

The World Technology Frontier

By FRANCESCO CASELLI AND WILBUR JOHN COLEMAN II*

We study cross-country differences in the aggregate production function when skilled and unskilled labor are imperfect substitutes. We find that there is a skill bias in cross-country technology differences. Higher-income countries use skilled labor more efficiently than lower-income countries, while they use unskilled labor relatively and, possibly, absolutely less efficiently. We also propose a simple explanation for our findings: rich countries, which are skilled-labor abundant, choose technologies that are best suited to skilled workers; poor countries, which are unskilled-labor abundant, choose technologies more appropriate to unskilled workers. We discuss alternative explanations, such as capital-skill complementarity and differences in schooling quality. (JEL E13, E23, J31, O14)

An important question in macroeconomics is how to accurately describe the relationship between aggregate inputs and aggregate output—the aggregate production function—and how this relationship varies across countries. Currently, most research focuses on the model

$$(1) \quad y = k^\alpha (Ah)^{1-\alpha},$$

where y , k , and h are, respectively, output, physical capital, and human capital per worker. This aggregate production function is generally allowed to vary across countries via the total factor productivity (TFP) term $A^{1-\alpha}$. The typical finding is that TFP is higher in high-income countries.¹

* Caselli: Department of Economics, London School of Economics, London WC2A, United Kingdom (e-mail: f.caselli@lse.ac.uk); Coleman: Fuqua School of Business, Duke University, Box 90120, Durham, NC 27708 (e-mail: coleman@duke.edu). We are grateful to Daron Acemoglu, Ravi Bansal, Chris Foote, Jess Gaspar, Chad Jones, Pete Klenow, Sam Kortum, Aart Kraay, Per Krusell, Greg Mankiw, Philippe Martin, Pietro Peretto, Richard Rogerson, Lee Ohanian, James Schmitz, Nikola Spatafora, Silvana Tenreyro, Jaume Ventura, and Gianluca Violante for useful discussions and suggestions. We also thank Robert Barro and Jong-Wha-Lee for sharing their data on duration of schooling. Caselli thanks the University of Chicago Graduate School of Business, and Coleman thanks the Center for International Business and Economic Research, for financial support.

¹ The literature based on (1) is vast. Caselli (2005) presents a partial survey.

In constructing h , most of the work using production function (1) assumes that workers with different educational achievement (henceforth, skill level) are perfect substitutes in production. This assumption clashes with considerable evidence to the contrary. In particular, the empirical labor literature consistently documents elasticities of substitution between skilled and unskilled workers between 1 and 2, i.e., well short of infinity.² In addition, current practice tends to use only data on output and input quantities. But such variables do not exhaust the available sources of evidence that may be relevant in characterizing how the production function varies across countries: factor prices may also be informative.

In this paper we investigate the implications of relaxing the assumption of perfect substitutability of different types of labor, as well as of bringing to bear cross-country evidence on factor prices—particularly skill premia. This is achieved by generalizing (1) to a production function of the form:

$$(2) \quad y = k^\alpha [(A_u L_u)^\sigma + (A_s L_s)^\sigma]^{(1-\alpha)/\sigma},$$

where L_u is unskilled labor and L_s is skilled labor. Here, the labor input into production can be thought of as a constant-elasticity-of-

² See, among many others, the surveys of Daniel S. Hamermesh (1993) and Lawrence F. Katz and David H. Autor (1999).

substitution (CES) aggregate of unskilled and skilled labor in which the elasticity of substitution between skilled and unskilled labor is $1/(1 - \sigma)$. The two types of labor are imperfect substitutes as long as $\sigma < 1$ (the perfect-substitutability case is the special case where $\sigma = 1$). The parameters A_u and A_s convert raw quantities of the two labor types into efficiency units. In analogy to the standard practice of allowing A to vary across countries in (1), we allow A_u and A_s to vary across countries in (2). And in analogy to the practice of backing out A from (1), we present a simple methodology to back out each country's efficiency pair (A_u, A_s) when the production function is (2). The methodology uses data on output, factor inputs, and factor prices. All the results presented in this paper are based on the CES aggregate of labor types just described, but we also argue that our results are not driven by functional form assumptions.

In order to interpret cross-country differences in A_u and A_s , it is first useful to recall what such differences mean in a *cross-time* context. When A_s (A_u) increases over time, technical change is said to be skilled-labor (unskilled-labor) augmenting: the economy is becoming more efficient at using skilled (unskilled) workers. When the ratio A_s/A_u is constant over time, technical change is defined as skill neutral. Finally, when A_s/A_u increases (decreases) over time, technical change is skilled (unskilled) biased, and the economy is becoming relatively more efficient at using skilled (unskilled) labor.³ In order to adapt this time-series terminology to a cross-section of countries, we can replace the time index with an index of per capita income. We can then say that cross-country technology differences are skilled-labor (unskilled-labor) augmenting if A_s (A_u) tends to be higher in

higher-GDP countries, i.e., if richer countries use skilled labor (unskilled labor) more efficiently than poor countries. Further, cross-country technology differences are skill neutral if all countries are characterized by the same ratio A_s/A_u , and skilled biased (unskilled biased) if A_s/A_u tends to be higher (lower) in higher-GDP countries.

The central finding of this paper is that A_s and A_u do not move in lockstep across countries. While A_s rises steeply with y , the relationship between A_u and y is much weaker. Hence, the ratio A_s/A_u is systematically higher in rich countries, implying *skilled-biased cross-country technology differences*. This pattern of skill bias is an extremely robust result across definitions of "skilled," choices of calibrated parameter values, and alternative functional forms. Under our preferred set of assumptions, however, we also find some suggestion of a stronger form of bias: not only are A_s/A_u and A_s higher in rich countries, but A_u is actually absolutely lower in these countries. To distinguish between the weaker and the stronger version of the result, we refer to the tendency of A_s/A_u to be higher in rich countries as *relative skill bias*, and to the tendency of A_s to be higher and A_u to be lower in rich countries as *absolute skill bias*.

Our finding of skill bias suggests that cross-country differences in technology are not merely a matter of some countries having an overall higher level of technical efficiency than others, as assumed in most of the theories that aim to explain cross-country income differences. Rather, these theories may need to be enriched to account for the fact that poor countries seem to use certain factors relatively, and perhaps even absolutely more efficiently than rich ones. With that goal in mind, while uncovering the evidence of skill bias in cross-country technology differences is the main objective and contribution of the paper, we also sketch a possible theoretical explanation for our empirical finding.

Our suggested explanation for the cross-country pattern is best motivated by a simple example. Suppose that there exist two methods to produce output. One is with an assembly line where unskilled workers, supervised by a few skilled workers, wield hand tools; the other is with a computer-controlled and -operated facility that is mainly run by skilled workers and

³ More precisely, technical change is skilled biased if it increases the marginal productivity of skilled labor relative to unskilled labor. Under (2), the relative marginal productivity of skilled labor increases in A_s/A_u if $\sigma > 0$, i.e., if the elasticity of substitution is greater than one, and decreases in A_s/A_u if $\sigma < 0$. As we already mentioned, and as we further argue below, the range $\sigma > 0$ is the empirically relevant case. The definitions of factor augmenting, neutral, and biased technical change go back to John R. Hicks (1939). Their application to skilled and unskilled labor in the context of (2) is discussed, among others, by Autor et al. (1998), Katz and Autor (1999), and Daron Acemoglu (2002).

where unskilled workers play the role of janitors. Since the first method makes the most of unskilled workers, it seems fairly plausible that—faced with this choice—firms in unskilled-labor-abundant countries (which happen to be low-GDP countries) will tend to choose assembly-line production. Since the second method uses skilled workers more efficiently, firms in skilled-labor-abundant countries (i.e., high-income countries) will tend to choose the computerized facility.

To see how this example relates to our empirical findings, notice that firms are choosing between two possible production functions, say $f_1(K, L_u, L_s)$ and $f_2(K, L_u, L_s)$. Suppose that both f_1 and f_2 are as in (2), and what makes them two different production functions is that they have different parameters (A_u, A_s). In particular, the assembly-line production function, which uses unskilled labor relatively more efficiently, has low A_s/A_u , while the IT-based production function, which makes the most of skilled labor, has high A_s/A_u . Since poor countries choose the former and rich countries choose the latter, we therefore observe skill bias in cross-country technology differences.⁴ We present a simple model of endogenous technology choice that generalizes this example, checks the conditions under which this intuition works, and establishes when we should observe relative, and when absolute, skill bias. The model also shows how our evidence can be reconciled with the idea that some countries face barriers to technology adoption, and links our results to the literature on development accounting.

Having advanced one possible explanation for our empirical findings, we also consider alternative ones. We first discuss alternative functional forms of the production function, most notably ones allowing for capital-skill complementarity, which is not featured in our baseline specification. We argue that our results are not driven by functional-form assumptions. We then tackle the possibility that our results are driven by cross-country differences in the quality of schooling. We argue that our model of endogenous technology

choice provides a more plausible interpretation of the evidence.

Related Literature.—As is clear from the discussion above, our empirical result of a relative skill bias in cross-country technology differences has a time-series analog in a large body of evidence of skilled-bias technical change. This literature is comprehensively reviewed in, e.g., Katz and Autor (1999). A particularly strong connection exists with the paper by Katz and Kevin M. Murphy (1992), who use equation (3) below to estimate the *time trend* of A_s/A_u in the United States. To back out A_s/A_u , however, we follow a calibration approach, so that we do not need to impose structure on its pattern of variation across countries (i.e., we do not need to impose the analog of a time trend, such as a “GDP trend”). More important, with our methodology we go one step further and back out the actual levels of A_s and A_u . This allows us to investigate the possibility of absolute skill bias.⁵

Our proposed model of endogenous technology choice belongs primarily in the appropriate-technology literature, which goes back at least to Anthony B. Atkinson and Joseph E. Stiglitz (1969) (who called it “localized technology”), and has recently been further explored theoretically by Ishac Diwan and Dani Rodrik (1991), Susanto Basu and David N. Weil (1998), and Acemoglu and Fabrizio Zilibotti (2001). The key idea in this literature that is shared by our model is that countries with different factor endowments should choose different technologies. The Acemoglu and Zilibotti paper is particularly closely related in that it focuses on skilled and unskilled labor, as ours, in order to interpret patterns in cross-country data. However, a central result of their model is that A_s/A_u is constant across countries, which our evidence directly contradicts. On the empirical side, supportive evidence for appropriate technology has recently been developed by Caselli and Coleman (2001a) and Caselli and Daniel J. Wilson (2004), who found that cross-country diffusion of R&D-intensive technologies is strongly influenced by factor endowments.

