If comparative advantages remain sufficiently heterogeneous, specialized agents continue to generate substantial productivity gains even when general-purpose capability is high.

4.4 Coordination Costs

We next consider the role of coordination.

Let

$$C(N) = \kappa\phi(N), \tag{57}$$

where $\phi'(N) > 0$ and $\phi''(N) \geq 0$. An increase in κ raises the cost of organizing production across separate agents without directly affecting their technological productivity.

The threshold condition for unification is

$$\Delta_N(G) \leq \kappa[\phi(N) - \phi(1)]. \tag{58}$$

Thus

$$G^* = G^*(\kappa, \sigma_\varepsilon), \tag{59}$$

with

$$\frac{\partial G^*}{\partial \kappa} < 0. \tag{60}$$

An economy with costly communication, contracting, monitoring, or synchronization therefore has a stronger incentive to internalize heterogeneous tasks within the same productive agent.

This result provides an important qualification to the usual interpretation of general-purpose technologies. The economic effect of a generalist depends on the organizational environment into which it is introduced. The same technological capability can generate very different equilibrium structures in economies with different coordination technologies.

Indeed, the model predicts that generalization and low coordination costs are substitutes. When communication and coordination are nearly frictionless, the productivity gains from specialization can be realized without much organizational cost. When coordination is expensive, a general-purpose agent can generate substantial value by internalizing activities even if it is somewhat less productive than a specialist at individual tasks.

4.5 The Phase Diagram

The preceding comparative statics can be summarized using the (G, σ_ε) plane. For each level of comparative-advantage heterogeneity, there exists a critical level of general-purpose capability,

$$G^*(\sigma_\varepsilon, \kappa), \tag{61}$$

separating the specialization and generalization regimes.

For low G and high σ_ε, specialization is optimal. For high G and low σ_ε, unification is optimal. Between these extremes lies a region of partial generalization.

The equilibrium can therefore be represented schematically as

	Low G	Intermediate G	High G
High σ_ε	Specialization	Partial specialization	Mixed organization
Low σ_ε	Specialization	Partial generalization	Unification

(62)

The important feature is that the boundary between regimes is endogenous. There is no universal level of AI capability at which specialization becomes obsolete. The transition depends on the structure of the economy in which the technology operates.

This observation also suggests why different industries may experience very different organizational consequences from the same general-purpose technology. Industries with highly differentiated comparative advantages and low coordination costs may remain specialized even at high levels of general-purpose capability. Industries characterized by substantial coordination frictions and relatively substitutable task capabilities may generalize much earlier.

4.6 The Transition Path

The most informative case is one in which general-purpose capability rises continuously over time. Let

$$G_t \tag{63}$$

denote the general-purpose capability available in period t, with

$$G_{t+1} > G_t. \tag{64}$$

The equilibrium degree of generalization is

$$\theta_t = \theta^*(G_t). \tag{65}$$

Under the conditions above,

$$\theta_{t+1} \geq \theta_t. \tag{66}$$

The transition can therefore be represented as a sequence:

$$\text{specialization} \to \text{partial generalization} \to \text{unification}. \tag{67}$$

This sequence is not inevitable. It occurs only if general-purpose capability crosses the relevant thresholds. But when it does, the change in organization is qualitatively different from conventional productivity growth.

Under a task-specific technological improvement, productivity rises while the allocation of tasks may remain unchanged:

$$a_S(j) \uparrow \quad \Rightarrow \quad Y \uparrow, \qquad \theta^* \text{ unchanged}. \tag{68}$$

Under general-purpose technological progress,

$$a_G(j; G) \uparrow \quad \text{across a broad task set}, \tag{69}$$

and therefore

$$\theta^* \uparrow. \tag{70}$$

The distinction is consequently between a technology that makes an existing organization more productive and a technology that changes which organization is productive.

4.7 A Discrete Threshold and a Continuous Transition

The model permits both continuous and discontinuous transitions.

Suppose first that the task-specific residual advantage $r(j, G)$ varies continuously across tasks. Then as G rises, tasks cross the assignment boundary one at a time:

$$r(j, G) = \lambda(G). \tag{71}$$

The generalized share $\theta(G)$ therefore changes continuously.

If instead there are substantial masses of tasks with similar productivity characteristics, several tasks may cross the boundary simultaneously. The generalized share can then exhibit discrete jumps.

More generally, let the distribution of residual specialization advantages be

$$F_R(r; G). \tag{72}$$

Then

$$\theta(G) = F_R(\lambda(G); G). \tag{73}$$

Differentiating,

$$\frac{d\theta}{dG} = f_R(\lambda(G); G)\lambda'(G) + \frac{\partial F_R(\lambda(G); G)}{\partial G}. \tag{74}$$

The first term captures changes in the organizational cost of separating tasks; the second captures the technological compression of residual comparative advantage.

This formulation shows that the speed of generalization depends on the density of tasks near the organizational boundary. If many tasks are close to indifferent between specialization and generalization, small improvements in general-purpose capability can generate large organizational changes.

4.8 Generalization as a Change in the Elasticity of Task Breadth

The comparative statics can also be summarized using the elasticity of task breadth with respect to general-purpose capability:

$$\varepsilon_{B,G} = \frac{dB}{dG}\frac{G}{B}. \tag{75}$$

A technology with a high $\varepsilon_{B,G}$ is highly generalizing: relatively small increases in technological capability produce large expansions in the range of tasks an agent can perform.

By contrast, a technology with

$$\varepsilon_{B,G} \approx 0 \tag{76}$$

is primarily task-specific. Such a technology can generate substantial productivity growth without materially changing the organization of production.

This provides a natural taxonomy of technological progress:

$$\begin{aligned} &\text{Task-specific progress} && \Rightarrow \quad \text{productivity within tasks,} \\ &\text{General-purpose progress} && \Rightarrow \quad \text{breadth across tasks,} \\ &\text{Highly general-purpose progress} && \Rightarrow \quad \text{potential unification of labor.} \end{aligned} \tag{77}$$

The distinction is useful because the same aggregate productivity growth can have very different implications for economic organization depending on which margin generates it.

4.9 Implications for the Number of Specializations

Let $K^*(G)$ denote the effective number of distinct specialized occupations or task portfolios in equilibrium. As general-purpose capability rises, the number of tasks internalized by a generalist increases. Holding the underlying production process fixed, this implies

$$\frac{dK^*(G)}{dG} \leq 0. \tag{78}$$

This result should again be interpreted as an organizational comparative static rather than a forecast for aggregate employment.

The distinction matters because the economy can respond to productivity gains by expanding the set of goods, services, and activities that are produced. In a richer general-equilibrium environment, the number of economically valuable tasks may therefore increase even while the number of distinct tasks performed by each productive agent also increases.

The model consequently predicts a change in the *mapping from agents to tasks*, not necessarily a contraction in the total set of economically meaningful activities.

4.10 A Generalization of the Classical Division of Labor

The comparative statics suggest a broader interpretation of the classical division of labor.

Suppose that the economy begins in a regime in which

$$G \ll G^*. \tag{79}$$

Then the productivity frontier is sufficiently heterogeneous that specialization is optimal. Agents divide production according to comparative advantage, and coordination costs are justified by the resulting productivity gains.

Now suppose that technological progress increases G. The economy does not immediately abandon specialization. Instead, the residual productivity advantage of specialists falls for some tasks. Those tasks become internalized by broader agents. The economy enters a regime of partial generalization.

Finally, if

$$G \geq G^*, \tag{80}$$

the residual gains from specialization are insufficient to justify organizational separation. Production becomes unified.

The classical division of labor is therefore recovered as the low-general-purpose-capability equilibrium of a more general model.

This can be summarized by the mapping

$$\boxed{\begin{array}{ccc} \text{Narrow capabilities} & \longrightarrow & \text{Division of labor} \\ \downarrow & & \downarrow \\ \text{Broad capabilities} & \longrightarrow & \text{Unification of labor.} \end{array}} \tag{81}$$

The fundamental object of economic organization is therefore not specialization itself, but the optimal allocation of task breadth across productive agents.

4.11 The AGI Limit

The limiting case clarifies the interpretation of the theory.

Suppose that there exists a sequence of general-purpose technologies indexed by G such that

$$\lim_{G \to \infty} \rho(j; G) = 1 \tag{82}$$

for almost every economically relevant task j.

Then

$$\lim_{G \to \infty} r(j, G) = 0. \tag{83}$$

Consequently, for any strictly positive organizational cost of dividing production,

$$\lambda > 0, \tag{84}$$

there exists sufficiently large G such that

$$r(j, G) < \lambda \tag{85}$$

for almost every task.

Hence

$$\lim_{G\to\infty} \theta(G) = 1. \tag{86}$$

[General-Purpose Limit] If a general-purpose technology converges to the task-specific productivity frontier over almost the entire task space, and coordination costs remain strictly positive, the equilibrium organization of a fixed production process converges to unification of labor.

The corollary is the strongest version of the paper's central idea. Under sufficiently general productive capability, specialization ceases to be valuable not because comparative advantage disappears, but because the residual productivity gains from exploiting comparative advantage become arbitrarily small relative to the organizational cost of dividing production.

In this limit, the economic problem changes character. Under specialization, the central allocation problem is

$$\text{“Who should perform each task?”} \tag{87}$$

Under unification, it becomes

$$\text{“Which tasks should be performed by the same agent?”} \tag{88}$$

The distinction is more than semantic. It represents a change in the dimension along which production is optimally organized.

4.12 Discussion

The comparative statics establish three broad conclusions.

First, general-purpose capability is organizationally distinct from conventional productivity growth. Its defining effect is to expand task breadth and reduce the residual value of dividing production across agents.

Second, the transition to generalization is mediated by economic structure. Strong comparative advantage delays the transition, while costly coordination accelerates it. The effects of AI or other general-purpose technologies therefore depend on the environment in which they are deployed.

Third, the generalized economy is not necessarily an economy without specialization. The relevant prediction is a decline in the optimal degree of division of labor. Specialized agents continue to exist whenever their residual productivity advantages exceed the organizational costs of coordinating them with other agents.

The theory therefore replaces a binary question—whether specialization survives—with a continuous one: *how much specialization is optimal given the breadth of productive capability?*

This distinction will be important for the applications developed below. In particular, the implications for firms, occupations, wages, and markets depend not simply on whether AI can perform a task, but on whether a single productive system can profitably internalize a sufficiently broad set of tasks that were previously distributed across agents.

The next section turns to these organizational implications and studies how generalization changes the boundary of the firm, the composition of occupations, and the allocation of productive activity across agents.

5 Extensions: Heterogeneity, Coordination, and Partial Generalization

The baseline model deliberately isolates the central mechanism by separating task-specific comparative advantage from general-purpose capability. This section relaxes that separation along three dimensions. First, agents may differ in both their general-purpose capability and their comparative advantages. Second, coordination costs may themselves respond to technological progress, including the possibility that general-purpose technologies reduce the cost of coordinating specialized agents. Third, specialization and generalization may coexist endogenously, generating a richer distribution of task breadth than the polar cases considered above.

These extensions serve two purposes. They establish that the central results do not depend on a representative generalist, and they clarify the conditions under which general-purpose technological progress produces partial rather than complete unification. They also identify an important distinction between two effects of artificial intelligence: AI may make an individual agent capable of performing more tasks, but it may simultaneously make it cheaper for many specialized agents to coordinate. The organizational consequences of AI therefore depend on the relative strength of its *breadth* and *coordination* effects.

5.1 Heterogeneous General-Purpose Capability

The baseline model assumes that agents share a common general-purpose capability. We now allow agents to differ in both their general-purpose productivity and their task-specific comparative advantages.

Let agent i have general-purpose capability G_i and task-specific productivity

$$a_{ij} = A(G_i)g_{ij}, \tag{89}$$

where $A(G_i)$ is increasing in G_i and g_{ij} captures the agent's comparative advantage in task j.

The distribution of productive capability is therefore two-dimensional. An agent may be highly capable in general but possess little comparative advantage in any particular task, while another may have lower general-purpose capability but an unusually strong advantage in a narrow domain.

This distinction matters because generalization need not occur uniformly across agents. A sufficiently capable generalist may internalize tasks that remain specialized for other agents.

Let

$$G_{(1)} \geq G_{(2)} \geq \cdots \tag{90}$$

denote the order statistics of general-purpose capability. For any task j, the relevant productivity frontier is

$$M(j) = \max_i \{A(G_i)g_{ij}\}. \tag{91}$$

The task allocation problem now has two margins. The first determines which agent has the highest productivity in each task. The second determines whether that agent should perform additional tasks for which another agent has a comparative advantage.

This distinction generates a richer form of specialization. Agents with high G_i tend to have broad task portfolios, while agents with low G_i can remain specialized in tasks for which their g_{ij} is sufficiently high.

5.1.1 Selection into Generalization

Consider two agents, i and k, with

$$G_i > G_k. \tag{92}$$

Suppose that agent k has a comparative advantage in task j such that

$$g_{kj} > g_{ij}. \tag{93}$$

The task will nevertheless be assigned to agent i whenever

$$A(G_i)g_{ij} - A(G_k)g_{kj} > \lambda, \tag{94}$$

where λ is the organizational cost associated with maintaining separate task assignments.

Thus, the generalist need not possess the highest task-specific productivity. A sufficiently high level of general-purpose capability can compensate for a comparative disadvantage.