⁴ Needless to say, our example is chosen to once again evoke the parallel with the skilled-biased technical-change literature. The adoption of IT-based production methods is the canonical source of increases in A_s/A_u over time.

⁵ Absolute skill biased in the U.S. time series has recently been documented by Marta Ruiz-Arranz (2002). See also Caselli and Coleman (2002).

Like all appropriate technology models, ours is also related to the literature on induced innovation/directed technical change, which studies the analogous problem of how factor endowments determine whether technical change will be biased toward certain factors rather than others. Important contributions in this tradition are Hicks (1932), Charles Kennedy (1964), Paul A. Samuelson (1965, 1966), Acemoglu (1998, 2002), and Charles I. Jones (2005). Formally, our model is closest to Samuelson's. Our argument that the cross-country skill bias we document is driven by endogenous technology choice dictated by skilled-labor endowments parallels Acemoglu's (1998) idea that skilled-biased technical change in recent years is driven by endogenous responses of R&D to changes in the relative supply of skilled labor.⁶

I. The Skill Bias in Cross-Country Technology Differences

When working with equation (1), one typically needs only solve for the unknown A . Our version of the exercise is slightly more complicated because equation (2) has two unknowns, A_u and A_s . We solve this problem by noting that, if factors of production are paid their marginal productivity, the skill premium is

$$(3) \quad \frac{w_s}{w_u} = \left(\frac{A_s}{A_u}\right)^\sigma \left(\frac{L_s}{L_u}\right)^{\sigma-1}.$$

The idea, then, is that (3) can be used as a second equation to combine with (2) to solve for the two unknowns.⁷ Hence, we back out each country's technology pair (A_u, A_s) so that measured inputs to production are exactly consistent

⁶ As mentioned, the model also makes contact with the literature on barriers to technology adoption. It is impossible to cite all, or even most, of the contributions in this vein. Some recent examples include Robert S. Barro and Xavier Sala-i-Martin (1997), Robert E. Hall and Jones (1999), Peter Howitt (2000), Stephen L. Parente and Edward C. Prescott (2000), Jonathan Eaton and Samuel S. Kortum (2001), Douglas Gollin et al. (2001), Philippe Aghion et al. (2005), Caselli and Nicola Gennaioli (2003), and Peter J. Klenow and Andrés Rodríguez-Clare (2006).

⁷ The closed-form solution is:

with measured output and skill premia.⁸ In order to execute this plan we need data on y , k , L_u , L_s , and w_s/w_u , as well as calibrated values for α and σ .⁹

A. The Data

Due to limitations in the availability of skill-premium data over time, we focus on a single cross-section of countries. Data for y and k for the year 1988 are obtained from Hall and Jones (1999); y is GDP per worker in international dollars (i.e., PPP adjusted); and k is an estimate of the real per-worker capital stock, obtained through a version of the perpetual-inventory method. The underlying data for both series come from Robert Summers and Alan Heston (1991).

Central to our exercise is the construction of the labor aggregates L_u and L_s , and the skill premia w_s/w_u . We build these variables up from three underlying datasets. The first dataset, from

$$A_u = \frac{y^{1/(1-\alpha)} k^{-\alpha/(1-\alpha)}}{L_u} \left(\frac{w_u L_u}{w_u L_u + w_s L_s} \right)^{1/\sigma},$$

$$A_s = \frac{y^{1/(1-\alpha)} k^{-\alpha/(1-\alpha)}}{L_s} \left(\frac{w_s L_s}{w_u L_u + w_s L_s} \right)^{1/\sigma}.$$

⁸ It is important for our methodology that relative wages are informative about relative marginal productivities. If developing countries had more egalitarian labor market institutions, the observed skill premium in these countries would underestimate the difference between the marginal productivity of skilled and unskilled labor, potentially leading to a spurious evidence of skill bias. Of course, however, it is well known that—if anything—social and political pressures for containing wage dispersion are much more severe in rich than in poor countries (with the possible exception of the United States), so, if anything, this type of measurement error biases the results against our finding of skill bias.

⁹ Our methodology is to allow A_u and A_s to vary across countries, while σ is constant, much as in the skilled-biased technical change literature. Needless to say, there is a certain amount of arbitrariness in the choice of which parameters vary, and which don't, across countries. This arbitrariness is inescapable: changes in σ cannot be separately identified from changes in A_s and A_u , as shown in the classic paper by Peter Diamond et al. (1978). It would, however, be possible to fix A_u , or A_s , or A_u/A_s , and let σ vary across countries. One would again be solving two equations in two unknowns, but one of the unknowns would now be σ . We leave the exploration of this alternative exercise for future work. See also John Duffy and Chris Papageorgiou (2000).

Barro and Jong-Wha Lee (2001), reports for each country the share of the labor force into each of seven categories of educational achievement: no education, some primary, completed primary, some secondary, completed secondary, some higher, and completed higher education. The second dataset, from Mark Bilal and Klenow (2000), reports each country's Mincerian coefficient, i.e., the coefficient on the number of years of education in a log-wage regression. The third dataset is an unpublished dataset by Barro and Lee which, for each country, reports the duration in years of primary and secondary schooling. Barro and Lee report attainment data at five-year intervals, so we pick 1985 to match the data on output and capital as close as possible.

In order to construct L_u and L_s , we must first decide which of the seven attainment subgroups to classify as "unskilled" and which as "skilled." For reasons discussed below, our preferred classification is that everyone who has completed a primary cycle of schooling is skilled, and those who have not are unskilled. Hence, L_u is a weighted sum of the first two subgroups, no education and some primary, while L_s is a weighted sum of the other five subgroups, from primary completed to completed higher education and above.

In order to identify the appropriate weight for each subgroup, we follow the standard convention according to which relative wages equal relative efficiency units. In particular, for each of the two aggregates, we choose the subgroup with least education as the "base" subgroup, and then weight all other subgroups by their wages relative to the base subgroup. Hence, for example, defining $L_{u,0}$ as the share of the labor force with no education, $L_{u,1}$ as the share of the labor force with only some primary education, and $w_{u,1}$ as the ratio of the wage of workers with some primary education to the wage of workers with no education, L_u is constructed as $L_{u,0} + w_{u,1}L_{u,1}$. Thus, L_u is measured in "no schooling equivalents." Similarly L_s is measured in "primary completed equivalents."¹⁰

In order to estimate the wages of the various

¹⁰ In other words L_u and L_s would sum to 100 (percent) if these two groups were constituted exclusively by workers at the respective "base" level of education (no education and primary completed, respectively).

subgroups relative to the base subgroup in each of the two labor aggregates, we use the Mincerian coefficients and the duration in years of the various schooling levels. From the duration of primary and secondary schooling we estimate the difference in years of schooling between different subgroups. For example, if secondary schooling takes five years, the difference in schooling years between workers who have completed secondary education (and not gone beyond) and workers who have completed primary schooling is five. Now the Mincerian coefficient is the percentage wage gain associated with an extra year spent in school, so that if β is the Mincerian rate of return, and n is the difference in schooling years between two workers, the ratio of their wages is $\exp(\beta n)$.¹¹

After completing the steps above, we have L_u and L_s in units of "no education" and "primary completed" equivalents. An additional correction is required, however, because in the data there is some (albeit minimal) cross-country variation in the duration of primary schooling. Hence, L_s is not fully comparable across countries, as the base worker may have different years of schooling (typically either four or five). In order to make L_s comparable across countries, therefore, we apply an additional rescaling that converts all workers in L_s into "four-years-of-schooling equivalents." In particular, if n_p is the duration of primary schooling, we multiply L_s in "primary completed equivalents" by $\exp[\beta(n_p - 4)]$.

The previous paragraphs describe the construction of labor aggregates based on a "primary-completed" definition of skilled. We also report results based on two alternative thresholds: completed secondary schooling and completed college. The construction of the labor aggregates and the skill premia for these alternative thresholds follows the same criteria as above. Hence, when we report results for the

¹¹ For subgroups with only partial completion of a certain education level (partial primary, partial secondary, or partial tertiary), we assume that they have completed exactly half of the overall duration of that course of study. So if primary schooling takes four years, workers with partial primary schooling have two years more schooling than their base group (no education). We do not have cross-country data on the duration in years of "higher education and above," so we assume that it lasts five years everywhere.

second definition of “skilled,” L_s is in “nine-years-of-schooling equivalents” (since across countries the modal number of years to complete secondary education is nine), and when we report results for the third threshold it is in “fourteen-years-of-schooling equivalents.” L_u is always in “no-schooling equivalents.”

Clearly there is no obvious way to establish a priori which of the three splits is the most empirically relevant. Workers within each of the two subaggregates are assumed to be perfect substitutes (though of course with different efficiency units), while workers across subaggregates are assumed to be imperfect substitutes. Heuristically, differences within groups are “quantitative”—some workers are more productive than others—but differences between groups are “qualitative”: some workers are fundamentally different. Reality is obviously much more nuanced, and drawing an arbitrary line to classify workers in these two categories is a subjective judgment. Having said that, our own intuition is that the definition of “skilled” based on primary schooling completed is the one that most closely captures this distinction. This definition roughly separates out the completely illiterate and innumerate from those who can at least read a simple text (e.g., a simple set of instructions or a newspaper article) and perform some basic arithmetic operations. We perceive this difference as qualitative: there are many tasks that no number of completely illiterate agents will be able to perform. Beyond the literacy threshold, most increases in education seem to us to have more of an incremental effect on skills, in the sense that most (though admittedly not all) production-relevant tasks that require literacy are accessible to all literate workers—though the less educated will need more time to perform them. Hence, the assumption that all workers who are at least literate are perfect substitutes is possibly more defensible than the assumption that the completely illiterate are perfectly substitutable with, say, those with some high-school education (but not with college).