This yields a selection condition for generalization:

$$\frac{A(G_i)}{A(G_k)} > \frac{g_{kj}}{g_{ij}} + \frac{\lambda}{A(G_k)g_{ij}}. \tag{95}$$

The first term captures the technological tradeoff between general-purpose capability and comparative advantage. The second captures the organizational cost of maintaining separate production.

As G_i increases, the set of tasks for which agent i is the efficient internal producer expands.

[Selection into Generalization] Suppose that $A'(G) > 0$ and that the organizational cost of maintaining separate task assignments is nonnegative. Then, holding task-specific comparative advantages fixed, the set of tasks optimally assigned to an agent with general-purpose capability G_i is weakly increasing in G_i.

The proof follows directly from the monotonicity of $A(G_i)$ and is omitted from the main text. A complete characterization with continuous agent heterogeneity is provided in Appendix **??**.

The proposition implies that generalization is not necessarily an economy-wide transformation affecting all agents equally. Instead, it can begin with a subset of highly capable agents and spread as the distribution of general-purpose capability shifts.

This distinction is important for interpreting AI. If artificial agents differ in capability, cost, or access to complementary inputs, the emergence of generalization will initially be concentrated among the most capable systems. The relevant equilibrium object is therefore not simply the average task breadth, but the joint distribution

$$F(G, B). \tag{96}$$

5.2 Heterogeneous Agents and the Distribution of Task Breadth

Let B_i denote the equilibrium task breadth of agent i. Under heterogeneous general-purpose capability,

$$B_i = B(G_i, g_i, \lambda), \tag{97}$$

where g_i denotes the vector of task-specific comparative advantages.

The cross-sectional distribution of task breadth is therefore endogenous:

$$F_B(b) = \Pr(B_i \leq b). \tag{98}$$

An increase in general-purpose capability can affect this distribution along both its intensive and extensive margins.

On the intensive margin, existing agents perform more tasks:

$$\frac{\partial B_i}{\partial G_i} > 0. \tag{99}$$

On the extensive margin, more agents become sufficiently general to operate across broad task sets:

$$\frac{\partial}{\partial G} \Pr(B_i > \bar{B}) \geq 0. \tag{100}$$

The distinction allows for a gradual diffusion of generalization rather than an instantaneous economy-wide transition.

In particular, an economy can move through the following sequence:

$$\begin{array}{c} \text{few generalists} \\ \downarrow \\ \text{generalists coexist with specialists} \\ \downarrow \\ \text{broad task portfolios become common} \\ \downarrow \\ \text{predominantly generalized production.} \end{array} \tag{101}$$

The transition therefore need not be represented by a single representative agent crossing a threshold. It can instead occur through a changing distribution of productive capabilities.

5.3 Endogenous Coordination Costs

The baseline model treats coordination costs as technologically fixed. This assumption is useful for isolating the effect of task breadth, but it is restrictive for studying general-purpose technologies.

Artificial intelligence may itself reduce coordination costs. Systems capable of translating between specialized domains, maintaining shared representations, monitoring workflows, or automating communication can make it easier to coordinate a collection of specialized agents.

We therefore allow coordination costs to depend on technology:

$$C(N, G) = \kappa(G)\phi(N), \tag{102}$$

where

$$\kappa(G) \geq 0. \tag{103}$$

A coordination-enhancing technology satisfies

$$\kappa'(G) < 0. \tag{104}$$

General-purpose technological progress can therefore have two opposing effects on the organization of production.

The first is the *breadth effect*: agents become capable of performing more tasks themselves.

The second is the *coordination effect*: separate agents become cheaper to coordinate.

These effects push the organization of production in opposite directions.

5.3.1 Breadth versus Coordination

Let $\Delta_N(G)$ denote the gross productivity advantage of specialization over unification and let

$$K_N(G) = C(N, G) - C(1, G) \tag{105}$$

denote the incremental coordination cost of using N agents rather than one.

Specialization is optimal whenever

$$\Delta_N(G) > K_N(G), \tag{106}$$

while unification is optimal whenever

$$\Delta_N(G) \leq K_N(G). \tag{107}$$

Differentiating the difference between the two sides with respect to G gives

$$\frac{d}{dG}\left[\Delta_N(G) - K_N(G)\right] = \frac{\partial \Delta_N}{\partial G} - \frac{\partial K_N}{\partial G}. \tag{108}$$

The first term is negative when general-purpose capability erodes comparative advantage. The second term is also negative when AI lowers coordination costs, but it enters with a minus sign. Consequently, the coordination effect increases the attractiveness of specialization.

This produces the following result.

[Competing Effects of General-Purpose Technology] Suppose that an increase in G simultaneously:

1. reduces the specialization gain, so that
$$\frac{\partial \Delta_N}{\partial G} < 0,$$
and

2. reduces coordination costs, so that
$$\frac{\partial K_N}{\partial G} < 0.$$

Then the effect of G on the equilibrium degree of specialization is ambiguous. Specifically, specialization increases with G if

$$-\frac{\partial K_N}{\partial G} > -\frac{\partial \Delta_N}{\partial G}, \tag{109}$$

and decreases with G if

$$-\frac{\partial \Delta_N}{\partial G} > -\frac{\partial K_N}{\partial G}. \tag{110}$$

This result qualifies the central prediction of the paper in an economically useful way. A general-purpose technology does not mechanically imply less specialization. What matters is which of its organizational effects dominates.

If AI primarily expands the productive breadth of individual agents, generalization follows. If instead AI primarily reduces the costs of coordinating specialized agents, specialization can become more attractive.

This distinction suggests that the economic consequences of AI cannot be inferred from its raw task-level performance alone. One must also ask how the technology changes the organization of production.

5.4 Coordination as an Endogenous Choice

The previous analysis treats coordination costs as a reduced-form function of the number of agents. We now consider the possibility that agents can invest in coordination.

Let $x \geq 0$ denote coordination effort. The effective coordination cost is

$$C(N, x, G) = \kappa(G)\phi(N)e^{-x}, \tag{111}$$

while coordination effort itself has cost

$$D(x). \tag{112}$$

The planner chooses x jointly with the task allocation.

The first-order condition is

$$D'(x) = \kappa(G)\phi(N)e^{-x}. \tag{113}$$

An increase in G that lowers $\kappa(G)$ therefore reduces the marginal value of additional coordination effort. At the same time, however, lower coordination costs can increase the optimal number of agents.

This generates an additional feedback:

$$G \uparrow \Rightarrow \text{coordination becomes cheaper} \Rightarrow N^* \uparrow \Rightarrow \text{specialization becomes more attractive.} \tag{114}$$

The breadth effect generates the opposite feedback:

$$G \uparrow \Rightarrow \text{task breadth increases} \Rightarrow N^* \downarrow \Rightarrow \text{unification becomes more attractive.} \tag{115}$$

The net effect is therefore fundamentally an empirical and technological question about the relative elasticity of these two margins.

The model consequently yields a more general condition for the transition to generalization:

$$-\frac{\partial \Delta_N}{\partial G} > -\frac{\partial K_N}{\partial G}. \tag{116}$$

The left-hand side measures the rate at which general-purpose capability erodes the productivity advantage of specialization. The right-hand side measures the rate at which the same technology lowers the organizational cost of specialization.

5.5 Partial Generalization

We next relax the polar assumption that tasks are either fully specialized or fully unified.

Let a general-purpose agent perform a set of tasks $\mathcal{U}$ and let specialized agents perform the complementary set $\mathcal{S}$:

$$\mathcal{S} \cup \mathcal{U} = [0,1], \qquad \mathcal{S} \cap \mathcal{U} = . \tag{117}$$

Define

$$\theta = \mu(\mathcal{U}) \tag{118}$$

as before.

The planner chooses θ as well as the composition of each set. Let the productivity difference between the best specialist and the generalist in task j be

$$r(j,G) = a_S(j) - a_G(j;G). \tag{119}$$

The marginal value of assigning task j to a specialist is

$$r(j,G) - \lambda, \tag{120}$$

where λ is the marginal organizational cost of maintaining the additional specialization structure.

Hence

$$j \in \mathcal{S} \iff r(j,G) > \lambda. \tag{121}$$

The equilibrium generalized set is therefore

$$\mathcal{U}^*(G) = \{j : r(j,G) \leq \lambda\}. \tag{122}$$

This result has an important implication: partial generalization is not an ad hoc intermediate case. It is the generic outcome whenever residual comparative advantage differs across tasks.

[Interior Generalization] Suppose that the distribution of residual specialization advantages $r(j,G)$ has support on an interval containing λ. Then

$$0 < \theta^*(G) < 1.$$

Moreover, if $r(j, G)$ is strictly decreasing in G, then

$$\frac{d\theta^*(G)}{dG} > 0.$$

The proof follows from the task-level assignment condition and is provided in Appendix **??**.

The proposition is useful because it changes the interpretation of the transition. The economy does not generally move directly from Smithian specialization to complete unification. Instead, it progressively reallocates tasks across the organizational boundary.

Tasks with large residual comparative advantages remain specialized. Tasks for which general-purpose agents are nearly as productive as specialists become internalized.

5.6 Nested Specialization

Partial generalization can also produce nested organizational structures.

Suppose that tasks are divided into three classes:

$$\mathcal{T} = \mathcal{T}_H \cup \mathcal{T}_M \cup \mathcal{T}_L, \tag{123}$$

where

$$j \in \mathcal{T}_H \quad \text{if} \quad r(j, G) > \lambda_H, \tag{124}$$

$$j \in \mathcal{T}_M \quad \text{if} \quad \lambda_L < r(j, G) \leq \lambda_H, \tag{125}$$

$$j \in \mathcal{T}_L \quad \text{if} \quad r(j, G) \leq \lambda_L. \tag{126}$$

The highest-comparative-advantage tasks remain assigned to specialists. Intermediate tasks may be performed by a small number of broadly capable agents, while low-residual-advantage tasks are fully internalized.

This generates a hierarchy:

$$\text{specialists} \rightarrow \text{semi-generalists} \rightarrow \text{generalists}. \tag{127}$$

The hierarchy provides a natural interpretation of contemporary production systems in which some activities remain highly specialized while other activities are increasingly bundled within general-purpose workers or systems.

5.7 Endogenous Breadth and Specialization Quality

The baseline analysis treats task breadth as the relevant margin, but an agent may broaden its task portfolio at the cost of becoming less proficient at each task. We therefore consider an explicit breadth-productivity tradeoff.

Let agent i choose breadth b_i and let productivity in each assigned task be

$$a_i(b_i) = A_i h(b_i), \tag{128}$$

where

$$h'(b) < 0. \tag{129}$$

The function $h(b)$ captures the possibility that broader portfolios dilute attention, learning, or task-specific expertise.

An agent therefore faces a tradeoff:

$$\text{broader capability} \quad \leftrightarrow \quad \text{lower task-specific intensity.} \tag{130}$$

General-purpose technology can be represented as a shift in this tradeoff:

$$h'_G(b) > h'_{G'}(b) \qquad \text{for } G' > G, \tag{131}$$

so that technological progress makes breadth less costly.

The agent's optimal breadth solves

$$\max_{b_i} \left\{ A_i h_G(b_i) - \mathcal{C}(b_i) \right\}, \tag{132}$$

where $\mathcal{C}(b_i)$ represents the organizational or cognitive cost of maintaining a broader portfolio.

The first-order condition is

$$A_i h'_G(b_i) = \mathcal{C}'(b_i). \tag{133}$$

An increase in G that makes breadth less costly therefore raises optimal task breadth whenever the usual second-order condition holds.

This extension is important because it shows that generalization does not require assuming that a general-purpose agent is equally productive across all tasks. It is sufficient that technological progress reduce the productivity penalty associated with expanding the task portfolio.

5.8 Learning and Dynamic Generalization

The static model can also be extended to allow task breadth to generate future productivity.

Suppose that an agent accumulates knowledge from performing tasks. Let H_i denote the agent's accumulated productive knowledge and assume

$$\dot{H}_i = \psi(B_i) - \delta H_i, \tag{134}$$

where $\psi'(B) > 0$ and δ is the depreciation rate of knowledge.

Generalization can therefore generate a dynamic complementarity. A broader task portfolio produces more varied experience, which raises future productivity across tasks:

$$B_i \uparrow \Rightarrow H_i \uparrow \Rightarrow a_{ij} \uparrow . \tag{135}$$

In such an environment, a temporary increase in general-purpose capability can have persistent organizational consequences.

Conversely, specialization can generate its own dynamic advantage through learning-by-doing. Let specialist productivity satisfy

$$\dot{a}_{ij} = \chi(e_{ij}) - \delta_a a_{ij}, \qquad \chi' > 0. \tag{136}$$

The economy then contains competing dynamic forces:

$$\begin{aligned} \text{specialization} &\quad \rightarrow \quad \text{task-specific learning}, \\ \text{generalization} &\quad \rightarrow \quad \text{cross-task learning}. \end{aligned} \tag{137}$$

The resulting equilibrium can exhibit hysteresis: an economy that begins highly specialized may remain specialized even after general-purpose capability reaches a level at which a newly created generalist would otherwise be competitive. Conversely, early adoption of general-purpose technology can place the economy on a trajectory toward broader task portfolios.

A complete dynamic treatment is beyond the scope of the main text; Appendix **??** develops the corresponding two-state model and characterizes the conditions for multiple steady states.

5.9 Market Equilibrium

The baseline analysis uses the social planner's problem to isolate the technological mechanism. It is natural to ask whether the results survive decentralized market allocation.