The construction of the skill premia w_s/w_u is consistent with the construction of the labor aggregates. Hence, when defining skill as primary completed, the skill premium w_s/w_u is $\exp(\beta_4)$. When skill is defined as secondary completed, the skill premium is $\exp(\beta_9)$. And

TABLE 1—SUMMARY STATISTICS OF THE DATA

Variable	Mean	Std. dev.	Minimum	Maximum
y	13,506	9,717	1,854	35,439
k	32,271	28,991	1,218	107,870
L_s	89	41	30	229
L_u	61	28	6	115
w_s/w_u	1.50	.33	1.10	3.16

Correlation matrix

	$\log(y)$	$\log(k)$	$\log(L_s)$	$\log(L_u)$	$\log\left(\frac{w_s}{w_u}\right)$
$\log(y)$	1				
$\log(k)$	0.96	1			
$\log(L_s)$	0.62	0.66	1		
$\log(L_u)$	-0.74	-0.74	-0.66	1	
$\log(w_s/w_u)$	-0.38	-0.32	0.06	0.67	1

Notes: y and k are per-worker levels of real GDP and capital; L_s and L_u are supplies of skilled and unskilled labor; w_s/w_u is the skilled/unskilled wage premium.

when skill is defined by the completion of college, the skill premium is $\exp(\beta_{14})$.

There are 52 countries with complete data for y , k , L_u , L_s , and w_s/w_u (this dataset is reproduced in Appendix Table A.1). Table 1 reports some basic statistics from the dataset. For L_s , L_u , and w_s/w_u we report only the values corresponding to our preferred definition of skilled (alternative values are available on request). Output per worker in the richest country is 19 times higher than that in the poorest country. The supplies of skilled and unskilled workers also vary widely across countries (the implied ratio between L_s and L_u ranges from 0.32 to 36.11). The skilled wage premium ranges from 10 percent to 300 percent. Output is strongly positively correlated with both capital and the supply of skilled labor, while it is strongly negatively correlated with the supply of unskilled labor. As Bils and Klenow have documented, output is also negatively correlated with the skilled wage premium. Not surprisingly, then, the relative supply of skilled labor is negatively correlated with the skilled wage premium.

B. Calibration

In order to solve (2) and (3) for A_s and A_u we need to calibrate two parameters, α and σ . The parameter α measures, of course, the capital

TABLE 2—REGRESSION COEFFICIENTS OF A_s AND A_u ON y

$1/(1 - \sigma)$	Literacy			High school			College		
	A_s	A_u	diff	A_s	A_u	diff	A_s	A_u	diff
1.1	3.45	-5.26	8.71*	4.62	-1.13	5.75*	3.90	0.55	3.35*
1.4	1.41	-0.70	2.11*	1.62	0.33	1.29*	1.35	0.75	0.60
1.7	1.12	-0.05	1.17*	1.19	0.54	0.65*	0.99	0.78	0.21
2	1.00	0.21	0.79*	1.02	0.62	0.40*	0.84	0.78	0.06

Notes: The A_s column reports the coefficient of a regression of $\log(A_s)$ on $\log(y)$. The A_u column reports the coefficient of a regression of $\log(A_u)$ on $\log(y)$. The “diff” column reports the difference between the two coefficients. The symbol * indicates that this difference is statistically significantly different from zero.

share in GDP. For ease of comparability with previous results in the literature, we stick to the standard convention of setting $\alpha = 1/3$, which matches the U.S. historical value for this variable.¹²

The parameter σ is related to the elasticity of substitution between skilled and unskilled labor, $1/(1 - \sigma)$. This elasticity is the object of considerable focus in the labor-economics literature. After conducting their own review of the evidence, Autor et al. (1998) conclude that the elasticity of substitution is very unlikely to fall outside of the interval between 1 and 2. Hence, we experiment with a variety of values within this range.¹³

In the (1, 2) interval, the most popular estimate appears to be that of Katz and Murphy (1992), who set $1/(1 - \sigma)$ at 1.4. They arrive at this value by estimating equation (3) on U.S. time-series data between 1963 and 1997, with a time trend to control for changes in A_s/A_u . If deviations of A_s/A_u from the trend are not systematically related to changes in L_s/L_u , this seems a plausible approach to generating an estimate of σ . Accordingly, Katz and Murphy’s 1.4 will be our “preferred” value for $1/(1 - \sigma)$.¹⁴

¹² Recent cross-country estimates of the capital share by Gollin (2002) and Ben S. Bernanke and Refet S. Gürkaynak (2002) actually do show considerable cross-country variation, but this variation is not systematically related to income. It is unlikely, therefore, that setting a common value for this parameter will bias our results in any particular direction.

¹³ An ingenious recent addition to this literature is Antonio Ciccone and Giovanni Peri (2005), whose estimates of $1/(1 - \sigma)$ are well within the consensus bounds.

¹⁴ An important caveat is that the existing estimates of $1/(1 - \sigma)$ are based on datasets where skilled workers are identified with the college educated, which leads to a slight mismatch between some of our definitions of skilled and the

C. The Result

For each choice of labor aggregates and each choice of the parameter σ , we solve equations (2) and (3) for the two unknowns A_s and A_u . Table 2 reports the coefficients of regressions of $\log(A_s)$ on $\log(y)$ (first entry) and of $\log(A_u)$ on $\log(y)$ (second entry), implied by different choices of σ and different placements of the unskilled-skilled boundary. A “*” on the “diff” column indicates that the two slope coefficients are statistically significantly different from each other (at the 5-percent level). As is readily seen, in all cases the relation between A_s and y is stronger than the relation between A_u and y , in the sense that a 1-percent increase in y is typically accompanied by a larger percent increase in A_s than in A_u . This is our *relative skill bias* result. In 10 cases (out of 12), the difference between coefficients is economically huge. In nine cases it is also statistically significant.

In four cases, A_u actually declines with income, or we get *absolute skill bias*. As already noted, one of the cases in which we get absolute skill bias is our preferred case, where the skill threshold is literacy (or primary completed) and the elasticity of substitution is 1.4. Figures 1 and 2 show scatterplots against $\log(y)$ of $\log(A_s)$ and $\log(A_u)$, respectively, in this benchmark case. The negative association between A_u and y depicted in Figure 2 is statistically significant (P -value 0.012), and becomes more so if we drop the two seeming outliers, USA and Jamaica (P -value 0.007). If we omit the four richest and poorest countries, however, the re-

calibrated parameters. This is why we report results for a broad range of possible elasticities.




FIGURE 1. EFFICIENCY OF SKILLED LABOR

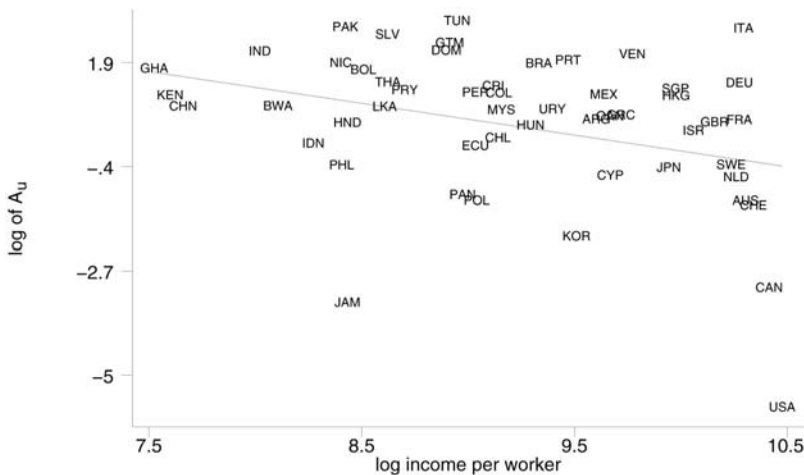


FIGURE 2. EFFICIENCY OF UNSKILLED LABOR

relationship becomes borderline insignificant (P -value 0.10). Because of this fragility, more conclusive support for the absolute bias property will have to await further work.

D. Deconstructing the Result

Before we move to a theoretical explanation for our finding of skill bias, it is worth digging a little deeper and assess what are the features of the data that drive our result. A glance at the two equations we are solving, equations (2) and

(3), immediately reveals that for each country the relative efficiency A_s/A_u is entirely pinned down by equation (3). In other words, A_s/A_u is chosen to fit the theoretical relationship between observed relative wages and observed relative labor supplies. We plot this theoretical relationship as the line “Model” in Figure 3 for a fixed choice of A_s/A_u . Clearly the relationship is negative, as an increase in relative employment of skilled labor leads to a fall in relative marginal productivities. Changes in A_s/A_u cause the line to shift: for example, a lower A_s/A_u implies lower skill premia for each value of




FIGURE 3. RELATION BETWEEN RELATIVE WAGE AND RELATIVE EMPLOYMENT

L_s/L_u , as skilled labor becomes relatively less productive. Changes in σ , i.e., changes in the elasticity of substitution, cause the line to tilt: for example, a higher elasticity of substitution implies that relative wages are less sensitive to changes in L_s/L_u .

Our main result is driven by the fact that—given empirically plausible values of σ —the relationship between w_s/w_u and L_s/L_u in the data is flatter than the theoretical relation, as summarized in the figure by the line “Data.” In order to reconcile the data with (3), then, low L_s/L_u countries must have lower A_s/A_u , as depicted by the line “Model’.” Recalling now that low L_s/L_u countries are also low-income countries, we have uncovered the sources of our relative skill bias result.

This discussion also allows us to understand why in Table 2 the skill bias becomes less pronounced when we increase the elasticity of substitution. As mentioned, as we increase the elasticity of substitution the theoretical relationship between L_s/L_u and w_s/w_u becomes flatter (the wage premium becomes less responsive to changes in relative employment), and hence closer to the empirical one. Hence, less of a shift (less of a difference in A_s/A_u) is required to match facts with theory.

So much for the relative bias. As for the absolute bias, it should be clear that—once equation (3) has dictated the value of A_s/A_u —the absolute levels are those necessary to fit the output equation, i.e., equation (2). In particular, note that we can rewrite equation (2) as

$$(4) \quad y = k^\alpha A_u^{1-\alpha} \left[L_u^\sigma + \left(\frac{A_s}{A_u} \right)^\sigma L_s^\sigma \right]^{(1-\alpha)/\sigma},$$

which pins down A_u for given A_s/A_u . As we just discussed, the wage and employment data call for a positive relation between A_s/A_u and L_s/L_u (and hence y), the more so the lower the elasticity of substitution, σ . The steeper the profile of A_s/A_u against y , the flatter the profile of A_u required to match the output data, y . When the A_s/A_u profile is at its steepest, as in our preferred case, matching the output data actually requires a declining profile for A_u .¹⁵

II. Explaining Our Findings: A Simple Model of Technology Choice

The previous subsection explains mechanically what features of the data give rise to the different patterns, vis-à-vis income per worker, in the observed efficiency units of skilled and unskilled labor. As such, the previous section does not provide an economic explanation for these patterns. The goal of this section is to sketch one possible explanation. Our explanation has its roots in the “appropriate-technology” tradition, which stresses that technology choice depends on factor endowment. To fully account for the patterns in the data, however, the idea of appropriate technology needs to be combined with the

¹⁵ While on the topic of reparametrizations of (2), we also note that our production function can be rewritten as:

$$y = A^{1-\alpha} k^\alpha \{ (\phi L_u)^\sigma + [(1-\phi)L_s]^\sigma \}^{(1-\alpha)/\sigma},$$

where ϕ is what is usually (mis)named a share parameter. Clearly, then, $A_s = (1-\phi)A$ and $A_u = \phi A$. In principle, therefore, one could read our evidence as indicating that “TFP” (A) is higher in rich countries, but that ϕ is much higher in poor countries, and the latter effect swamps the former, so that ϕA ends up being no higher in rich countries. It should be clear that this reinterpretation makes no substantive difference. In particular, one would still require the basic ingredients of the model in the next section to rationalize these findings: the higher A in rich countries would have to be explained by barriers to technology adoption (akin to the B term in the model below), while the higher ϕ in poor countries would continue to call for an appropriate-technology explanation. Similar observations would apply to any other reparametrization of (2).

idea of “barriers to technology adoption,” i.e., with cross-country differences in the overall ability of countries to absorb and implement technological improvements.

A. The Idea

As mentioned in the introduction, our proposed explanation is partly motivated by the recent literature on skilled-biased technical change. This literature has documented substantial increases in A_s/A_u over the last few decades in the United States and in a few other industrialized countries. A canonical example of skilled-biased technical change is the transition from an assembly line manned by unskilled workers, and supervised by a few skilled workers, to a computer-controlled facility operated by skilled workers, and where unskilled workers are at best retained as janitors (if not entirely displaced). In particular, the widely held view is that the shift from assembly-line-type methods to computer-based methods is strongly skilled-labor augmenting, i.e., it leads to a big increase in A_s . At the same time, since unskilled workers are demoted to janitorial roles, if not entirely displaced (to resurface elsewhere in menial jobs), it is plausible that the same shift leads to a decline in A_u . A decline in A_u over time is documented in Ruiz-Arranz (2002), and is consistent with the fact that absolute wages in the lower half of the wage distribution have actually declined in the United States over much of the last few decades.

Now the switch to the computer-controlled plant is of course a choice by the firm, since it could have decided to stick to the assembly line. But the fact that rich-country producers seem largely to have embraced the switch to computer-controlled production does not mean that firms in poor countries should necessarily make the same choice. In a country that is skilled-labor abundant, such as the United States, it makes sense to expect firms to adopt more skilled-biased technologies. But in countries abundant in unskilled labor, we may expect firms to stick to the old technology, and avoid the loss in the efficient use of the abundant factor. In this case, we will observe the cross-country skill bias we document: the skilled-abundant country will have

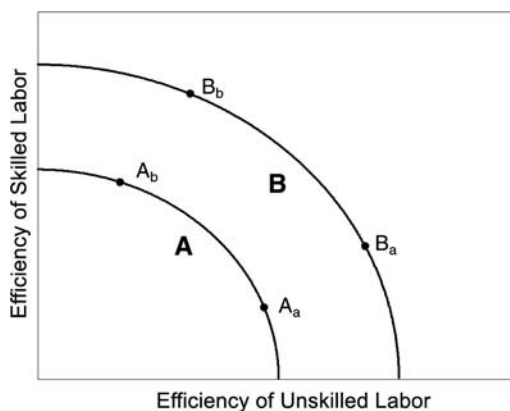


FIGURE 4. TECHNOLOGY CHOICE AND BARRIERS TO ADOPTION

Notes: A (B) is the technology frontier of country A (B). A_a and B_a (A_b and B_b) are appropriate choices of technology for unskilled-labor (skilled-labor) rich countries.

high A_s , and low A_u , relative to the unskilled-abundant country.

Our model generalizes this example simply by allowing a choice from a large number of technologies, instead of just the two of the example. The basic idea is that in each country firms choose from a menu of different production methods that differ in the use they make of skilled and unskilled labor. Each of these methods is a different production function. To capture the idea that different production functions use different inputs more or less efficiently, we assume that all production functions are of the form (2), but they differ in the parameters A_u and A_s . Hence, we can represent the menu of possible choices of production function by a set of possible (A_u, A_s) pairs. Clearly, no country will use a production function characterized by a certain pair (A_u, A_s) when another production function exists such that both A_u and A_s are higher, so only *nondominated* (A_u, A_s) pairs are relevant. We call this set of nondominated (A_u, A_s) pairs a “technology frontier.” We illustrate a possible frontier in Figure 4, where the axes measure the efficiencies of skilled labor and unskilled labor. The locus labelled A is the technology frontier for country A.¹⁶

¹⁶ Clearly, this reasoning relies on cross-country technology differences as being entirely characterized by differ-

The profit-maximizing choice of production function depends of course on factor prices. Since factor prices depend on factor endowments, firms in countries with different endowments will operate different production functions. If country A is unskilled-labor abundant, skilled labor will be relatively expensive, so we might expect firms in this country to choose a technology such as the one represented by point A_a , i.e., a relatively unskilled-complementary technology. If, instead, this country is skill abundant, firms may choose a technology such as A_b . In terms of the existing literature, A_a is an *appropriate technology* for an unskilled-labor-abundant country, while A_b is an appropriate technology for a skilled-labor-abundant country.¹⁷

Aside from the example we opened this subsection with, another way to motivate the idea of a technology frontier is suggested in a recent paper by Jones (2004). Jones argues that a new invention is essentially a draw from the distribution of possible (but yet uninvented) production functions. Suppose that production functions all have the functional form (2), but differ in the parameters A_u and A_s . Then a new idea—a newly invented production function—can be represented as a point in (A_u, A_s) space. Hence, technical change is nothing but the progressive “filling up” of the (A_u, A_s) space with newly available technologies. At any given point in time, firms will choose their production function from this set of feasible possibilities. Clearly, again, no country will choose a dominated technology, so only the subset of non-dominated production functions will be relevant. Such a set may look like a downward sloping curve in (A_u, A_s) space: a technology frontier.

An important question is how this appropriate-technology idea can be reconciled with the (more mainstream) view that poor countries face barriers to technology adoption. This is

important because the evidence on TFP differences is so compelling that one would not want to abandon the latter in order to embrace the former. In order to combine our appropriate-technology model with the “barriers” view of technology differences, we let the technology frontiers be country specific. The idea is that countries with more severe barriers face a more limited set of choices. In Figure 4 we illustrate this by drawing a separate frontier for country B . Since country B 's frontier is higher than country A 's, country B has fewer barriers to technology adoption. On its frontier, country B will choose B_a if it is unskilled-labor abundant, and B_b if it is skilled-labor abundant.

The following metaphor may be helpful in thinking about our framework. Suppose that in each country there is a library, containing blueprints, or recipes to turn inputs into output. Each blueprint is associated with a different realization of the efficiency vector. For example, there is a blueprint entitled “computer-controlled processing,” which leads to high skilled-labor efficiency and low unskilled-labor efficiency; and one called “assembly line” which is associated with an opposite pattern of efficiencies. The different country-specific frontiers can further be interpreted as library sizes. Some countries have just a handful of blueprints that fit on a short shelf, while some others have roomfuls of them.

It should be clear now how combining the “appropriate-technology” and the “barrier-to-adoption” ideas can rationalize our basic findings. Consider again the world of Figure 4, and imagine that country A is unskilled-labor abundant (and hence uses A_a) and country B is skill-labor abundant (uses B_b). If the frontiers are relatively close to each other, the appropriate-technology effect will dominate, and we will observe absolutely higher A_s in country B (the rich country), and absolutely higher A_u in country A (the poor country). In other words we will find absolute skill bias. This is the case depicted in the figure. If, instead, the frontiers are relatively far apart, the barriers effect will dominate, and A_s and A_u will be both higher in the rich country. In either case, however, the *ratio* of A_s to A_u is higher in the rich country, i.e., we always have relative skill bias.

ences in (A_u, A_s) . If we considered additional sources of heterogeneity in aggregate production functions, it would no longer be the case that the technology frontier had to be downward sloping in (A_u, A_s) space.

¹⁷ Also, as mentioned in the introduction, this reasoning is analogous to models of induced innovation/directed technical change, where an increase in the relative supply of skilled labor induces a skill bias in R&D activities (Acemoglu, 1998).

We conclude this discussion by noting that our framework implicitly defines a *world technology frontier*. This can be thought of as the “highest” frontier, or the frontier of a country that faces no barriers. It represents the set of nondominated (A_u, A_s) combinations dreamed up to date by scientists and management gurus, i.e., it reflects the current state of human technical knowledge. By introducing new technologies that dominate a subset of the pre-existing ones on the frontier, technological progress shifts this locus (locally) up.¹⁸

B. The Model

The following simple model formalizes the ideas set out in the previous subsection and establishes the conditions under which the intuition that countries will choose technologies that augment the abundant factor go through. We will see that the key parameter is the elasticity of substitution between skilled and unskilled labor.

We consider an economy with a large number of competitive firms. Each firm generates output using a production function of the form (2), which we reproduce here for convenience as equation (5):

$$(5) \quad y = k^\alpha [(A_u L_u)^\sigma + (A_s L_s)^\sigma]^{(1-\alpha)/\sigma}.$$

Firms hire the two labor types and capital, taking as given the rental rates w_u , w_s , and r . The novel element is that—besides optimally choosing factor inputs—firms also optimally choose the production function. In particular, they can choose from a menu of production functions that differ by the parameters A_u and A_s . The menu of feasible technology choices is given by:

$$(6) \quad (A_s)^\omega + \gamma(A_u)^\omega \leq B,$$

where ω , γ , and B , all strictly positive, are exogenous parameters. This says that, on the boundary of the feasible menu—on the technology frontier—changing production function involves a trade-off between the efficiency of unskilled labor, on the one hand, and the efficiency of skilled labor, on the other. The parameters γ and ω govern the trade-off; the parameter B determines the “height” of the technology frontier. The particular functional form of equation (6) is dictated by technical convenience, but it is rather flexible, and it does get at the central idea that there are trade-offs associated with technology choice.