Suppose that agents own their productive capabilities and firms purchase task services competitively. Let w_j denote the price of task j. A specialized agent chooses

$$\max_{S_i} \int_{S_i} w_j a_{ij} \, dj - C_i(S_i), \tag{138}$$

where $C_i(S_i)$ captures the private cost of maintaining a broader task portfolio.

A generalist enters a task whenever

$$w_j a_{ij} - C_i'(S_i) \geq w_j a_{kj} \tag{139}$$

for the relevant specialist k.

Under competitive task prices and transferable utility, the decentralized allocation reproduces the planner's task-assignment condition whenever coordination costs are internalized in the relevant organizational decision.

This observation is useful because it clarifies that the main mechanism is not an artifact of centralized allocation. The division of labor emerges from decentralized comparative advantage when task specialization is valuable; generalization emerges when broader capability makes internalization privately profitable.

The details of firm formation and bargaining over organizational rents are more naturally treated in a separate model of the firm. The main text therefore focuses on the efficient allocation, while Appendix **??** establishes sufficient conditions under which the same allocation can be decentralized.

5.10 The Boundary of the Firm

The preceding extensions provide a direct implication for theories of the firm.

Suppose that a firm can organize a set of tasks internally at cost

$$C_F(N, G), \tag{140}$$

while market exchange across independent agents incurs

$$C_M(N,G). \tag{141}$$

The organizational boundary is determined by

$$C_F(N,G)C_M(N,G). \tag{142}$$

A general-purpose technology can affect this inequality through both productivity and coordination.

If AI makes an individual agent capable of performing more tasks internally, it reduces the number of distinct contractual relationships required for a given production process. The boundary of the firm may therefore move inward in one sense: fewer agents need to be coordinated within a production chain.

But if AI simultaneously lowers the costs of coordinating independent agents, the boundary can move outward. The technology can make market-based specialization more viable.

Thus the theory does not imply a universal expansion or contraction of firms. It predicts that the organizational response depends on whether AI's breadth effect or coordination effect dominates.

5.11 Robustness to Alternative Production Technologies

The Leontief production function used in the baseline model emphasizes complementarity across tasks. The central results do not require this particular functional form.

Consider instead a CES aggregator,

$$Y = \left[\int_0^1 q_j^{\frac{\eta-1}{\eta}}\, dj\right]^{\frac{\eta}{\eta-1}}, \tag{143}$$

where η is the elasticity of substitution across tasks.

When η is finite, a productivity advantage in one task can partially compensate for lower productivity in another. The value of generalization therefore depends not only on task-level productivity but also on the substitutability of tasks.

For low η, tasks are strong complements and coordination across specialized agents is particularly important. Generalization can therefore generate large organizational gains.

For high η, tasks are more substitutable, reducing the need for every agent to perform every task and potentially weakening the force toward complete unification.

The central result survives provided that there remains a positive organizational cost to maintaining separate task assignments and that general-purpose capability reduces the productivity gap between generalists and specialists over a sufficiently broad task set.

[Robustness to Task Aggregation] Suppose that the production aggregator is monotone and differentiable in task outputs and that general-purpose capability weakly reduces the residual productivity advantage of specialists. Then the equilibrium degree of generalization is weakly increasing in general-purpose capability whenever the marginal organizational cost of task separation is positive.

The proof is provided in Appendix **??**. The result demonstrates that the mechanism does not depend on perfect complementarity among tasks.

5.12 What Generalization Does Not Imply

The extensions also help clarify several interpretations that the theory does *not* support.

First, generalization does not imply that every agent performs every task. Heterogeneous capabilities naturally generate an interior distribution of task breadth.

Second, generalization does not imply that specialists cease to have comparative advantage. They can retain large productivity advantages in particular tasks while becoming less important to the organization of production.

Third, generalization does not necessarily imply fewer jobs or lower aggregate employment. The model concerns the allocation of a given production process. In general equilibrium, higher productivity can expand the set of goods and services produced and create new tasks.

Fourth, generalization does not imply that firms disappear. It changes the technological tradeoff underlying organizational boundaries, but the equilibrium size and structure of firms depend on contracting, coordination, ownership, and market power.

Finally, generalization does not require that artificial intelligence be literally capable of performing every conceivable human activity. The relevant condition is economic rather than philosophical: whether a sufficiently broad set of economically relevant tasks can be performed by the same productive system at productivity close enough to the task-specific frontier that specialization gains no longer justify organizational separation.

5.13 Summary of the Extensions

The extensions establish four conclusions.

First, when agents are heterogeneous in general-purpose capability, generalization occurs first among the most capable agents and can diffuse gradually through the economy.

Second, when general-purpose technology also reduces coordination costs, the organizational consequences of AI become theoretically ambiguous. Breadth pushes toward unification, while cheaper coordination pushes toward specialization.

Third, partial generalization is a generic equilibrium outcome when residual comparative advantage differs across tasks. Complete unification is therefore a limiting case rather than a necessary prediction.

Fourth, the core mechanism survives alternative production technologies, endogenous breadth, learning, and decentralized allocation under standard regularity conditions.

The central comparative static can therefore be stated more generally than in the baseline model. Let $\mathcal{G}$ denote the breadth effect of a technology and $\mathcal{K}$ its coordination effect. Then the direction of organizational change depends on

$$\underbrace{-\frac{\partial \Delta}{\partial \mathcal{G}}}_{\text{erosion of specialization}} \quad \text{versus} \quad \underbrace{-\frac{\partial K}{\partial \mathcal{K}}}_{\text{reduction in coordination costs}} . \tag{144}$$

When the breadth effect dominates, the economy generalizes. When the coordination effect dominates, specialization can deepen. When the two effects are comparable, the equilibrium can exhibit coexistence between specialized and generalized production.

This distinction is particularly important for artificial intelligence. AI is not a single economic shock. It is a collection of technological changes affecting task productivity, task breadth, coordination, learning, and organizational design. The theory identifies task breadth as the margin capable of changing the fundamental logic of the division of labor.

The next section uses the framework to examine the implications of this organizational transition for firms, occupations, wages, and the allocation of productive rents.

6 Implications for Firms, Labor Markets, and Technological Change

The preceding sections characterize the transition from specialization to generalization as an equilibrium response to an expansion in the breadth of productive capability. This section considers the broader economic implications of that transition. The central observation is that the division of labor is not only a theory of occupations. It is simultaneously a theory of firms, markets, human capital, and technological change. Once the optimal allocation of tasks across agents changes, each of these objects changes with it.

Four implications follow. First, generalization changes the boundary of the firm by altering the relative cost of internalizing tasks and acquiring them through specialized exchange. Second, it changes the structure of occupations by reducing the optimal degree of task concentration and increasing the value of broad capabilities. Third, it changes the distribution of productive rents because comparative advantage becomes less valuable when general-purpose capability approaches the task-specific frontier. Fourth, it changes the direction in which technological progress is economically valuable: technologies that increase the breadth of productive capability can have organizational effects that task-specific innovations do not.

These implications distinguish the theory from a conventional model in which artificial intelligence is simply another productivity-enhancing technology. The central prediction is not that AI raises productivity. It is that sufficiently general-purpose AI changes the economic value of organizing production around specialized agents in the first place.

6.1 The Boundary of the Firm

The division of labor has long been inseparable from the theory of the firm. When tasks are performed by different agents, production requires some mechanism for coordinating their activities. That mechanism may be a market, a contract, or an organizational hierarchy. The optimal boundary of the firm therefore depends on the relative costs of internal coordination and market exchange.

The present framework introduces an additional margin. The number of distinct productive agents required to perform a given production process is itself endogenous to the breadth of individual productive capability.

Let N denote the number of agents involved in a production process. Let the cost of coordinating N agents within a firm be

$$C_F(N, G), \tag{145}$$

and let the cost of obtaining the corresponding task services through external markets be

$$C_M(N, G). \tag{146}$$

The firm internalizes a set of tasks whenever

$$C_F(N, G) < C_M(N, G). \tag{147}$$

In a conventional theory of the firm, technological change affects this inequality primarily through the costs of contracting, monitoring, and coordination. In the present model, there is an additional channel:

$$G \uparrow \quad \Rightarrow \quad N^*(G) \downarrow . \tag{148}$$

If general-purpose capability allows a single agent to perform tasks previously distributed across several agents, fewer productive relationships are required even before the firm chooses whether those relationships should be internal or external.

This creates a distinction between two forms of organizational change. The first concerns the *boundary* of an existing production organization:

$$\text{market} \quad \leftrightarrow \quad \text{firm}. \tag{149}$$

The second concerns the *number of productive agents* required to execute the underlying production process:

$$N^*(G). \tag{150}$$

The theory predicts that general-purpose capability directly affects the second margin and only indirectly affects the first.

6.1.1 Firm Size

Suppose that firm size is measured by the number of productive agents employed within the firm,

$$L_F. \tag{151}$$

If the production process requires $N^*(G)$ productive agents and the firm internalizes a fraction $\omega(G)$ of the associated tasks, then

$$L_F(G) = \omega(G)N^*(G). \tag{152}$$

Differentiating,

$$\frac{dL_F}{dG} = \omega'(G)N^*(G) + \omega(G)N^{*\prime}(G). \tag{153}$$

Since

$$N^{*\prime}(G) < 0 \tag{154}$$

under the breadth-dominant case, the effect on firm employment is ambiguous.

If general-purpose technology also reduces the cost of internal coordination relative to markets, then

$$\omega'(G) > 0, \tag{155}$$

and firms may become larger in organizational scope even as fewer productive agents are required for a given production process.

Conversely, if general-purpose technology makes external coordination cheaper, then

$$\omega'(G) < 0, \tag{156}$$

and firms may become smaller even as their remaining workers perform broader sets of tasks.

This yields a useful distinction:

$$\boxed{\text{firm size} \neq \text{task breadth} \neq \text{production complexity.}} \tag{157}$$

A theory of generalization therefore does not predict a simple relationship between AI adoption and firm employment.

6.1.2 The Declining Importance of Organizational Interfaces

There is, however, a sharper prediction.

Let I denote the number of interfaces between agents that must be coordinated in a production process. In a simple network representation,

$$I = I(N). \tag{158}$$

For a connected production network,

$$I(N) \geq N - 1, \tag{159}$$

with equality for a tree and larger values under richer interaction structures.

As general-purpose capability increases,

$$N^*(G) \downarrow \quad \Rightarrow \quad I^*(G) \downarrow. \tag{160}$$

Thus generalization reduces the number of organizational interfaces required to execute a fixed production process.

This is a distinctive implication of the theory. Even if aggregate output and the number of firms remain unchanged, the internal structure of production can become substantially less divided.

The economic significance is potentially large because many costs of specialization arise not from the tasks themselves but from the interfaces between tasks: handoffs, communication, monitoring, translation, scheduling, and reconciliation. A sufficiently general productive system internalizes some of these interfaces automatically.

6.2 Occupations and Task Bundling

The labor market provides the most direct empirical setting for the theory.

Let an occupation be represented by a set of tasks

$$\mathcal{O}_i \subseteq [0,1]. \tag{161}$$

Define occupational breadth as

$$B_i = \mu(\mathcal{O}_i). \tag{162}$$

In a specialized economy, occupations are narrow:

$$B_i \approx 0 \tag{163}$$

relative to the overall task space. In a generalized economy, occupations encompass a larger fraction of the production process.

The theory therefore predicts

$$\frac{\partial B_i}{\partial G} > 0 \tag{164}$$

for occupations whose tasks are sufficiently substitutable by general-purpose capability.

The relevant labor-market transformation is consequently not simply the disappearance of occupations. It is the re-bundling of tasks across occupations.

6.2.1 Occupational Recomposition

Suppose that an occupation initially consists of

$$\mathcal{O}_i = \{j_1, j_2, \ldots, j_k\}. \tag{165}$$

A general-purpose technology may make some subset

$$\mathcal{A}_i \subset \mathcal{O}_i \tag{166}$$

more efficiently performed by a generalist.

The occupation is then transformed into

$$\mathcal{O}_i' = \mathcal{O}_i \setminus \mathcal{A}_i \tag{167}$$

unless the released capacity is redeployed toward other tasks.

The latter possibility is important. If workers can expand their task portfolios, then automation of a subset of tasks does not necessarily imply a proportional reduction in employment.

Instead,

$$\text{automation of tasks} \rightarrow \text{task reallocation} \rightarrow \text{occupational recomposition.} \tag{168}$$

The theory therefore predicts that task-level exposure to AI is not sufficient to determine occupational outcomes. One must also know whether the remaining tasks can be bundled with other activities performed by the same worker.

6.2.2 The Declining Value of Narrow Comparative Advantage

Let the wage of worker i reflect the value of the worker's marginal contribution to production. A narrow specialist may earn a premium when

$$a_{ij} - a_{kj} \gg 0 \tag{169}$$

for some task j.

As general-purpose capability rises, the relevant productivity gap becomes

$$a_{ij} - a_G(j; G). \tag{170}$$

If

$$\lim_{G\to\infty} [a_{ij} - a_G(j;G)] = 0, \tag{171}$$

then the scarcity rent associated with narrow comparative advantage vanishes.

This does not imply that all wage inequality disappears. Workers may differ in general-purpose capability, ownership of complementary assets, or access to scarce inputs. But the source of rents changes.

Under specialization, rents are generated by

$$\text{task-specific scarcity.} \tag{172}$$

Under generalization, rents increasingly arise from

$$\text{breadth} + \text{complementary assets} + \text{organizational position.} \tag{173}$$

This is a fundamental redistribution in the source of productive advantage.

6.3 Human Capital

The model also changes the interpretation of human capital.

In a specialized economy, human capital investment is naturally directed toward task-specific productivity. Let h_j denote human capital specific to task j. The return to investment satisfies

$$R_j = \frac{\partial Y}{\partial h_j}. \tag{174}$$

When comparative advantage is strong, investment in h_j is highly productive because it increases the worker's advantage in a narrow task.

Under generalization, however, the value of human capital may increasingly depend on its transferability across tasks.

Let h^G denote general-purpose human capital and let

$$R_G = \frac{\partial Y}{\partial h^G}. \tag{175}$$

The transition toward generalization raises the relative return to broad capabilities whenever

$$\frac{\partial R_G}{\partial G} > \frac{\partial R_j}{\partial G} \tag{176}$$

for sufficiently broad sets of tasks.

This yields a prediction about education and training:

$$G\uparrow \quad \Rightarrow \quad \frac{R_G}{R_j}\uparrow. \tag{177}$$

The relevant distinction is not between cognitive and non-cognitive skills, nor between technical and nontechnical skills. It is between capabilities that are highly task-specific and capabilities that remain productive across a broad task space.

6.3.1 The Value of Learning to Learn

Suppose that an individual's productivity in a new task depends on a transferable learning capability ℓ_i:

$$a_{ij} = a_{ij}^0 + \ell_i q_j. \tag{178}$$

The value of ℓ_i increases with the number of tasks performed because its return is realized across the task portfolio:

$$\frac{\partial^2 Y}{\partial \ell_i \partial B_i} > 0. \tag{179}$$

Generalization therefore creates a complementarity between broad capability and learning.

This suggests that the economic value of education may shift toward capabilities that facilitate movement across task boundaries. Again, this is not a claim that specialized knowledge becomes worthless. Rather, the opportunity cost of acquiring extremely narrow knowledge rises when general-purpose systems can reproduce narrow task performance at low cost.

6.4 Wages and the Distribution of Rents

The model provides a simple decomposition of productive rents.

Let worker i's earnings be

$$w_i = \underbrace{R_i^{\text{spec}}}_{\text{specialization rent}} + \underbrace{R_i^{\text{gen}}}_{\text{generalization rent}} + \underbrace{R_i^{\text{asset}}}_{\text{asset rent}}. \tag{180}$$

The specialization component depends on the worker's comparative advantage in particular tasks. The generalization component depends on the breadth of tasks the worker can perform. The asset component captures ownership of complementary capital or technology.

As G increases,

$$R_i^{\text{spec}} \downarrow \tag{181}$$

for tasks in which general-purpose capability approaches the specialist frontier.

At the same time,

$$R_i^{\text{gen}} \uparrow \tag{182}$$

for workers whose general-purpose capability is sufficiently high.

The distributional consequence is therefore ambiguous in the aggregate but structured across dimensions of skill.

In particular, if workers differ in general-purpose capability,

$$G_i > G_k \quad \Rightarrow \quad B_i > B_k \tag{183}$$

and the return to general-purpose capability may increase with technological progress:

$$\frac{\partial w_i}{\partial G_i} > \frac{\partial w_k}{\partial G_k} \tag{184}$$

for sufficiently large differences in breadth.

This can generate a new form of inequality: not simply inequality between skilled and unskilled workers, but inequality between workers with different capacities to operate across task boundaries.

6.5 Technological Change as an Expansion of Task Breadth

The framework also suggests a broader taxonomy of technological progress.

Let technological change be represented by a vector

$$T = (T_1, T_2, \ldots, T_J), \tag{185}$$

where each component affects a subset of tasks.

A task-specific innovation increases

$$a_S(j) \tag{186}$$

for a particular j.

A general-purpose innovation instead increases

$$a_G(j) \tag{187}$$

across a broad set of j.

The distinction can be expressed using the Jacobian of productivity with respect to technology:

$$J_{j\ell} = \frac{\partial a_j}{\partial T_\ell}. \tag{188}$$

Task-specific technologies have a relatively sparse Jacobian. General-purpose technologies have a comparatively dense one.

But the organizational significance of a technology depends on more than the density of this Jacobian. It depends on whether the technology changes the *cross-task covariance* of productivity.

Define

$$\Gamma_{jk} = \mathrm{Cov}\left(a_{ij}, a_{ik}\right). \tag{189}$$

Under specialization, productivity advantages are highly task-specific. Under generalization, productivity becomes more positively correlated across tasks within an agent.

Thus one can characterize general-purpose technological change as a technological increase in cross-task capability:

$$G \uparrow \quad \Rightarrow \quad \Gamma_{jk} \uparrow. \tag{190}$$

This provides a formal interpretation of "general intelligence" that does not require a philosophical definition of intelligence. Economically, generality means that productive capability transfers across tasks.

6.6 Directed Technical Change

The framework also generates implications for the direction of innovation.

Suppose that innovators choose between a task-specific technology T_S and a general-purpose technology T_G. Let expected profits be

$$\Pi_S(T_S, G) \tag{191}$$

and

$$\Pi_G(T_G, G). \tag{192}$$

The direction of innovation is determined by

$$\Pi_G - \Pi_S. \tag{193}$$

When general-purpose capability is low, specialization creates large rents because task-specific comparative advantages are valuable. Innovation directed toward specialized technologies can therefore be attractive.

As G rises, however, the marginal value of improving an already general capability may increase because each improvement can be deployed across many tasks.

The return to general-purpose innovation can be approximated by

$$R_G = \int_0^1 \frac{\partial a_G(j;G)}{\partial G} \frac{\partial Y}{\partial a_G(j)} dj. \tag{194}$$

The return to a task-specific innovation in task j is

$$R_j = \frac{\partial a_S(j)}{\partial T_j} \frac{\partial Y}{\partial a_S(j)}. \tag{195}$$

A general-purpose technology becomes disproportionately valuable when

$$R_G > \max_j R_j. \tag{196}$$

The model therefore admits an endogenous transition in the direction of innovation itself.

6.7 Complementarity and the Pace of Generalization

General-purpose technology may also generate increasing returns through task complementarity.

Suppose that output is

$$Y = F(q_1, \ldots, q_J), \tag{197}$$

with

$$\frac{\partial^2 F}{\partial q_j \partial q_k} > 0 \tag{198}$$

for complementary tasks.

If general-purpose capability raises productivity across many tasks simultaneously, the marginal value of improving the technology in one task increases with its productivity in other tasks.

Thus

$$\frac{\partial^2 Y}{\partial G^2} > 0 \tag{199}$$

is possible even when the underlying task-level productivity function is approximately linear.

This creates a potential acceleration mechanism:

$$G \uparrow \to B \uparrow \to \text{more productive tasks internalized} \to \text{larger return to further } G. \tag{200}$$

The transition from specialization to generalization can therefore be gradual, but the rate of organizational change need not be constant.

6.8 General-Purpose Technology and the Speed of Structural Change

A conventional productivity shock can be absorbed without large organizational disruption when the mapping from agents to tasks remains approximately unchanged.

A general-purpose shock is different. If the equilibrium task allocation is

$$A^*(G), \tag{201}$$

then structural change is governed by

$$\frac{dA^*}{dG}. \tag{202}$$

The model predicts that this derivative is largest near the specialization-generalization boundary.

To see this, let the density of tasks near the assignment threshold be

$$f_R(\lambda; G). \tag{203}$$

From Section IV,

$$\frac{d\theta}{dG} = f_R(\lambda; G)\left|\frac{dr}{dG}\right| - f_R(\lambda; G)\lambda'(G), \tag{204}$$

up to the sign convention for the assignment boundary.

Thus economies with many tasks close to indifference between specialization and generalization experience the largest structural response to a common technological shock.

This creates a potentially important empirical prediction:

$$\boxed{\text{AI exposure} \times \text{task-level indifference} \Rightarrow \text{organizational change.}} \tag{205}$$

Task exposure alone is therefore insufficient. Two economies can have identical exposure to a general-purpose technology but experience very different organizational transitions because their task distributions differ.

6.9 Generalization and the Measurement of Productivity

The theory also has implications for productivity measurement.

Suppose that measured productivity is

$$\widehat{A} = \frac{Y}{X}, \tag{206}$$

where X is a conventional input aggregate.

If general-purpose technology causes agents to internalize tasks previously purchased from other agents or firms, some production that was previously recorded as an intermediate input may become an internal activity.

The measured allocation of inputs can therefore change even when the underlying production process changes less dramatically.

Moreover, conventional productivity statistics may understate the value of generalization when they do not capture reductions in coordination costs.

Let total production cost be

$$C = C_{\text{task}} + C_{\text{coord}}. \tag{207}$$

A general-purpose technology can raise measured task productivity while simultaneously reducing

$$C_{\text{coord}}. \tag{208}$$

The latter improvement may be poorly captured by conventional task-level productivity measures.

The theory therefore suggests that empirical evaluation of general-purpose technologies should distinguish

$$\text{task productivity} \quad \text{from} \quad \text{organizational productivity.} \tag{209}$$

The second includes the value created by reducing the number of productive interfaces and broadening the task portfolios of agents.

6.10 A New Margin of Technological Progress

The preceding results suggest a useful reinterpretation of technological progress.

Traditional growth theory focuses on technologies that increase the productivity of inputs. The present framework adds a distinct margin:

$$\text{technology} \rightarrow \text{task breadth.} \tag{210}$$

Let

$$\mathcal{B}(G) \tag{211}$$

denote the set of tasks that can be performed above a given productivity threshold by a single agent. General-purpose progress expands this set:

$$\frac{d\mu(\mathcal{B})}{dG} > 0. \tag{212}$$

This expansion changes the feasible organizational set.

In a specialized economy,

$$\mathcal{B}_i \cap \mathcal{B}_k \tag{213}$$

is relatively small for distinct agents i and k. In a generalized economy, these sets become larger:

$$\mu(\mathcal{B}_i \cap \mathcal{B}_k) \uparrow . \tag{214}$$

The technological transformation is therefore one from *complementary specialization* toward *capability overlap.*

This provides a more precise interpretation of the unification of labor. It does not mean that agents become identical. Rather, technological progress makes the productive capability sets of different agents increasingly overlap.

6.11 From Division of Labor to Unification of Labor

The theory can now be summarized in terms of the fundamental economic problem introduced at the beginning of the paper.

Under specialization, productive capability is distributed across agents:

$$\mathcal{B}_1, \mathcal{B}_2, \ldots, \mathcal{B}_N, \tag{215}$$

with relatively narrow capability sets.

The economic system gains from combining these capabilities because each agent possesses a comparative advantage in a subset of tasks.

Under generalization,

$$\mu(\mathcal{B}_i) \uparrow \tag{216}$$

and the overlap among capability sets increases.

As the general-purpose frontier approaches the task-specific frontier,

$$\mathcal{B}_i \to \mathcal{B}, \tag{217}$$

where $\mathcal{B}$ encompasses most economically relevant tasks.

The source of productivity then changes from the division of complementary capabilities across agents to the breadth of capabilities within agents.

This yields the central organizational transition:

$$\boxed{\begin{aligned} &\text{Division of labor:} && \text{agents specialize across tasks;} \\ &\text{Generalization:} && \text{agents internalize broader task sets;} \\ &\text{Unification of labor:} && \text{one productive system can span most relevant tasks.} \end{aligned}} \tag{218}$$

The third regime should be understood as a limiting case, not as an assertion that every economy will literally contain one agent. Even with arbitrarily general productive capability, scarcity of capital, physical inputs, information, ownership rights, and institutions can sustain multiple agents and firms. What disappears in the limiting case is not economic organization itself, but the *productivity-based necessity of dividing the production process into narrow human or artificial specialties.*

6.12 The Scope of the Theory

The theory therefore makes a deliberately narrower claim than a theory of an "AI economy" in general.

It does not predict the level of aggregate employment, the distribution of income, the rate of economic growth, or the equilibrium number of firms without additional assumptions. Those outcomes depend on demand, capital accumulation, market structure, innovation, institutions, and the creation of new tasks.

What the theory does predict is the response of the *organization of production* to an increase in general-purpose capability.

The fundamental sufficient condition is

$$\frac{\partial}{\partial G}\left[a_S(j) - a_G(j;G)\right] < 0 \tag{219}$$

over a sufficiently broad set of tasks, combined with a positive cost of maintaining separate task assignments.

Under these conditions,

$$G \uparrow \Rightarrow B^* \uparrow \Rightarrow N^* \downarrow \tag{220}$$

on the breadth-dominant margin.

The labor-market, firm-level, and technological implications are then consequences of this organizational transformation rather than independent assumptions.

6.13 Testable Predictions

The framework yields several empirical predictions.

Prediction 1: Task breadth should increase. Conditional on task composition, exposure to sufficiently general-purpose technology should increase the number of economically valuable tasks performed by a productive agent:

$$\frac{\partial B_i}{\partial G} > 0. \tag{221}$$

Prediction 2: Narrow comparative-advantage rents should decline. Workers whose productivity advantage is concentrated in tasks that become accessible to general-purpose systems should experience a decline in their task-specific scarcity rents:

$$\frac{\partial R_i^{\text{spec}}}{\partial G} < 0. \tag{222}$$

Prediction 3: Broad capabilities should become more valuable. The return to transferable capabilities should increase relative to the return to narrow task-specific capabilities:

$$\frac{\partial}{\partial G}\left(\frac{R_G}{R_j}\right) > 0. \tag{223}$$

Prediction 4: Organizational interfaces should decline. For a fixed production process, generalization should reduce the number of handoffs, coordination relationships, and distinct productive agents required:

$$\frac{\partial I^*}{\partial G} < 0. \tag{224}$$

Prediction 5: Organizational effects should be largest near the specialization boundary. The response of task allocation to general-purpose technology should be greatest where residual specialization advantages are concentrated near the organizational threshold.