In sum, in each country the representative firm maximizes profits $(y - w_u L_u - w_s L_s - rk)$ with respect to L_u , L_s , k , and A_u and A_s , subject to (5) and (6), the latter with equality. Here r is the firms’ cost of capital. We close the model by assuming that the economy’s endowments of k , L_u , and L_s are all inelastically supplied.¹⁹ An equilibrium is a situation where all firms maximize profits and all inputs are fully employed.

In the Appendix, we prove the following.

PROPOSITION: *An equilibrium exists and is unique. If $\omega > \sigma/(1 - \sigma)$ the equilibrium is symmetric, in the sense that all firms choose the same technology (A_u, A_s) , and the same factor ratios, L_s/k and L_u/k . If $\omega < \sigma/(1 - \sigma)$ the equilibrium is asymmetric, with some firms setting $A_u = 0$ and employing only skilled labor, and some others setting $A_s = 0$ and employing only unskilled labor.*

The proposition says that condition $\omega > \sigma/(1 - \sigma)$ is what is needed to rule out deviations from the symmetric equilibrium, deviations in which a firm chooses a corner with either $A_s = 0$ or $A_u = 0$.²⁰ Its meaning is rather

¹⁸ We do not take a stand (because we don’t need to) on two questions that are implicit in the foregoing discussion. First, we are agnostic about the determinants of the position of the world technology frontier in (A_u, A_s) space. Acemoglu and Zilibotti (2001) and Jones (2004) present two possible approaches to this question. Second, we are also agnostic on the sources of country-specific barriers to technology adoption. The literature exploring the possible culprits is huge, and growing, and we eschew a laundry list here. See footnote 6 for some citations.

¹⁹ None of the results we care about would change if we assumed that capital freely flows in and out of the country at some given world cost of capital r .

²⁰ Note that a symmetric equilibrium is always *interior*, in the sense that it features $A_s > 0$, $A_u > 0$. To see this, notice that a firm choosing $A_s = 0$ ($A_u = 0$) would also always choose $L_s = 0$ ($L_u = 0$). But then there must be some other firm making a different technology choice.

intuitive. When σ is low the two inputs are poor substitutes, and firms will want to operate production functions with positive quantities of both L_s and L_u . But if one is going to employ both inputs, it better be the case that the respective efficiency units A_s and A_u are strictly positive. As σ becomes larger, however, and L_s and L_u become better and better substitutes, it makes more and more sense to use only one of the inputs, and then maximize the efficiency of that input. For example a firm may choose to set $L_u = 0$ and then maximize A_s by also setting $A_u = 0$. The condition says that this will happen when σ becomes sufficiently large relative to ω ; ω regulates the concavity of the technology frontier—a higher ω pushes the frontier further away from the origin, i.e., it makes interior technology choices more attractive relative to the corners. Hence, it makes firms more reluctant to move to the corners. Notice that the condition for a symmetric equilibrium is always satisfied if $\sigma < 0$.

We now assume that the condition for existence of a symmetric equilibrium is satisfied, and examine this equilibrium's properties. Each firm's first-order conditions include

$$(7) \quad \left(\frac{L_s}{L_u}\right)^{1-\sigma} = \left(\frac{A_s}{A_u}\right)^\sigma \frac{w_s}{w_u},$$

$$(8) \quad \left(\frac{A_s}{A_u}\right)^{\omega-\sigma} = \gamma \left(\frac{L_s}{L_u}\right)^\sigma.$$

The first equation is of course just (3) rearranged. It combines the first-order conditions for L_u and L_s . It obviously says that the optimal choice of L_s/L_u is decreasing in w_s/w_u . For $\sigma > 0$ (good substitutability between skilled labor and unskilled labor) it also says that the greater the relative efficiency of L_s , the greater the desired relative employment of L_s . For $\sigma < 0$ (poor substitutability), L_s/L_u decreases in A_s/A_u , as the firm tries to boost the input of the inefficient (and hence effectively scarce) input.

The second equation is the first-order condition with respect to A_u . It describes how technology choice depends on the quantities of inputs employed. For $\sigma > 0$, the symmetric-equilibrium condition $\omega > \sigma/(1 - \sigma)$ implies $\omega - \sigma > 0$. Hence, equation (8) implies that

firms that employ a lot of skilled labor tend to choose technologies that augment skilled labor relative to unskilled labor. Conversely, if $\sigma < 0$, firms tend to direct technology choice toward the scarce input.²¹

Straightforward algebra combining the first two conditions leads to the following solution to the firm's problem:

$$(9) \quad \frac{A_s}{A_u} = \left(\frac{w_s}{w_u}\right)^{\sigma/(\omega\sigma - (\omega - \sigma))} \gamma^{(1-\sigma)/((\omega - \sigma) - \omega\sigma)}$$

$$(10) \quad \frac{L_s}{L_u} = \left(\frac{w_s}{w_u}\right)^{(\omega - \sigma)/(\omega\sigma - (\omega - \sigma))} \gamma^{\sigma/((\omega - \sigma) - \omega\sigma)}.$$

Of course the condition $\omega > \sigma/(1 - \sigma)$ can be rewritten as $\omega\sigma - (\omega - \sigma) < 0$. Hence, if $\sigma > 0$ firms increase the relative efficiency of the relatively cheap factor, while for $\sigma < 0$ firms focus on increasing the efficiency of the relatively expensive factor. Also, irrespective of σ , relative demand for skilled labor decreases in the relative skilled wage.²²

It is straightforward now to move from the firm's problem to the general equilibrium of the economy. Since the equilibrium is symmetric, equation (8) holds for L_s/L_u equal to the economy's endowment. Hence, with $\sigma > 0$ —i.e., when inputs are relatively good substitutes—countries with abundant unskilled labor will choose relatively unskilled-labor-augmenting technologies, while with $\sigma < 0$ —or when inputs are poor substitutes—countries with abundant unskilled labor will try to boost the productivity of skilled labor. In other words, when inputs are good substitutes, countries make the most of the abundant

²¹ There is, of course, also a first-order condition with respect to capital: $r = \alpha k^{\alpha-1} [(A_s L_s)^\sigma + (A_u L_u)^\sigma]^{(1-\alpha)/\sigma}$; but it plays no role in our subsequent analysis.

²² An additional insight on the properties of the firm's optimal technology choice is afforded by the equation:

$$w_u L_u = w_s L_s \gamma \left(\frac{A_u}{A_s}\right)^\omega,$$

which is a direct implication of combining (7) with (8). This relationship says that, in equilibrium, firms facing a relatively large skilled-labor share in output will tend to choose relatively skill-complementary technologies, or that firms will try to implement technologies that augment the factors that absorb a large share of income.

input, while when they are poor substitutes, it is optimal to increase the effective supply of the scarce factor. Now recall that empirically the elasticity of substitution $1/(1 - \sigma)$ is greater than 1, implying that $\sigma > 0$. Hence equation (8)—together with the fact that L_u/L_s is higher in poor countries—is therefore the rationalization of our basic finding or relative skill bias.

Indeed, if all countries shared the same technology frontier, i.e., if B were the same in all countries, it would follow directly from (8) and (6) that we should always find absolute skill bias: A_u should be absolutely higher in poor countries. However, the central message of the barriers-to-adoption literature is surely right: there are impediments to the diffusion of technology across countries. As already mentioned, one can nest this idea in the model by allowing the technology frontier in equation (6) to be country-specific. In particular, suppose that the height of the frontier, B , varies from country to country. It is straightforward to show that in this case one gets the relative version of skill bias—it's equation (8)!—without necessarily implying the absolute version. In particular, if B is much higher in rich countries, the absolute levels of both A_s and A_u will be higher in those countries. This can be seen formally by combining equations (8) and (6) to get:

$$A_s = \left(\frac{B}{1 + \gamma^{\sigma/(\sigma - \omega)} (L_s/L_u)^{\omega\sigma/(\sigma - \omega)}} \right)^{1/\omega}$$

$$A_u = \left(\frac{B/\gamma}{1 + \gamma^{\sigma/(\omega - \sigma)} (L_s/L_u)^{\omega\sigma/(\omega - \sigma)}} \right)^{1/\omega}.$$

Recalling that $\omega > \sigma$ is implied by our condition for an interior optimum, this says that A_s is increasing in both B and L_s/L_u , while A_u is increasing in B and decreasing in L_s/L_u (as long as $\sigma > 0$). We will soon explore some quantitative implications of our simple model. First, however, we briefly consider some alternative explanations.

III. Alternative Explanations

A. Alternative Functional Forms

The advantage of (2) is that it features only one parameter, σ , to capture the elasticity of substitu-

tion between skilled and unskilled labor. Moreover, a broad consensus exists in the literature on the likely magnitude (or at least a reasonably tight range) of this parameter. Hence, as we have seen, equation (2) is relatively easy to calibrate. Nevertheless, we also argue that our basic message would go through under alternative specifications for the aggregate production function.

Cobb-Douglas.—Since our stated aim was to explore the implications of aggregate production functions where skilled and unskilled labor are allowed to be imperfect substitutes, it would have been possible for us, instead of working with (2), to study the production function

$$(11) \quad y = Ak^\alpha L_u^\beta L_s^{1 - \alpha - \beta}.$$

In particular, we could have allowed technology to vary across countries through variation in A and in β (with α always fixed at 0.33, as in our current exercise). Each country's A and β would again be chosen to simultaneously fit the production function (11) and the expression for the skill premium equivalent to (3).

The important observation to make about this possible alternative exercise is that it is based on a counterfactual choice of the elasticity of substitution between skilled and unskilled labor. Equation (11) imposes this elasticity to be 1. As discussed above, empirical evidence puts it in the neighborhood of 1.4. Our model (2) is thus consistent with empirical evidence in a way that (11) is not.

With that caveat, when we perform the exercise described above, we find high β and low A in poor countries, and low β and high A in rich countries.²³ This is substantially the same result that obtained with our CES specification. Notice, in particular, that it still involves exactly the same trade-off between the contributions to output of unskilled and skilled labor: higher β gives more weight to unskilled labor and less to skilled labor. Ultimately, then, one still ends

²³ These patterns are very marked for the "primary" and "secondary" definition of skill, while the negative association between β and y is quite weak in the case of "college." Note that in the case of college we are particularly confident that the elasticity of substitution is greater than 1.

up writing an endogenous technology-choice model where unskilled-abundant countries choose unskilled-complementary technologies (high β) and skilled-abundant countries choose skilled-complementary technologies (low β). One cannot use the language of trading off efficiency units of skilled versus unskilled labor because by forcing the production function to be Cobb-Douglas there is no such thing as a separate measure of efficiency units. But the *economics* of what is going on is the same and there is no substantive difference in message.