Prediction 6: The effect on firm size is ambiguous. Generalization reduces the number of agents needed for a fixed production process, but changes in internal versus external organization can offset this effect. Hence

$$\frac{\partial L_F}{\partial G} \tag{225}$$

is theoretically ambiguous.

Prediction 7: AI can increase or decrease specialization. If the coordination effect dominates the breadth effect,

$$-\frac{\partial K}{\partial G} > -\frac{\partial \Delta}{\partial G}, \tag{226}$$

specialization may increase even as general-purpose capability improves.

These predictions provide a natural empirical agenda. Rather than asking only whether AI substitutes for particular occupations, one can ask whether AI changes the breadth of workers' task portfolios, the number of organizational interfaces, the distribution of comparative-advantage rents, and the boundaries between specialized and generalized production.

6.14 Section Conclusion

The implications of the model extend beyond the particular technology that motivates the analysis.

The division of labor is efficient when productive capability is sufficiently heterogeneous across agents and coordination costs are sufficiently low relative to the gains from exploiting comparative advantage. Generalization becomes efficient when technological progress makes productive capability sufficiently transferable across tasks.

The transition therefore changes several foundational economic objects simultaneously:

$$\begin{array}{ccc} \text{tasks} & \to & \text{broader task bundles,} \\ \text{occupations} & \to & \text{broader capabilities,} \\ \text{firms} & \to & \text{different organizational boundaries,} \\ \text{human capital} & \to & \text{greater value of transferability,} \\ \text{technological change} & \to & \text{expansion of capability breadth.} \end{array} \tag{227}$$

The deeper implication is that general-purpose technological progress can alter not merely the productivity of an economic system but the *unit around which production is organized*. When productive capability is narrow, the natural unit of production is the specialist. As capability becomes broader, the natural unit becomes the generalist. In the limiting case, the distinction between separate task specialists ceases to be economically fundamental.

The next section returns to the central theoretical result and places it in the context of the classical literature on the division of labor, comparative advantage, the theory of the firm, general-purpose technologies, and recent work on artificial intelligence.

7 Conclusion

The division of labor is one of the oldest organizing principles in economics. From Smith's account of specialization and the extent of the market to modern theories of comparative advantage, human capital, and the firm, a central premise of economic organization has been that productive capabilities are distributed across agents. Different agents perform different tasks because specialization allows the economy to exploit comparative advantage.

This paper asks what happens when that premise changes.

The starting point is deliberately simple. An economy consists of a set of productive tasks, and agents differ in the sets of tasks they can perform productively. When capability is sufficiently task-specific, specialization is efficient. When agents can perform broad sets of tasks at sufficiently high productivity, the gains from specialization decline. The optimal allocation of tasks therefore depends not only on the productivity of agents within tasks, but on the *breadth* of their productive capabilities across tasks.

This observation generates a theory of the transition from specialization to generalization.

The baseline model establishes that the division of labor is an equilibrium outcome rather than a technological primitive. Specialization emerges when comparative advantages are sufficiently large relative to the organizational costs of coordinating separate agents. As general-purpose capability increases, comparative-advantage gaps narrow and the equilibrium breadth of individual task portfolios expands. Above a threshold, the economy moves from specialization toward generalization; with heterogeneous tasks and agents, the generic outcome is partial generalization, in which some tasks remain specialized while others are internalized within broader productive systems.

The resulting framework nests the classical division of labor as one equilibrium and generalized production as another:

$$\boxed{\begin{array}{ccc} \text{specialization} & \longleftrightarrow & \text{generalization} \\ \text{narrow capability sets} & & \text{broad capability sets} \\ \text{comparative advantage} & & \text{capability breadth} \\ \text{many specialized agents} & & \text{fewer, broader productive agents.} \end{array}} \tag{228}$$

The distinction is conceptual as well as quantitative. Under specialization, the productive capabilities of different agents are complementary because each agent performs a different subset of tasks. Under generalization, productive capabilities increasingly overlap. The source of efficiency consequently shifts from combining narrow capabilities across agents toward deploying broad capabilities within agents.

This shift provides a different way of thinking about artificial intelligence.

Much of the economic analysis of AI naturally begins at the task level: whether a technology can perform a particular task, whether it substitutes for a particular occupation, or whether it complements a particular type of labor. These questions are important, but they take the mapping between agents and tasks as given. The theory developed here instead treats that mapping as endogenous.

The relevant technological variable is therefore not simply how productive an AI system is at any particular task. It is how broadly the same productive capability can be deployed across economically relevant tasks.

Let

$$B(G) \tag{229}$$

denote the breadth of the task set that can be performed by a productive system at a given technological frontier. The defining feature of a sufficiently general-purpose technology is

$$\frac{dB(G)}{dG} > 0. \tag{230}$$

When this increase in breadth is sufficiently large, the technology changes the optimal organization of production.

This distinction is important because a technology can raise productivity without fundamentally altering the division of labor. A machine that makes a specialist substantially more productive strengthens specialization if it does not increase the range of tasks that the specialist can perform. By contrast, a technology that allows the same productive system to perform many previously distinct tasks changes the economic value of specialization itself.

The difference can be represented schematically as

$$\begin{aligned} \text{task-specific innovation} &\rightarrow \text{higher productivity within specialization,} \\ \text{general-purpose innovation} &\rightarrow \text{higher productivity across tasks,} \\ \text{sufficiently broad innovation} &\rightarrow \text{change in the optimal division of labor.} \end{aligned} \tag{231}$$

The last step is the central object of the paper.

7.1 The Unification of Labor

The phrase "unification of labor" is intended to describe this limiting organizational transformation.

It does not mean that all economic agents become identical, that all production occurs within a single firm, or that specialization disappears in every dimension. Nor does it imply that every task can or should be performed by one system. Rather, unification refers to a reduction in the extent to which productive tasks must be divided across distinct agents because the same productive system can perform a sufficiently broad set of tasks.

In this sense, unification is the counterpart to the classical division of labor.

Under division of labor,

$$\mathcal{T} = \bigcup_{i=1}^{N} \mathcal{T}_i, \tag{232}$$

where each $\mathcal{T}_i$ is relatively narrow and the economy gains from assigning different subsets of $\mathcal{T}$ to different agents.

Under generalization,

$$\mu(\mathcal{T}_i) \uparrow, \tag{233}$$

and the overlap

$$\mu(\mathcal{T}_i \cap \mathcal{T}_k) \tag{234}$$

increases across agents.

In the limiting case,

$$\mathcal{T}_i \approx \mathcal{T}_k \approx \mathcal{T} \tag{235}$$

for the economically relevant task set.

The economic system has not ceased to organize production. It has changed the object being organized.

This distinction is central. The theory does not predict the disappearance of organization. It predicts a change in the technological basis of organization: from coordinating complementary specialists toward allocating and coordinating broad productive capabilities.

7.2 The Role of Coordination

The paper also emphasizes an important qualification.

General-purpose capability does not operate in isolation. The same technologies that increase task breadth may reduce the costs of coordinating specialized agents. If AI improves communication, monitoring, translation, scheduling, or information processing, it can make specialization more attractive even while increasing individual task breadth.

The organizational effect of AI is therefore determined by the relative strength of two forces:

$$\boxed{\text{breadth effect} \quad \text{versus} \quad \text{coordination effect.}} \tag{236}$$

The breadth effect reduces the productivity gains from specialization by allowing individual agents to internalize more tasks. The coordination effect reduces the organizational cost of specialization by making it easier to combine distinct agents.

This distinction prevents the theory from making the stronger and unnecessary claim that all improvements in artificial intelligence must reduce specialization. They need not.

The prediction is instead conditional:

$$\left|\frac{\partial \Delta}{\partial G}\right| > \left|\frac{\partial K}{\partial G}\right| \quad \Rightarrow \quad \text{generalization,} \tag{237}$$

where Δ denotes the productivity advantage of specialization and K denotes the organizational cost of maintaining it.

This condition identifies a measurable empirical object. The economic consequences of AI depend not only on its ability to perform tasks, but on how it changes the relative economics of *internalizing* versus *dividing* those tasks.

7.3 Implications for Economic Organization

The framework consequently has implications for several objects traditionally studied separately.

For firms, generalization can reduce the number of productive agents and organizational interfaces required to execute a given production process. Whether firms themselves become larger or smaller remains ambiguous because AI may simultaneously change the relative costs of internal organization and market exchange.

For labor markets, the key transformation is task recomposition. Occupations need not simply disappear when particular tasks become automatable. Workers can instead reallocate toward broader portfolios of tasks. The theory predicts an increasing relative value of capabilities that transfer across tasks and a declining value of narrow comparative advantage where general-purpose systems approach the specialist frontier.

For human capital, the relevant margin shifts from specialization alone toward the capacity to acquire, integrate, and deploy knowledge across domains. This does not make specialized knowledge irrelevant. It changes the opportunity cost of acquiring knowledge that has value in only a narrow region of the task space.

For technological change, the theory identifies task breadth as a distinct dimension of technological progress. An innovation that raises productivity in one task and an innovation that raises productivity across many tasks are not organizationally equivalent, even if their effects on measured productivity are initially similar.

Finally, for the measurement of productivity itself, the theory suggests that organizational gains may be as important as task-level gains. A technology can create value by eliminating handoffs, reducing coordination requirements, and allowing a smaller number of productive systems to perform a larger set of activities. These effects need not be fully captured by conventional measures of task productivity.

7.4 A Different Interpretation of General-Purpose Technology

The theory also suggests a refinement of the concept of a general-purpose technology.

A technology is conventionally described as general purpose when it has applications across many sectors or uses. The framework here identifies a different, though related, dimension of generality: the ability of the same productive system to perform many economically distinct tasks.

These two notions need not coincide.

A technology can be broadly applicable across industries while remaining highly task-specific within each industry. Conversely, a technology can have enormous organizational significance if it allows a single productive system to substitute for many specialized capabilities, even if its sectoral diffusion initially appears narrow.

The economically relevant measure for the present theory is therefore not simply the number of industries in which a technology is used, but the expansion of the task set available to a given productive system.

This suggests that future empirical work should distinguish

$$\text{sectoral generality} \quad \text{from} \quad \text{task generality}. \tag{238}$$

The latter is the margin that directly enters the theory of the division of labor.

7.5 The Broader Theoretical Implication

The deeper contribution of the framework is to treat the division of labor as a contingent equilibrium.

Economic theory has accumulated a large body of results explaining why agents specialize: comparative advantage, increasing returns to specialization, learning-by-doing, matching, coordination, and the extent of the market all contribute to the gains from dividing production across agents.

But these theories generally take sufficiently narrow productive capabilities as given.

The present framework asks what happens when the distribution of capability itself changes.

Once agents can perform broader sets of tasks, the economic question is no longer simply

$$\text{Who is best at task } j? \tag{239}$$

It becomes

$$\text{How many distinct agents are needed to perform the production process at all?} \tag{240}$$

That is a different question.

The answer depends on the shape of the technological frontier across tasks, the distribution of comparative advantage, the cost of coordinating agents, and the breadth of general-purpose capability. The division of labor is therefore one possible solution to an underlying allocation problem, not the allocation problem itself.

This perspective places the classical theory of specialization inside a broader theory of organizational breadth.

7.6 Limits and Directions for Future Research

Several important questions remain outside the scope of the present paper.

First, the model abstracts from the creation of entirely new tasks. If general-purpose technology generates new goods, services, and occupations, the task space itself becomes endogenous. A complete theory of technological change would therefore need to study the joint evolution of task creation and task generalization.

Second, the model does not fully characterize equilibrium market power. If general-purpose AI is owned by a small number of firms, the distribution of rents may differ substantially from the competitive benchmark.

Third, the dynamic extension could be developed further. Learning-by-doing, human-capital accumulation, technological adoption, and organizational adjustment costs may generate multiple equilibria, path dependence, or prolonged periods of coexistence between specialization and generalization.

Fourth, the analysis has treated general-purpose capability primarily as an increase in productive breadth. Future work could incorporate information acquisition, experimentation, and scientific discovery, where the ability to move across domains may itself create new forms of comparative advantage.

Finally, a complete welfare analysis would need to incorporate the endogenous creation of new tasks, capital accumulation, demand responses, and institutional adaptation. The present paper deliberately isolates the organizational mechanism through which technological breadth changes the division of labor.

These extensions are important precisely because the central mechanism is sufficiently general to survive them. The question is not whether specialization will literally disappear. The question is how far the economy moves along the spectrum between narrow specialization and broad productive capability.

7.7 Final Remarks

The history of economic development can be viewed, in part, as a history of increasingly sophisticated ways of combining human capabilities.

The division of labor was transformative because it allowed production to exploit differences among agents. It converted individual specialization into collective productivity. Its importance follows not from specialization being intrinsically desirable, but from the fact that productive capabilities were sufficiently heterogeneous and task-specific that separating activities created value.

A sufficiently general-purpose technology changes that condition.

If a single productive system can perform many tasks at productivity close to the best specialized alternative, then the economic rationale for dividing those tasks across separate agents weakens. The relevant unit of productive organization becomes broader. Tasks that once had to be coordinated across specialists can increasingly be performed within the same productive system.

The resulting transformation is not simply another increase in productivity. It is a change in the organization of production.

The classical question was:

$$\boxed{\text{Why divide labor?}} \tag{241}$$

The theory developed here adds its complement:

$$\boxed{\text{When should labor no longer be divided?}} \tag{242}$$

The answer is not that specialization ceases to matter. It is that specialization ceases to be the necessary organizing principle when productive capability becomes sufficiently general.

In that sense, the economic significance of artificial general intelligence may ultimately lie not only in what it can do, but in what it makes unnecessary.

The division of labor is an equilibrium of scarcity in productive breadth. If technological progress relaxes that scarcity, economics must allow for a different equilibrium—one in which the fundamental unit of production is no longer the specialist, but the general-purpose producer.

That possibility is the central subject of the theory of the unification of labor.

References

[1] Acemoglu, Daron, and Pascual Restrepo. 2018. "Artificial Intelligence, Automation and Work." NBER Working Paper No. 24196. Published in *The Economics of Artificial Intelligence: An Agenda*, edited by Ajay Agrawal, Joshua Gans, and Avi Goldfarb, Chicago: University of Chicago Press, 197–236, 2019.

[2] Acemoglu, Daron, and Pascual Restrepo. 2019. "Automation and New Tasks: How Technology Displaces and Reinstates Labor." *Journal of Economic Perspectives*, 33(2), 3–30.

[3] Autor, David H., Frank Levy, and Richard J. Murnane. 2003. "The Skill Content of Recent Technological Change: An Empirical Exploration." *Quarterly Journal of Economics*, 118(4), 1279–1333.

[4] Becker, Gary S., and Kevin M. Murphy. 1992. "The Division of Labor, Coordination Costs, and Knowledge." *Quarterly Journal of Economics*, 107(4), 1137–1160.

[5] Bresnahan, Timothy F., and Manuel Trajtenberg. 1995. "General Purpose Technologies 'Engines of Growth'?" *Journal of Econometrics*, 65(1), 83–108.

[6] Coase, Ronald H. 1937. "The Nature of the Firm." *Economica*, 4(16), 386–405.

[7] Gong, Jie, and I. P. L. Png. 2024. "Automation Enables Specialization: Field Evidence." *Management Science*, 70(3), 1580–1595.

[8] Grossman, Sanford J., and Oliver D. Hart. 1986. "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration." *Journal of Political Economy*, 94(4), 691–719.

[9] Grossman, Gene M., and Esteban Rossi-Hansberg. 2008. "Trading Tasks: A Simple Theory of Offshoring." *American Economic Review*, 98(5), 1978–1997.

[10] Hart, Oliver, and John Moore. 1990. "Property Rights and the Nature of the Firm." *Journal of Political Economy*, 98(6), 1119–1158.

[11] Smith, Adam. 1776. *An Inquiry into the Nature and Causes of the Wealth of Nations*. London: W. Strahan and T. Cadell.

[12] Williamson, Oliver E. 1975. *Markets and Hierarchies: Analysis and Antitrust Implications*. New York: Free Press.

[13] Williamson, Oliver E. 1985. *The Economic Institutions of Capitalism*. New York: Free Press.

Appendices

A Baseline Model and Equilibrium Characterization

This appendix provides the formal derivations underlying the baseline model. We first establish the planner's problem, characterize the optimal assignment of tasks to agents, and derive the conditions under which specialization and unification are optimal. We then establish the threshold characterization used throughout the paper.

A.1 The Production Environment

There is a unit mass of tasks indexed by

$$j \in [0,1].$$

A production process requires one unit of every task. Let q_j denote the quality-adjusted output of task j. Aggregate output is

$$Y = \min_{j\in[0,1]} q_j. \tag{243}$$

The Leontief specification is useful because it makes explicit that a production process is a collection of complementary activities. None of the results concerning the assignment of tasks requires the minimum operator itself; Appendix **??** establishes analogous results for a general class of monotone production aggregators.

There are N productive agents. Agent i has productivity

$$a_{ij}(G) = s_{ij} + \gamma(G), \tag{244}$$

in task j, where s_{ij} is task-specific comparative advantage and $\gamma(G)$ is a common general-purpose component satisfying

$$\gamma'(G) > 0. \tag{245}$$

The additive specification is not essential. What matters for the principal results is that increases in G compress productivity differences across agents. A multiplicative formulation is treated in Appendix **??**.

Let

$$a_j^*(G) = \max_i a_{ij}(G) \tag{246}$$

denote the task-level productivity frontier.

For any assignment $A : [0,1] \to \{1,\ldots,N\}$, output is determined by

$$q_j = a_{A(j),j}(G). \tag{247}$$

The cost of operating the assignment is

$$C(A,G) = \Phi(N_A, G), \tag{248}$$

where N_A denotes the number of distinct agents assigned a positive measure of tasks.

We assume throughout the baseline appendix that

$$\Phi_N(N,G) > 0, \qquad \Phi_{NN}(N,G) \geq 0. \tag{249}$$

The first condition states that maintaining a larger number of productive agents entails greater organizational cost. The second imposes weak convexity.

A.2 The Planner's Problem

The planner chooses an assignment A to maximize net output:

$$\max_A \left\{ \min_{j\in[0,1]} a_{A(j),j}(G) - \Phi(N_A, G) \right\}. \tag{250}$$

Because the Leontief aggregator implies that the least productive required task determines output, the planner's problem can be decomposed into two stages.

First, conditional on the set of active agents, the planner assigns each task to the most productive active agent. Second, the planner chooses the number and identity of active agents.

Let $\mathcal{I} \subseteq \{1,\ldots,N\}$ denote the set of active agents. Conditional on $\mathcal{I}$, the optimal task assignment is

$$A^*(j;\mathcal{I}) \in \arg\max_{i\in\mathcal{I}} a_{ij}(G). \tag{251}$$

This follows immediately because assigning task j to an agent with lower productivity cannot increase output and does not reduce the number of active agents.

Thus the planner's problem reduces to

$$\max_{\mathcal{I}\subseteq\{1,\ldots,N\}} \left\{ \min_j \max_{i\in\mathcal{I}} a_{ij}(G) - \Phi(|\mathcal{I}|, G) \right\}. \tag{252}$$

This representation makes clear that the organizational problem is fundamentally a problem of selecting the set of productive capabilities that will be active.

A.3 Task-Level Assignment

We first consider the case in which coordination costs are separable across additional agents. Let

$$\Phi(N, G) = \lambda(G)(N-1), \tag{253}$$

where $\lambda(G) > 0$ is the marginal cost of maintaining an additional productive agent.

For a given task j, suppose agent i is the best general-purpose producer and agent k is the best specialist. Define the residual specialization advantage

$$r_j(G) = a_{kj}(G) - a_{ij}(G). \tag{254}$$

The task is assigned to the specialist if and only if the productivity gain exceeds the organizational cost of maintaining the additional agent:

$$r_j(G) > \lambda(G). \tag{255}$$

Conversely, the task is assigned to the generalist if

$$r_j(G) \le \lambda(G). \tag{256}$$

This yields the basic task-assignment lemma.

[Task Assignment] Suppose that the marginal organizational cost of activating a specialist is $\lambda(G)$. A task j is optimally assigned to a specialist rather than a generalist if and only if

$$a_{Sj}(G) - a_{Gj}(G) > \lambda(G).$$

Proof. Assigning task j to the specialist rather than the generalist changes productive output by

$$a_{Sj}(G) - a_{Gj}(G).$$

It also requires the organizational cost $\lambda(G)$. The specialist assignment therefore increases net output if and only if

$$a_{Sj}(G) - a_{Gj}(G) - \lambda(G) > 0.$$

Rearranging yields the result. □

The importance of Lemma A.3 is that it converts the organizational problem into a task-level comparison. Generalization occurs not when the generalist becomes literally best at every task, but when the residual advantage of maintaining a specialist becomes smaller than the cost of specialization.

A.4 Complete Specialization

Consider the polar case in which each task has a distinct specialist and no agent can profitably internalize another agent's tasks. Let

$$\underline{r}(G) = \inf_j r_j(G). \tag{257}$$

Complete specialization is optimal whenever

$$\underline{r}(G) > \lambda(G). \tag{258}$$

[Condition for Complete Specialization] Suppose that every task has a specialist with residual productivity advantage $r_j(G)$ and that organizational costs are linear in the number of active agents. If

$$\inf_j r_j(G) > \lambda(G),$$

then every task is optimally assigned to its specialist.

Proof. By assumption,

$$r_j(G) > \lambda(G)$$

for every j. Lemma A.3 therefore implies that assigning every task to its specialist yields greater net output than assigning it to the generalist. Since this holds task by task, the globally optimal assignment is complete specialization. $\square$

A.5 Complete Generalization

Now let

$$\overline{r}(G) = \sup_j r_j(G). \tag{259}$$

If

$$\overline{r}(G) \leq \lambda(G), \tag{260}$$

then no task-specific specialist generates enough additional productivity to justify the organizational cost of maintaining a separate agent.

[Condition for Complete Generalization] Suppose that organizational costs are linear in the number of active agents. If

$$\sup_j r_j(G) \leq \lambda(G),$$

then the optimal allocation assigns all tasks to the general-purpose agent.

Proof. For every task j,

$$r_j(G) \leq \lambda(G).$$

By Lemma A.3, no task generates sufficient productivity gain from specialist assignment to compensate for the organizational cost of activating the specialist. Hence all tasks are optimally assigned to the generalist. $\square$

A.6 The Threshold

The preceding propositions imply a threshold characterization.

Suppose that $r_j(G)$ is continuous in G and weakly decreasing for every task:

$$\frac{\partial r_j(G)}{\partial G} \leq 0. \tag{261}$$

Suppose also that $\lambda(G)$ is weakly increasing or constant. Define the threshold

$$G_j^* = \inf \{G : r_j(G) \leq \lambda(G)\}. \tag{262}$$

Task j becomes generalized at G_j^*.

The economy-wide transition is governed by the distribution of these thresholds.

Let

$$F(G) = \mu\left(\{j : G_j^* \leq G\}\right). \tag{263}$$

Then $F(G)$ is the equilibrium measure of generalized tasks.

[Monotone Transition] Under the assumptions above, the measure of generalized tasks $F(G)$ is weakly increasing in G.

Proof. For every task j, $r_j(G)$ is weakly decreasing in G while $\lambda(G)$ is weakly increasing or constant. Hence once

$$r_j(G) \leq \lambda(G)$$

holds, it continues to hold for all larger G. Therefore the set

$$\{j : G_j^* \leq G\}$$

is weakly increasing in G, implying that its measure $F(G)$ is weakly increasing. □

This proposition formalizes the transition described in the main text. General-purpose technological progress does not necessarily produce an immediate discrete shift from specialization to unification. Instead, tasks cross their organizational thresholds at different levels of general-purpose capability.

A.7 Continuous Task Heterogeneity

Suppose that the residual specialization advantage can be represented by a random variable

$$R_G \tag{264}$$

with cumulative distribution function $F_R(r; G)$.

A task is generalized if

$$R_G \leq \lambda(G). \tag{265}$$

Hence the equilibrium share of generalized tasks is

$$\theta(G) = F_R(\lambda(G); G). \tag{266}$$

If F_R is differentiable, then

$$\theta'(G) = f_R(\lambda(G); G)\lambda'(G) + \frac{\partial F_R(\lambda(G); G)}{\partial G}. \tag{267}$$

When technological progress compresses residual specialization advantages,

$$\frac{\partial F_R(r; G)}{\partial G} > 0, \tag{268}$$

and coordination costs do not rise sufficiently rapidly, the share of generalized tasks increases.

The density

$$f_R(\lambda(G); G) \tag{269}$$

plays an especially important role. It measures the mass of tasks near the organizational boundary.

Thus the response to technological progress is largest when many tasks are nearly indifferent between specialization and generalization.

[Density at the Organizational Margin] Suppose $\lambda'(G) = 0$ and $\partial F_R(r; G)/\partial G > 0$. Then

$$\theta'(G) = \frac{\partial F_R(\lambda; G)}{\partial G}.$$

If the technological compression of residual comparative advantage is represented by

$$R_G = R_0 - \delta(G), \qquad \delta'(G) > 0,$$

then

$$\theta'(G) = f_{R_0}\left(\lambda + \delta(G)\right)\delta'(G).$$

Proof. The first expression follows directly from equation (267). Under

$$R_G = R_0 - \delta(G),$$

we have

$$F_R(r; G) = F_{R_0}(r + \delta(G)).$$

Therefore

$$\theta(G) = F_{R_0}(\lambda + \delta(G)).$$

Differentiating gives

$$\theta'(G) = f_{R_0}(\lambda + \delta(G))\,\delta'(G).$$

$\square$

This result provides the formal basis for the claim in Section VI that structural change should be concentrated near the specialization boundary.

A.8 Discrete Agents and the Number of Productive Systems

The preceding continuous formulation focuses on the share of tasks generalized. We now establish the corresponding result for the number of productive agents.