Capital-Skill Complementarity.—Another alternative choice of functional form could have been the two-level CES

$$(12) \quad y = \{(A_u L_u)^\sigma + [(A_s L_s)^\rho + (A_k k)^\rho]^{\sigma/\rho}\}^{1/\sigma}.$$

As recently emphasized by Per Krusell et al. (2000), the potential advantage of this functional form is to allow for a version of capital-skill complementarity. In particular, if $\sigma > \rho$, an increase in the supply of physical capital increases the skill premium. Readers of Krusell et al. may wonder whether our finding that A_s/A_u is higher in high-income countries is driven by not having taken into account this capital-skill complementarity effect.

We used (12) to perform an exercise similar to that performed in this paper (results available upon request). In particular, we backed out not only A_u and A_s , but also A_k . This required a third equation, besides the production function and the skill premium. We used an international no-arbitrage condition on the return to capital. We experimented with a wide range of values for σ and for ρ , finding overwhelming evidence of nonneutrality and skill bias. For example, when using the Krusell et al. estimates of the parameters we found that A_s and A_k are higher in rich countries, and A_u is higher in poor countries. This is the same result we presented here. Note that the Krusell et al. parameters imply capital-skill complementarity, so it is clearly

not the case that our results are driven by the omission to account for capital-skill complementarity. In sum, even accounting for possible capital-skill complementarity we obtain the skill bias result.

B. *Quality of Schooling*

Another potential concern is that A_s/A_u is picking up differences in the quality of schooling across countries. This is best illustrated by referring to the production function as rewritten in (4). Since workers classified as unskilled are largely unschooled, it is plausible to assume that they are homogeneous across countries. Instead, workers included in L_s , which is based on the quantity of schooling, could still be heterogeneous if the quality of schooling differed systematically across countries. In this case, the term A_s/A_u would pick up such quality differences, and the result of a relative skill bias would really be reflecting higher quality of schooling in rich countries.

While we acknowledge this possibility, we think it implausible that unmeasured schooling quality explains more than a small fraction of the patterns we uncover. Caselli (2005) examines in detail the question of how much of the unexplained variation in income can be attributed to schooling quality. He looks at data on pupil/teacher ratios, educational expenditures, test scores, etc., and tries to calibrate how variations in these measures can increase the cross-country variance of human capital (and thereby reduce the unexplained component of the variance in GDP). The conclusion is that, to register even a small effect, one needs elasticities of human capital to these measures of schooling quality that are vastly larger than those in the empirical labor literature. Admittedly, there may be other dimensions of schooling quality not captured by these measures, but the difference in their impact would have to be enormous. Admittedly, also, Caselli's calculations are based on model (1). But given how little impact these corrections have on that version of the production function, it is hard to imagine that similar corrections would do much to

flatten the upward sloping pattern of A_s/A_u vis-à-vis y in model (4).

C. Exogenous Technology

In our model we took the skill distribution as given, and showed how (and when) high L_s/L_u countries choose high (A_s/A_u) technologies. Another view may be that some countries have exogenously high (A_s/A_u)—perhaps for Ricardian reasons—and the resulting higher skill premia lead to greater human capital investment, thereby providing an alternative explanation for the correlation between L_s/L_u and (A_s/A_u). The problem with this alternative view is that, as implied by Table 1, (L_s/L_u) and (w_s/w_u) are negatively correlated: countries in which skilled workers are relatively efficient generally have a low skilled-wage premium. This is clearly inconsistent with the alternative interpretation. Caselli and Coleman (2001b) exploit a similar observation to interpret the interactions between relative wages and relative labor supplies over the last century in the United States.

IV. Some Quantitative Implications of the Model

If the interpretation of the data laid out in Section III is correct, we can use the model to extract a number of interesting quantitative implications on the importance of appropriate technology, and on the size of barriers to technology adoption. How severe a trade-off (as dictated by γ and ω) do countries face in their technology choice? How large are the differences in the distance from the origin (as captured by B) of different countries' frontiers?

In this section, we first show how one can parameterize our simple model to infer each country's frontier. That is, the additional structure generated by the model allows us to go beyond the simple backing out of A_u and A_s of the first part of the paper, and to quantify the parameters ω , γ , and B with which we can plot the family of equations (6) that rationalizes each country's choice. With the country-specific frontiers at hand, we then assess the implications of both counterfactual movements from appropriate to inappropriate technology along these frontiers, and counter-

factual removal of barriers that lift the frontiers themselves. We can also revisit the literature on development accounting.

A. Backing out the Frontiers

Our approach to backing out each country's frontier is very simple. The first step is to relax our assumption that all countries face the same trade-off parameter γ , and allow each country's realization of γ to be a random variable uncorrelated with its endowments. With that assumption, we can rewrite (8) in log as

$$(13) \quad \log \left(\frac{A_s^i}{A_u^i} \right) = \frac{\sigma}{\omega - \sigma} \log \left(\frac{L_s^i}{L_u^i} \right) + \frac{1}{\omega - \sigma} \log \gamma^i,$$

where we have now introduced country superscripts to distinguish country-varying from invariant parameters or variables. Then, an estimate of $\sigma/(\omega - \sigma)$ can be obtained by regressing $\ln(A_s^i/A_u^i)$ on $\ln(L_s^i/L_u^i)$. From this estimate (and our calibrated value of σ), we can back out an estimate of ω . Furthermore, each country's trade-off coefficient γ^i can be recovered from the regression residual. With γ^i and ω at hand, we can then back out each country's B^i (and hence its technology frontier) from equation (6). In performing this exercise, we use our preferred definition of skilled and our preferred calibration of α and σ (and the values of A_s and A_u that these choices imply). It is important to note that in our approach we do not need to impose any restriction on the covariance between L_s/L_u and B . We do, however, have to make the admittedly strong assumption that γ^i is uncorrelated with (L_s^i/L_u^i).

Our estimate of ω is 0.41, with a standard error of (0.01).²⁴ To illustrate the implications of this estimate, Figure 5 depicts the country frontiers for (from highest to lowest) Italy, Argentina, and India, as well as every other country's observed

²⁴ These are backed out from a regression coefficient of 2.280 and a regression standard error (0.086). Recall that our preferred value for $1/(1 - \sigma)$ is 1.4, which implies $\sigma = 0.286$. The R^2 of the regression is 0.93.




FIGURE 5. TECHNOLOGY FRONTIERS OF ITALY (TOP), ARGENTINA (MIDDLE), AND INDIA (BOTTOM)

technology choice (A_u , A_s). Note that the poor country, India, is choosing from a frontier that is considerably inside the frontier of a rich country such as Italy, and the middle-income country, Argentina, is somewhere in the middle. This is, of course, consistent with the barriers-to-adoption component of our framework. Indeed, this observation generalizes: the correlation between the log of the “intercept” parameter B and observed log per-worker income is 0.9. Poor countries are very systematically on lower technology frontiers.

B. Counterfactual Calculations

With each country’s frontier at hand, we can identify the World Technology Frontier as the outer envelope of the country-specific frontiers. In other words, for each A_u , we maximize over all of the country frontiers to find the maximum possible value of A_s . As one of the authors is happy to report, and as is already discernible from Figure 5, this envelope coincides entirely with the frontier of Italy. As we have already seen, checking out each country’s observed position relative to the world technology frontier reveals that poor countries are typically farther away from the world frontier than rich ones.

With the world technology frontier at hand, we can turn to the counterfactual calculations. In order to assess the quantitative importance of appropriateness, we ask the following question: *holding constant the technology frontier*, by

how much would a country’s GDP change, should this country operate a technology (a pair of A_s) different from the appropriate one? In other words, we assess the output consequences of movements along a given technology frontier. For this experiment, we (counterfactually) assume that all countries have access to the world technology frontier. We then compute the level of GDP associated with an optimal (appropriate) choice of technology on the world technology frontier. Finally, we compare this number with the level of GDP the same country would have if forced to use the technology that is appropriate for the United States. In other words, for each country we compare two points on the world technology frontier: the one corresponding to that country’s optimal choice, and the one corresponding to the optimal choice of the United States.

The result of this experiment is plotted in Figure 6, where the vertical axis measures the ratio of U.S.-appropriate-technology GDP to appropriate-technology (on the world frontier) GDP; and the horizontal axis measures actual per capita output. As can be seen, the adoption of an inappropriate technology involves very large output losses—up to 50 percent of GDP—the more so the more different the levels of development (and hence factor endowments). We interpret this finding as indicating that appropriate-technology considerations could play an important role in determining technology differences across countries.

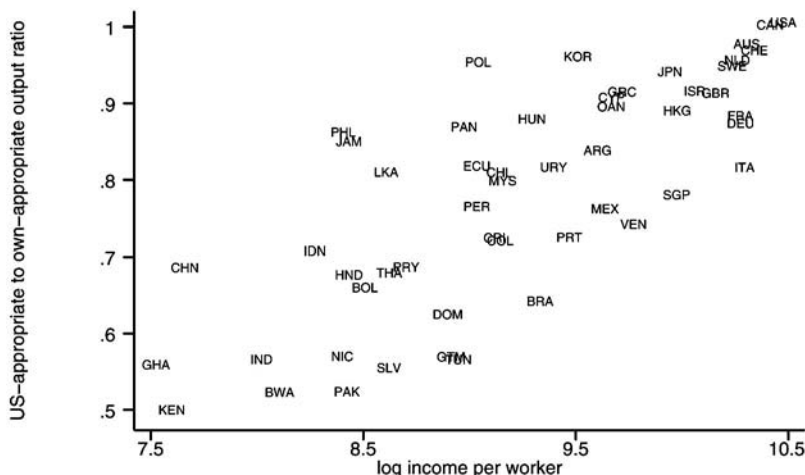


FIGURE 6. RELATIVE OUTPUT FROM USING U.S.-APPROPRIATE TECHNOLOGY

Next, we turn to a quantitative assessment of the “barriers” to technology adoption. If the role of appropriateness could be gleaned by movements along a constant frontier, looking at barriers involves some measure of the (economic) distance between different frontiers. Figure 7 compares each country’s observed level of GDP, with the level of GDP the country would obtain if it had access to the world technology frontier. Hence, we now compare two points on different technology frontiers: the one corresponding to that country’s optimal choice on the world frontier, and the one corresponding to its optimal choice on its own frontier. Both points are “appropriate,” but they are conditional on different choice sets. The Figure shows staggering effects from barriers to technology adoption, with output increasing by up to a factor of 6.5 if such barriers were removed.

C. Development Accounting

Several authors working with the skill neutral formulation in (1) have performed development-accounting exercises. Development accounting asks what fraction of the cross-country variance of income can be explained by observed factors of

production, and how much by the residual “TFP” term. The consensus view is that variation in TFP is roughly as important a source-of-income differences as the combined effects of variation in observed factor inputs, such as physical and human capital. As an illustration, using (1) and our data, we find that variation in A accounts for 40 percent of the variation in y .^{25, 26}

In our framework based on (5) and (6), the analogous to a higher TFP is a higher technology frontier, i.e., a higher value of B . The higher the value of B , the higher the overall technological prowess of the country, which is what the TFP term in standard accounting exercises is all about. The analogy between TFP and B is borne out quantitatively by our data: the correlation between $\log(B)$ as esti-

²⁵ To perform this “standard” development accounting exercise, we constructed h from the average number of years of education in the working-age population, s , as collected by Barro and Lee. We followed current practice in development accounting and computed $h = \exp(0.10s)$, where 0.10 is roughly the world average Mincerian coefficient on education in wage regressions.

²⁶ Other studies tend to find an even larger role for A . Some of the papers in the development-accounting literature are N. Gregory Mankiw et al. (1992), Robert G. King and Ross Levine (1994), Nazrul Islam (1995), Caselli et al. (1996), Klenow and Rodríguez-Clare (1997), Hall and Jones (1997), Parente and Prescott (2000), and Caselli (2005).




FIGURE 7. THE GAIN FROM ACCESSING THE WORLD TECHNOLOGY FRONTIER

imated in Section IVA and $\log(A)$ as estimated in the previous paragraph is 0.96! Hence, the same countries that are estimated to have a low TFP in the skill-neutral approach are estimated to have a low technology frontier in our framework. Hence, there is clearly strong support for the consensus view that poor countries are generally inside the world technology frontier.

Indeed, adding appropriate-technology considerations, as we have done, should deliver an upwardly revised estimate of the role of “barriers” (relative to factor endowments) in explaining cross-country income differences. The reason is that the appropriate choice of technology dampens the effects of differences in factor endowments: countries with “poor” endowments can “remedy” by tailoring their technology choice to their factor supplies, an option denied to them when all differences are factor neutral.²⁷

²⁷ Heuristically, in the appropriate-technology framework one asks what would income dispersion be if all countries had pairs of A_u and A_s chosen from the same frontier, while in the factor-neutral approach one asks what would the dispersion of incomes be if all countries used the same pair of A_u and A_s . If this “TFP pair” is close to the optimal choice of rich countries, the appropriate-technology framework will assign less of a role to factor differences than the TFP framework; while if the TFP pair is close to

We therefore close the paper by adapting the development-accounting calculation for our framework. In particular, we can compute the cross-country variation of per capita income that would be predicted by this model if all countries had access to the same set of potential technologies. This is analogous to the exercise in the development-accounting literature, where it is asked how GDP would vary if there were no differences in TFP. In our model, if each country could choose the point on the world frontier that maximizes its output—given its labor endowments—the standard deviation of the log of per capita GDP would be 0.41. This compares to a value of 0.8 in the data. Hence, differences in inputs explain just about 50 percent of the observed disparity of incomes, while the rest is explained by barriers to technology adoption—i.e., by the fact that different countries have different frontiers. As we have seen above, a model in which technological choice is factor neutral leads to a roughly 60–40 split of the responsibility for the variation of income between factor endowments and differences in technology, so we confirm that marrying “barriers” with “appropriateness” makes the former look even bigger.

the optimal choice of poor countries, the role of factor endowments will be magnified in the appropriate-technology approach.

V. Conclusions

A realistic generalization of the aggregate production function to allow for imperfect substitutability among labor types, combined with cross-country data on skill premia, implies that poor countries are relatively—and possibly also absolutely—better at using unskilled labor. The simple view that countries just differ in a multiplicative TFP level, and that all that is needed is to transfer to them the technologies observed in rich countries, is overly simplistic. Instead, the data are better rationalized if one allows for the possibility that at each point in time firms have access to a whole menu of feasible technologies, that some of these feasible technologies are complementary with skilled labor and others with unskilled labor, and that firms in poor countries will choose to make the most of their abundant factor, unskilled labor. It turns out, however, that accepting this appropriate-technology rationalization also implies that poor countries choose from a much narrower menu of feasible technologies than rich countries do, so in no way are our results inconsistent with the “barriers” view of technology differences.

Indeed, given that we find even larger barrier effects than previous contributions in development accounting, we hope that our paper will provide further spur to ongoing efforts to uncover the nature of these barriers. Given our evidence that deviations from appropriateness entail large output losses, however, such efforts should be mindful that pushing poor countries to adopt technologies used by rich countries may not be optimal, particularly if poor countries’ factor endowments are significantly different from those of rich nations. Removing barriers should be understood as widening poor countries’ choices of technology, not passively copying rich countries’ production processes.

The framework developed in this paper could be extended in a number of directions. First, it would be interesting to explore the relation between skilled-labor endowments and the location in space of each country’s technology frontier. The interaction between barriers and skill accumulation is a potentially important one, and it could shed light on the nature of the barriers themselves. Second, it would be interesting to

bring in a dynamic dimension and try to identify how the world and the country-specific frontiers have evolved over time. Again, this would shed further light on the nature of barriers. Furthermore, it would uncover potential factor-biases in technological change over time.

Third, one could attempt to unpack the aggregate data we used and look at cross-country barriers and appropriateness at the industry level. A potentially fruitful way to interpret the frontier we identify is not that countries are faced with different ways of producing the same good, but rather that goods may differ in the relative efficiency of different factors. This would provide our model of appropriate technology with roots in the Heckscher-Ohlin tradition. Repeating our estimates at the industry level would allow us to distinguish between the two interpretations. It would also link our work to the work of Treffer (1993), who has argued that country-specific augmentation of factor supplies (which can be interpreted as country and factor-specific efficiency levels, just as in our analysis) helps explain jointly the pattern of trade in factor services and cross-country differences in factor prices.

APPENDIX: EXISTENCE AND UNIQUENESS OF SYMMETRIC EQUILIBRIUM

Consider first the optimal choice of inputs for a firm that faces given factor prices w_s , w_u , and r , and has a given technology, A_u , A_s . The solution to the cost-minimization problem can be shown to give rise to the following cost function:

$$\begin{aligned} \text{Cost}(w_u, w_s, r; y) \\ = \beta r^\alpha \left[\left(\frac{w_u}{A_u} \right)^{\sigma(\sigma-1)} + \left(\frac{w_s}{A_s} \right)^{\sigma(\sigma-1)} \right]^{((\sigma-1)/\sigma)(1-\alpha)} y, \end{aligned}$$

where $\beta = 1/[\alpha^\alpha(1-\alpha)^{1-\alpha}]$. Note that this cost function also accurately describes minimized costs when A_u or A_s is zero. Now it is obvious that even if A_u and A_s are chosen by the firm, the choice of factors must still be cost-minimizing in the sense above. Furthermore, since the cost function is linear in output the optimal choice of technology must itself be cost minimizing. Hence, the choice of an optimal technology is a choice of $(A_u,$

TABLE A.1—DATA

Country	Code	y	k	L_u^P	L_s^P	$\left(\frac{w_s}{w_u}\right)^P$	L_u^S	L_s^S	$\left(\frac{w_s}{w_u}\right)^S$	L_u^H	L_s^H	$\left(\frac{w_s}{w_u}\right)^H$
Argentina	ARG	14,805	33,151	60	106	1.51	148	29	2.53	192	7	4.23
Australia	AUS	29,858	88,076	17	129	1.24	85	56	1.63	146	15	2.13
Bolivia	BOL	4,953	9,076	75	51	1.33	105	20	1.89	125	6	2.70
Botswana	BWA	3,316	9,885	115	41	2.15	174	5	5.58	190	1	14.50
Brazil	BRA	11,297	21,227	100	62	1.80	129	22	3.75	159	7	7.83
Canada	CAN	33,337	82,443	7	134	1.23	83	56	1.60	159	6	2.07
Chile	CHL	9,323	22,452	73	108	1.62	143	35	2.94	203	8	5.37
China	CHN	2,124	4,156	65	50	1.22	105	13	1.57	124	1	2.01
Colombia	COL	9,360	15,434	83	76	1.75	138	22	3.53	180	5	7.10
Costa Rica	CRI	9,118	16,695	83	77	1.55	126	28	2.67	156	10	4.60
Cyprus	CYP	15,805	37,046	48	144	1.55	125	54	2.69	211	13	4.66
Dom. Rep.	DOM	7,314	12,232	85	49	1.46	116	17	2.33	134	6	3.73
Ecuador	ECU	8,388	21,190	69	107	1.60	126	39	2.89	171	13	5.22
El Salvador	SLV	5,548	5,617	92	38	1.47	123	10	2.39	136	3	3.89
France	FRA	28,972	84,929	46	112	1.49	130	33	2.46	183	7	4.06
Ghana	GHA	1,854	1,218	84	35	1.40	123	5	2.15	130	1	3.29
Greece	GRC	16,607	42,802	27	85	1.11	88	26	1.28	109	9	1.46
Guatemala	GTM	7,431	7,773	98	43	1.81	133	11	3.82	151	3	8.05
Honduras	HND	4,597	6,175	102	75	2.02	152	21	4.87	217	3	11.75
Hong Kong	HKG	21,532	29,128	38	99	1.28	97	39	1.73	151	6	2.35
Hungary	HUN	10,869	33,857	37	89	1.19	111	22	1.47	128	8	1.83
India	IND	3,046	3,775	79	34	1.22	102	12	1.55	115	3	1.99
Indonesia	IDN	3,914	8,084	85	72	1.97	177	11	4.62	222	0	10.80
Israel	ISR	23,362	51,768	37	119	1.29	87	58	1.78	156	14	2.45
Italy	ITA	29,552	82,318	43	66	1.10	91	20	1.23	110	4	1.38
Jamaica	JAM	4,596	12,831	96	185	3.16	294	29	13.36	490	3	56.37
Japan	JPN	20,807	64,181	28	119	1.30	106	43	1.79	152	12	2.48
Kenya	KEN	1,998	2,748	109	34	1.93	159	4	4.38	166	1	9.93
Malaysia	MYA	9,472	23,543	59	82	1.46	127	22	2.33	169	2	3.73
Mexico	MEX	15,330	28,449	81	92	1.76	144	28	3.56	196	7	7.20
Netherlands	NLD	28,550	79,069	24	128	1.34	119	40	1.95	169	10	2.82
Nicaragua	NIC	4,453	8,762	91	40	1.47	114	15	2.39	126	6	3.89
Pakistan	PAK	4,552	3,793	85	30	1.47	107	9	2.39	120	2	3.89
Panama	PAN	7,898	19,794	63	139	1.73	142	47	3.43	227	11	6.81
Paraguay	PRY	6,015	9,689	88	68	1.58	141	19	2.82	170	5	5.00
Peru	PER	8,387	18,075	65	75	1.38	106	30	2.07	139	10	3.11
Philippines	PHL	4,473	8,042	46	97	1.38	103	37	2.05	141	13	3.06
Poland	POL	8,439	33,949	19	98	1.12	98	24	1.30	119	7	1.50
Portugal	PRT	12,960	29,437	64	59	1.49	118	14	2.46	142	3	4.06
S. Korea	KOR	13,483	24,651	28	159	1.53	114	61	2.60	219	12	4.41
Singapore	SGP	21,470	56,218	71	89	1.71	150	22	3.34	196	4	6.53
Sri Lanka	LKA	5,476	5,919	51	76	1.32	117	18	1.88	149	1	2.66
Sweden	SWE	27,886	72,777	28	133	1.31	78	67	1.83	170	13	2.55
Switzerland	CHE	30,965	107,870	22	142	1.37	87	64	2.04	192	8	3.02
Taiwan	OAN	15,787	26,240	36	97	1.27	95	37	1.72	143	7	2.32
Thailand	THA	5,558	7,477	87	65	1.52	143	16	2.55	152	8	4.29
Tunisia	TUN	7,696	10,823	82	36	1.38	104	14	2.05	122	3	3.06
UK	GBR	25,775	50,409	36	115	1.31	121	36	1.84	163	9	2.59
USA	USA	35,439	87,330	6	229	1.48	65	116	2.42	237	27	3.94
Uruguay	URY	12,036	23,398	62	97	1.47	139	28	2.39	176	7	3.89
Venezuela	VEN	17,529	42,713	69	71	1.40	111	27	2.13	142	8	3.24
W. Germany	DEU	28,992	89,368	41	94	1.22	122	22	1.55	145	5	1.99