Suppose there are N potential agents and that agent i can perform a set of tasks $\mathcal{T}_i(G)$. Let the value of activating agent i be

$$V_i(G) = \int_{\mathcal{T}_i(G)} [a_{ij}(G) - a_{Gj}(G)]\, dj. \tag{270}$$

The agent is activated if

$$V_i(G) > \lambda(G). \tag{271}$$

If general-purpose capability reduces the productivity advantage of specialists, then

$$\frac{dV_i(G)}{dG} < 0. \tag{272}$$

Hence the number of active specialists is weakly decreasing.

[Number of Active Specialists] Suppose that $V_i'(G) < 0$ for every potential specialist and $\lambda'(G) \geq 0$. Then the equilibrium number of active specialists is weakly decreasing in G.

Proof. An agent is active whenever

$$V_i(G) > \lambda(G).$$

Since $V_i(G)$ is decreasing and $\lambda(G)$ is weakly increasing, the inequality can change only from true to false as G increases. Therefore the set of active specialists is weakly shrinking, and so is its cardinality. $\square$

A.9 Alternative Cost Functions

The linear organizational cost assumption is convenient but stronger than necessary.

Let

$$\Phi(N, G) \tag{273}$$

be any increasing differentiable function of the number of active agents.

Consider adding specialist i to an existing organizational structure containing N agents. The marginal cost is

$$\Delta\Phi(N, G) = \Phi(N + 1, G) - \Phi(N, G). \tag{274}$$

The specialist is activated if the productivity gain from the tasks assigned to that specialist exceeds this marginal cost.

Thus the relevant task-level condition becomes

$$\int_{\mathcal{T}_i} [a_{ij}(G) - a_{Gj}(G)]\, dj > \Delta\Phi(N, G). \tag{275}$$

The principal mechanism therefore survives nonlinear coordination costs.

What changes is that the value of specialization depends on the number of specialists already present. Under convex costs,

$$\Delta\Phi(N + 1, G) > \Delta\Phi(N, G), \tag{276}$$

so the economy may optimally retain only the specialists with the largest comparative advantages.

This generates an endogenous ranking of specialists.

A.10 A General Characterization

The baseline model can therefore be summarized by the following general problem.

Let S denote the set of specialized agents and let $\mathcal{T}_i$ denote the tasks assigned to specialist i. The planner solves

$$\max_{\{S, \mathcal{T}_i\}} \left\{ Y\left(\{\mathcal{T}_i\}_{i \in S}, G\right) - \Phi(|S| + 1, G) \right\}, \tag{277}$$

subject to

$$\bigcup_{i \in S} \mathcal{T}_i \subseteq [0,1]. \tag{278}$$

The remaining tasks are assigned to the general-purpose agent.

The division of labor is therefore optimal if and only if the marginal productivity gains generated by additional specialization exceed the corresponding marginal organizational costs.

This gives the paper's fundamental equilibrium condition:

$$\boxed{\text{specialization} \iff \text{marginal comparative advantage} > \text{marginal organizational cost.}} \tag{279}$$

General-purpose technological progress changes the first term by compressing comparative-advantage differences and may change the second by altering coordination technology.

The transition from division of labor to generalization is therefore an equilibrium response to a change in the relative magnitude of these two objects.

B Endogenous Productive Breadth

The baseline model treats the breadth of a productive system as technologically determined. This appendix endogenizes breadth itself. Agents choose how many tasks to internalize, trading off the benefits of additional task coverage against the costs of maintaining sufficiently broad capability.

B.1 Choice of Productive Breadth

Let an agent choose a task breadth

$$b \in [0,1]. \tag{280}$$

A breadth level b permits the agent to perform a subset of tasks of measure b.

Let the gross productivity benefit from breadth be

$$P(b, G), \tag{281}$$

with

$$P_b > 0, \qquad P_{bb} \leq 0. \tag{282}$$

The cost of acquiring and maintaining breadth is

$$K(b, G), \tag{283}$$

with

$$K_b > 0, \qquad K_{bb} \geq 0. \tag{284}$$

The agent solves

$$\max_{b \in [0,1]} \{P(b, G) - K(b, G)\}. \tag{285}$$

B.2 Interior Breadth

For an interior solution,

$$P_b(b, G) = K_b(b, G). \tag{286}$$

The second-order condition is

$$P_{bb}(b, G) - K_{bb}(b, G) < 0. \tag{287}$$

[Unique Optimal Breadth] If

$$P_{bb} \leq 0$$

and

$$K_{bb} \geq 0,$$

with at least one inequality strict, then the optimal breadth is unique.

Proof. The objective

$$P(b, G) - K(b, G)$$

is strictly concave under the stated conditions. A strictly concave function on the compact interval $[0, 1]$ has a unique maximizer. □

B.3 Technological Effect on Breadth

Differentiating equation (286) with respect to G yields

$$P_{bG} + P_{bb} b_G = K_{bG} + K_{bb} b_G. \tag{288}$$

Therefore

$$b_G = \frac{P_{bG} - K_{bG}}{K_{bb} - P_{bb}}. \tag{289}$$

The denominator is positive by the second-order condition. Hence

$$b_G > 0 \tag{290}$$

if and only if

$$P_{bG} > K_{bG}. \tag{291}$$

Thus technological progress expands optimal breadth whenever it raises the marginal return to breadth more than it raises the marginal cost of maintaining breadth.

[General-Purpose Technology and Endogenous Breadth] Under the second-order condition, general-purpose technological progress increases optimal productive breadth if and only if

$$P_{bG} > K_{bG}.$$

Proof. Immediate from equation (289). □

B.4 Technological Progress as a Reduction in Breadth Cost

A particularly useful special case is

$$K(b, G) = K_0(b) - \tau(G)b, \tag{292}$$

where

$$\tau'(G) > 0. \tag{293}$$

General-purpose technology therefore lowers the effective cost of acquiring breadth.

Then

$$K_{bG} = -\tau'(G) < 0. \tag{294}$$

Since

$$P_{bG} \geq 0, \tag{295}$$

it follows that

$$b_G > 0. \tag{296}$$

This gives a particularly transparent mechanism for generalization: technology increases the optimal breadth of productive systems by lowering the marginal cost of operating across tasks.

B.5 Technological Progress as a Breadth Productivity Increase

Alternatively, suppose

$$P(b, G) = P_0(b) + \alpha(G)b, \tag{297}$$

where

$$\alpha'(G) > 0. \tag{298}$$

Then

$$P_{bG} = \alpha'(G) > 0. \tag{299}$$

If breadth costs are technologically invariant,

$$K_{bG} = 0, \tag{300}$$

then

$$b_G > 0. \tag{301}$$

The two formulations are observationally similar in the static model, but they have different implications for technological diffusion and investment incentives.

B.6 The Boundary Solutions

The interior condition need not hold.

At

$$b = 0, \tag{302}$$

the agent chooses zero breadth if

$$P_b(0, G) \leq K_b(0, G). \tag{303}$$

At

$$b = 1, \tag{304}$$

the agent chooses complete breadth if

$$P_b(1, G) \geq K_b(1, G). \tag{305}$$

Thus the three organizational regimes are:

$$\begin{aligned} b^* = 0 \quad &\iff \quad P_b(0, G) \leq K_b(0, G), \\ 0 < b^* < 1 \quad &\iff \quad P_b(b^*, G) = K_b(b^*, G), \\ b^* = 1 \quad &\iff \quad P_b(1, G) \geq K_b(1, G). \end{aligned} \tag{306}$$

These correspond respectively to specialization, partial generalization, and complete generalization.

B.7 Connection to the Task-Assignment Model

The endogenous-breadth model can be embedded in the task-assignment model by defining

$$P(b, G) = \int_0^b \left[a_G(x, G) - a_S(x, G) \right] dx \tag{307}$$

for tasks ordered from easiest to hardest to internalize.

Then the marginal return to breadth is

$$P_b(b,G) = a_G(b,G) - a_S(b,G). \tag{308}$$

The optimality condition becomes

$$a_G(b,G) - a_S(b,G) = K_b(b,G). \tag{309}$$

This is precisely the task-level assignment condition expressed at the margin of the agent's portfolio.

Hence the endogenous-breadth model is not a separate mechanism. It is the continuous version of the task-assignment problem.

B.8 Breadth and Specialization as Dual Margins

Let total task mass be normalized to one. If the generalist internalizes breadth b, then the remaining task mass

$$1 - b \tag{310}$$

is available for specialized production.

Thus

$$\frac{d(1-b)}{dG} = -b_G. \tag{311}$$

Whenever

$$b_G > 0, \tag{312}$$

generalization and specialization move in opposite directions.

This gives a direct equivalence:

$$\boxed{\text{generalization} = \text{expansion of productive breadth} = \text{contraction of the specialized task set.}} \tag{313}$$

B.9 Multiple General-Purpose Agents

Now suppose there are M general-purpose agents, each choosing breadth b_i.

Let the production contribution of agent i be

$$P_i(b_i, G) \tag{314}$$

and let coordination costs depend on the total breadth deployed:

$$C = C\left(\sum_{i=1}^{M} b_i, M, G\right). \tag{315}$$

The planner solves

$$\max_{\{b_i\}} \left\{\sum_{i=1}^{M} P_i(b_i, G) - C\left(\sum_i b_i, M, G\right)\right\}. \tag{316}$$

The first-order condition for agent i is

$$P_{i,b} = C_B, \tag{317}$$

where

$$C_B = \frac{\partial C}{\partial B}, \qquad B = \sum_i b_i. \tag{318}$$

Thus all active general-purpose producers face the same marginal organizational cost.

If

$$P_{i,b} \tag{319}$$

differs across agents, breadth is allocated toward agents with the highest marginal return to breadth.

[Efficient Allocation of General-Purpose Breadth] Suppose production functions are concave in breadth and coordination costs are convex in aggregate breadth. In an interior optimum,

$$P_{i,b_i} = P_{k,b_k} = C_B$$

for all active general-purpose agents i, k.

Proof. The first-order conditions for any two active agents imply

$$P_{i,b_i} - C_B = 0$$

and

$$P_{k,b_k} - C_B = 0.$$

Hence

$$P_{i,b_i} = P_{k,b_k} = C_B.$$

Concavity ensures that the first-order conditions characterize the optimum. □

The proposition implies that technological generalization need not cause every productive system to become equally broad. Instead, breadth is allocated according to marginal productivity.

B.10 Strategic Breadth

If general-purpose agents choose breadth non-cooperatively, each agent takes the breadth of other agents as given.

Let the payoff of agent i be

$$\pi_i = P_i(b_i, G) - C_i(b_i, B_{-i}, G). \tag{320}$$

The Nash first-order condition is

$$P_{i,b} = C_{i,b}. \tag{321}$$

The planner instead internalizes

$$C_B = C_{i,b} + \sum_{k \neq i} C_{k,b_i}. \tag{322}$$

Therefore decentralized breadth may differ from the social optimum whenever agents impose coordination externalities on one another.

If

$$C_{k,b_i} > 0, \tag{323}$$

private agents may choose excessive breadth.

If breadth generates positive externalities, so that

$$\frac{\partial P_k}{\partial b_i} > 0, \tag{324}$$

private agents may instead choose insufficient breadth.

Thus the emergence of generalized production does not by itself imply that the decentralized equilibrium is efficient.

B.11 Learning and Endogenous Breadth

Suppose breadth generates learning:

$$\frac{\partial P}{\partial G} = P_G(b, G), \tag{325}$$

with

$$P_{bG} > 0. \tag{326}$$

Then greater technological capability raises the return to breadth, while greater breadth raises the return to technological capability.

This creates a complementarity:

$$G \uparrow \Rightarrow b \uparrow \Rightarrow P_G \uparrow . \tag{327}$$

Such complementarity can generate amplification in adoption.

If the adoption decision is itself endogenous, there may exist multiple equilibria in which an economy is either trapped in a specialized organizational structure or transitions toward generalized production.

B.12 Thresholds with Endogenous Breadth

Suppose

$$P_b(b, G) - K_b(b, G) \tag{328}$$

is strictly increasing in G and strictly decreasing in b.

Define G^0 by

$$P_b(0, G^0) = K_b(0, G^0) \tag{329}$$

and G^1 by

$$P_b(1, G^1) = K_b(1, G^1). \tag{330}$$

Then

$$G < G^0 \Rightarrow b^* = 0, \tag{331}$$

$$G^0 < G < G^1 \Rightarrow 0 < b^* < 1, \tag{332}$$

and

$$G > G^1 \Rightarrow b^* = 1. \tag{333}$$

Hence the transition can have three distinct stages:

$$\boxed{\text{specialization} \rightarrow \text{partial generalization} \rightarrow \text{unification.}} \tag{334}$$

This is the endogenous-breadth counterpart to the task-level threshold model.

B.13 The Unification Limit

Suppose

$$\lim_{G\to\infty} P_b(1, G) > \limsup_{G\to\infty} K_b(1, G). \tag{335}$$

Then complete breadth eventually becomes optimal.

[Sufficient Condition for Eventual Unification] If

$$\liminf_{G\to\infty} \left[P_b(1, G) - K_b(1, G)\right] > 0,$$

then there exists $\bar{G} < \infty$ such that

$$b^*(G) = 1$$

for all $G > \bar{G}$.

Proof. By assumption, there exists $\epsilon > 0$ and $\bar{G}$ such that for every $G > \bar{G}$,

$$P_b(1, G) - K_b(1, G) > \epsilon.$$

Therefore the marginal return to increasing breadth at the boundary exceeds its marginal cost. Since the objective is concave in breadth, the maximizer is the upper boundary:

$$b^*(G) = 1.$$

□

The proposition provides a precise sufficient condition for the limiting unification regime.

B.14 What Generalization Does Not Imply

The endogenous-breadth framework makes explicit three results that are sometimes conflated.

First,

$$\text{technological capability} \uparrow \tag{336}$$

does not imply

$$b^* = 1. \tag{337}$$

Second,

$$b^* \uparrow \tag{338}$$

does not imply

$$N^* \downarrow . \tag{339}$$

An economy can simultaneously have broader productive agents and more productive agents if the technology sufficiently increases the value of complementary activities.

Third,

$$b^* = 1 \tag{340}$$

does not imply that all agents have identical capabilities.

Unification concerns the breadth of economically relevant productive portfolios. It does not require homogeneity in the underlying technology or in agents' preferences, skills, or comparative advantages.

B.15 Summary

The endogenous-breadth extension establishes that productive breadth can itself be an equilibrium choice. General-purpose technological progress expands optimal breadth when it raises the marginal return to cross-task capability relative to the marginal cost of acquiring and maintaining it.

This yields a unified interpretation of the paper's central transition:

$$\boxed{\text{division of labor} \to \text{partial generalization} \to \text{unification of labor.}} \tag{341}$$

The transition is therefore not a binary technological event. It is an endogenous movement in the optimal breadth of productive systems.

C Robustness, Limiting Cases, and Additional Results

This appendix collects several robustness results. The purpose is to establish that the central distinction between specialization and generalization does not depend on particular functional forms, a representative-agent assumption, or the existence of a single general-purpose producer.

C.1 No Technological Advantage

Suppose general-purpose technology does not alter relative productivity:

$$a_{Gj}(G) = a_{Gj}. \tag{342}$$

Then technological progress in G has no direct effect on task assignment.

If coordination costs are also invariant,

$$\lambda_G = 0, \tag{343}$$

the equilibrium allocation is unchanged.

Thus generalization requires a change in the economic value of breadth rather than merely a relabeling of an existing technology.

C.2 Uniform Comparative Advantage

Suppose

$$a_{Gj}(G) - a_{Sj}(G) = R(G) \tag{344}$$

for all tasks.

Then there is no cross-task heterogeneity.

The equilibrium is necessarily a corner:

$$R(G) > \lambda(G) \Rightarrow \text{unification}, \tag{345}$$

and

$$R(G) < \lambda(G) \Rightarrow \text{specialization}. \tag{346}$$

Partial generalization therefore requires heterogeneity in task-level comparative advantage.

C.3 Arbitrarily Small Heterogeneity

Suppose task heterogeneity is arbitrarily small but nonzero.

Let

$$R_j(G) = \bar{R}(G) + \epsilon\xi_j, \qquad \epsilon > 0. \tag{347}$$

As

$$\epsilon \to 0, \tag{348}$$

the partial-generalization region becomes increasingly narrow.

Thus the theory approaches the homogeneous benchmark continuously.

C.4 Arbitrarily Large Heterogeneity

Conversely, let

$$\mathrm{Var}(R_j) \to \infty. \tag{349}$$

Then even very high levels of general-purpose capability may leave a nontrivial set of tasks for which specialization remains optimal.

Hence technological progress alone does not guarantee complete generalization unless it sufficiently compresses the upper tail of task-specific comparative advantage.

[Persistence of Specialization] If there exists a positive-measure set of tasks for which

$$\liminf_{G\to\infty} R_j(G) > \limsup_{G\to\infty} \lambda(G),$$

then specialization persists on a positive-measure set in the limit.

Proof. For every task in the stated set, the specialization inequality

$$R_j(G) > \lambda(G)$$

holds for sufficiently large G. Because the set has positive measure, the measure of specialized tasks remains positive in the limit. □

This proposition is important because it establishes the precise boundary of the unification result.

C.5 Bounded General-Purpose Capability

Suppose

$$\lim_{G\to\bar{G}} a_{Gj}(G) = \bar{a}_{Gj} < \infty. \tag{350}$$

Then complete unification occurs only if

$$\bar{a}_{Gj} - a_{Sj} > \bar{\lambda}_j \tag{351}$$

for almost every task.

If instead a positive-measure set satisfies

$$\bar{a}_{Gj} - a_{Sj} < \bar{\lambda}_j, \tag{352}$$

those tasks remain specialized.

Thus bounded capability naturally generates a mixed equilibrium.

C.6 Vanishing Coordination Costs

Suppose

$$\lambda(G) \to 0. \tag{353}$$

Then the assignment condition becomes asymptotically

$$a_{Gj}(G) > a_{Sj}(G). \tag{354}$$

If general-purpose productivity eventually exceeds specialist productivity for almost every task, unification follows.

If not, residual comparative advantage preserves specialization.

Hence the unification result can arise either from increasing general-purpose productivity or from falling coordination costs.

C.7 Increasing Coordination Costs

Suppose instead

$$\lambda(G) \to \infty. \tag{355}$$

Then specialization may survive even if general-purpose capability improves substantially.

In particular, if

$$\limsup_{G\to\infty} [a_{Gj}(G) - a_{Sj}(G)] < \liminf_{G\to\infty} \lambda(G) \tag{356}$$

for a positive-measure set of tasks, those tasks remain specialized.

This result emphasizes that productive breadth alone cannot overturn arbitrarily large coordination costs.

C.8 Discrete Tasks

The continuous task-space formulation is convenient but not essential.

Suppose there are J discrete tasks. Define

$$q_j(G) = \begin{cases} 1, & R_j(G) \le \lambda_j(G), \\ 0, & R_j(G) > \lambda_j(G). \end{cases} \tag{357}$$

The number of generalized tasks is

$$Q(G) = \sum_{j=1}^{J} q_j(G). \tag{358}$$

Generalization is then a sequence of discrete transitions.

If thresholds satisfy

$$G_1^* < G_2^* < \cdots < G_J^*, \tag{359}$$

the economy experiences

$$0 \to 1 \to 2 \to \cdots \to J \tag{360}$$

generalized tasks as G increases.

The continuous model can therefore be interpreted as the large-task limit of a discrete assignment model.

C.9 Stochastic Productivity

Suppose task productivity is stochastic:

$$a_{Gj} = \bar{a}_{Gj} + \varepsilon_{Gj}. \tag{361}$$

Let the expected payoff from generalization be

$$E[\Pi_{Gj}] \tag{362}$$

and that from specialization be

$$E[\Pi_{Sj}]. \tag{363}$$

The generalized assignment condition becomes

$$E[\Pi_{Gj}] > E[\Pi_{Sj}]. \tag{364}$$

Risk aversion introduces a variance term.

If utility is approximately mean-variance,

$$U(\Pi) \approx E[\Pi] - \frac{\gamma}{2}\,\mathrm{Var}(\Pi), \tag{365}$$

then generalization occurs when

$$E[\Pi_{Gj}] - \frac{\gamma}{2}\,\mathrm{Var}(\Pi_{Gj}) > E[\Pi_{Sj}] - \frac{\gamma}{2}\,\mathrm{Var}(\Pi_{Sj}). \tag{366}$$

Thus uncertainty can preserve specialization even when expected general-purpose productivity is higher.

C.10 Learning Curves

Suppose general-purpose capability improves through experience:

$$G_{t+1} = G_t + \alpha Y_t. \tag{367}$$

Then the production allocation itself affects future general-purpose capability.

If

$$\frac{\partial G_{t+1}}{\partial b_t} > 0, \tag{368}$$

generalization creates a dynamic externality.

An economy that generalizes early can accumulate general-purpose capability more rapidly than an otherwise identical economy that remains specialized.

This creates the possibility of endogenous divergence between economies.

C.11 Cross-Economy Divergence

Consider two economies, A and B, with

$$G_0^A > G_0^B. \tag{369}$$

Suppose

$$G_{t+1} = G_t + \alpha b_t \tag{370}$$

and

$$b_t = b^*(G_t), \qquad b^*_G > 0. \tag{371}$$

Then

$$G^A_{t+1} - G^B_{t+1} = G^A_t - G^B_t + \alpha \left[b^*(G^A_t) - b^*(G^B_t) \right]. \tag{372}$$

Because b^* is increasing,

$$b^*(G^A_t) > b^*(G^B_t), \tag{373}$$

so

$$G^A_{t+1} - G^B_{t+1} > G^A_t - G^B_t. \tag{374}$$

The initial technological advantage therefore grows over time.

[Generalization-Induced Divergence] If general-purpose capability and productive breadth are mutually reinforcing, then sufficiently small initial differences in G can generate increasing differences in long-run productive capability.

Proof. The result follows directly from the transition equation and the monotonicity of $b^*(G)$. □

This result provides a potential macro-development implication of the theory: economies may differ not only in the level of technology they possess, but in the organizational structures through which technology generates further technological capability.

C.12 Complementary Human Capabilities

Suppose generalized production requires a complementary human input h.

Output is

$$Y = F(B(G), h, G). \tag{375}$$

The marginal product of complementary human capability is

$$F_h. \tag{376}$$

If

$$F_{hG} > 0, \tag{377}$$

general-purpose technology increases the productivity of complementary human inputs.

Hence

$$G \uparrow \not\Rightarrow F_h \downarrow . \tag{378}$$

Instead, some human capabilities may become more valuable precisely because general-purpose production expands.

This provides a theoretical basis for distinguishing between tasks directly substituted by general-purpose systems and capabilities that become complements to them.

C.13 Generalization and Innovation

Let innovation produce new tasks at rate

$$\dot{N} = \Xi(G, B), \tag{379}$$

with

$$\Xi_G > 0, \qquad \Xi_B > 0. \tag{380}$$

Then greater general-purpose capability may increase the rate at which new productive tasks are created.

Consequently, even if

$$\mu(\mathcal{S}(G)) \downarrow, \tag{381}$$

the total number of tasks in the economy can increase.

This distinguishes the *share* of tasks performed by specialized labor from the *absolute number* of productive activities.

C.14 A No-Arbitrage Representation

Let the productivity frontier of a productive unit be

$$\mathcal{A}(G) = \left\{ \mathbf{a} \in R_+^J : \mathbf{a} \text{ is technologically feasible at } G \right\}. \tag{382}$$

Specialization corresponds to production vectors lying near coordinate axes of the frontier.

Generalization corresponds to production vectors with support over many coordinates.

Unification occurs when the frontier contains a productive vector with positive support over essentially the entire task space.

Thus the central transition can be represented geometrically as an expansion in the support of feasible production vectors:

$$\operatorname{supp}(\mathbf{a}) \uparrow \quad \text{as} \quad G \uparrow . \tag{383}$$

This formulation makes clear that the theory concerns the geometry of the production set rather than a particular labor-market institution.

C.15 Generalized Production as a Limiting Case

Let

$$\mathcal{A}_S \tag{384}$$

denote the production set generated by specialized technologies and

$$\mathcal{A}_G(G) \tag{385}$$

the production set generated by general-purpose technology.

Suppose

$$\mathcal{A}_S \subseteq \mathcal{A}_G(G) \tag{386}$$

for sufficiently large G.

Then general-purpose technology weakly expands the feasible production set.

If the inclusion is strict on a positive-measure set of task combinations, generalization creates strictly new productive possibilities.

In the limiting case,

$$\mathcal{A}_G(G) \to \mathcal{A}_\infty, \tag{387}$$

where $\mathcal{A}_\infty$ contains productive vectors spanning the full task space.

Unification is then the limiting organizational response to an expanded production set.

C.16 Summary of Robustness Results

The preceding results establish that the central mechanism survives:

1. alternative aggregate production functions;
2. discrete rather than continuous tasks;
3. stochastic productivity;
4. heterogeneous comparative advantage;
5. multiple general-purpose producers;
6. endogenous breadth;
7. endogenous coordination costs;
8. dynamic adoption;
9. learning and technological feedback;
10. decentralized competitive equilibrium.

The common condition is always an allocation margin of the form

$$\boxed{\text{marginal value of breadthmarginal value of specialization.}} \tag{388}$$

Smithian specialization emerges when the latter dominates. Generalization emerges when the former dominates. Unification is the limiting case in which the marginal value of further breadth remains positive through the entire task space.

C.17 Final Theoretical Statement

The central theoretical object of the paper can therefore be summarized by the generalized production problem

$$\max_{\mathcal{P}} \left\{ \int_{\mathcal{P}} [a_G(j,G) - a_S(j,G)]\, dj - C(\mathcal{P}, G) \right\}, \tag{389}$$

where $\mathcal{P}$ is the set of tasks assigned to a broad productive unit.

The first-order allocation condition is

$$a_G(j, G) - a_S(j, G)C_j(\mathcal{P}, G). \tag{390}$$

This condition contains both canonical regimes.

When

$$a_G(j, G) - a_S(j, G) < C_j, \tag{391}$$

the task remains specialized.

When

$$a_G(j, G) - a_S(j, G) > C_j, \tag{392}$$

the task is generalized.

The economy therefore moves continuously through the space of organizational forms as the relative productivity of breadth changes.

In the limiting case in which

$$a_G(j, G) - a_S(j, G) > C_j \tag{393}$$

for almost every task j,

$$\mu(\mathcal{P}) \to 1. \tag{394}$$

This is the unification limit.

The result does not abolish the logic of comparative advantage. Rather, it places comparative advantage inside a more general theory in which the productive unit's task set is endogenous.

Consequently, division of labor and unification of labor are not competing axioms. They are alternative equilibrium configurations of the same underlying production problem.