Note: Superscripts "P" ("S", "H") identify variables computed under the "primary completed" ("secondary completed," "college completed") definition of skill.

A_s) on a country's technology frontier that minimizes this cost function.

Make the change of variables $D_u = (A_u)^\omega$ and $D_s = (A_s)^\omega$. To simplify the notation, also write $\theta = \sigma/[\omega(1 - \sigma)]$. We can then write the firm's problem as

$$\begin{aligned} \text{Min}_{\{D_s, D_u\}} \{ & \text{Cost}(w_u, w_s, r; y) \\ & = \beta r^\alpha [(w_u)^{\sigma(\sigma-1)} (D_u)^\theta \\ & + (w_s)^{\sigma(\sigma-1)} (D_s)^\theta]^{((\sigma-1)/\sigma)(1-\alpha)} y \}. \end{aligned}$$

$$\text{Subject to: } D_s + \gamma D_u = B.$$

Consider first the case where $\theta < 1$, or $\omega > \sigma/(1 - \sigma)$. It is clear in this case that the firm's problem has a unique interior solution. Hence, if this condition is satisfied, all firms choose the same interior technology. The particular technology choice depends on factor prices. From the first-order conditions for an interior optimum, we have (10)—which shows that if firms are in a symmetric equilibrium, there is a unique equilibrium wage ratio for given L_s/L_u . Hence, we have existence and uniqueness in the $\theta < 1$ case.

For the $\theta > 1$ case, it is immediate that the firm cost-minimization problem requires firms to be at a corner, with either $A_s = L_s = 0$ or $A_u = L_u = 0$. The zero-profit condition for firms choosing the former strategy is $r^\alpha (w_u/(B/\gamma))^{1/\omega} 1^{-\alpha} = 1$, where the left-side term is the unit production cost and the right side is the unit revenue. Similarly, for firms choosing the latter strategy we have $r^\alpha (w_s/B)^{1/\omega} 1^{-\alpha} = 1$. Finally, the marginal product of capital must equal the interest rate, or $\alpha k^{\alpha-1} (L_s/B)^{1/\omega} 1^{-\alpha} = r$. These three conditions identify unique equilibrium values of w_u , w_s , and r . Note that at these factor prices, firms are indifferent between hiring only skilled workers or only unskilled workers. This indifference guarantees full employment.

REFERENCES

- Acemoglu, Daron.** "Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality." *Quarterly Journal of Economics*, 1998, 113(4), pp. 1055–89.
- Acemoglu, Daron.** "Directed Technical Change." *Review of Economic Studies*, 2002, 69(4), pp. 781–809.
- Acemoglu, Daron and Zilibotti, Fabrizio.** "Productivity Differences." *Quarterly Journal of Economics*, 2001, 116(2), pp. 563–606.
- Aghion, Philippe; Howitt, Peter and Mayer-Foulkes, David.** "The Effect of Financial Development on Convergence: Theory and Evidence." *Quarterly Journal of Economics*, 2005, 120(1), pp. 173–222.
- Atkinson, Anthony B. and Stiglitz, Joseph E.** "A New View of Technological Change." *Economic Journal*, 1969, 79(315), pp. 573–78.
- Autor, David H.; Katz, Lawrence F. and Krueger, Alan B.** "Computing Inequality: Have Computers Changed the Labor Market?" *Quarterly Journal of Economics*, 1998, 113(4), pp. 1169–1213.
- Barro, Robert J. and Lee, Jong-Wha.** "International Data on Educational Attainment: Updates and Implications." *Oxford Economic Papers*, 2001, 53(3), pp. 541–63.
- Barro, Robert J. and Sala-i-Martin, Xavier.** "Technological Diffusion, Convergence, and Growth." *Journal of Economic Growth*, 1997, 2(1), pp. 1–26.
- Basu, Susanto and Weil, David N.** "Appropriate Technology and Growth." *Quarterly Journal of Economics*, 1998, 113(4), pp. 1025–54.
- Bernanke, Ben S. and Gürkaynak, Refet S.** "Is Growth Exogenous? Taking Mankiw, Romer, and Weil Seriously," in Ben S. Bernanke and Kenneth Rogoff, eds., *NBER macroeconomics annual 2001*. Cambridge, MA: MIT Press, 2002, pp. 11–57.
- Bils, Mark and Klenow, Peter J.** "Does Schooling Cause Growth?" *American Economic Review*, 2000, 90(5), pp. 1160–83.
- Caselli, Francesco.** "Accounting for Cross-Country Income Differences," in Philippe Aghion and Steven Durlauf, eds., *Handbook of economic growth*, Vol. 1. Amsterdam: North-Holland, 2005, chap. 10.
- Caselli, Francesco and Coleman, Wilbur John, II.** "Cross-Country Technology Diffusion: The Case of Computers." *American Economic Review*, 2001a (*Papers and Proceedings*), 91(2), pp. 328–35.
- Caselli, Francesco and Coleman, Wilbur John, II.** "The U.S. Structural Transformation and Regional Convergence: A Reinterpretation." *Journal of Political Economy*, 2001b, 109(3), pp. 584–616.

- Caselli, Francesco and Coleman, Wilbur John, II.** "The U.S. Technology Frontier." *American Economic Review (Papers and Proceedings)*, 2002, 92(2), pp. 148–52.
- Caselli, Francesco; Esquivel, Gerardo and Lefort, Fernando.** "Reopening the Convergence Debate: A New Look at Cross-Country Growth Empirics." *Journal of Economic Growth*, 1996, 1(3), pp. 363–89.
- Caselli, Francesco and Gennaioli, Nicola.** "Dynastic Management." National Bureau of Economic Research, Inc., NBER Working Papers: No. 9442, 2003.
- Caselli, Francesco and Wilson, Daniel J.** "Importing Technology." *Journal of Monetary Economics*, 2004, 51(1), pp. 1–32.
- Ciccone, Antonio and Peri, Giovanni.** "Long-Run Substitutability between More and Less Educated Workers: Evidence from U.S. States 1950–1990." *Review of Economics and Statistics*, 2005, 87(4), pp. 652–63.
- Diamond, Peter; McFadden, Daniel L. and Rodriquer, Miguel.** "Measurement of the Elasticity of Factor Substitution and Bias of Technical Change," in Melvyn Fuss and Daniel L. McFadden, eds., *Production economics: A dual approach to theory and applications. Volume II: Applications to the theory of production*. Amsterdam: North-Holland, 1978, chap. 5.
- Diwan, Ishac and Rodrik, Dani.** "Patents, Appropriate Technology, and North-South Trade." *Journal of International Economics*, 1991, 30(1–2), pp. 27–47.
- Duffy, John and Papageorgiou, Chris.** "A Cross-Country Empirical Investigation of the Aggregate Production Function Specification." *Journal of Economic Growth*, 2000, 5(1), pp. 87–120.
- Eaton, Jonathan and Kortum, Samuel S.** "Trade in Capital Goods." *European Economic Review*, 2001, 45(7), pp. 1195–1235.
- Gollin, Douglas.** "Getting Income Shares Right." *Journal of Political Economy*, 2002, 110(2), pp. 458–74.
- Gollin, Douglas; Parente, Steven and Rogerson, Richard.** "Structural Transformation and Cross-Country Income Differences." Unpublished Paper, 2001.
- Hall, Robert E. and Jones, Charles I.** "Why Do Some Countries Produce So Much More Output per Worker than Others?" *Quarterly Journal of Economics*, 1999, 114(1), pp. 83–116.
- Hamermesh, Daniel S.** *Labor demand*. Princeton: Princeton University Press, 1993.
- Hicks, John R.** *The theory of wages*. London: MacMillan, 1932.
- Hicks, John R.** *Value and capital*. 2nd ed. New York: Oxford University Press, 1939.
- Howitt, Peter.** "Endogenous Growth and Cross-Country Income Differences." *American Economic Review*, 2000, 90(4), pp. 829–46.
- Islam, Nazrul.** "Growth Empirics: A Panel Data Approach." *Quarterly Journal of Economics*, 1995, 110(4), pp. 1127–70.
- Jones, Charles I.** "The Shape of Production Functions and the Direction of Technical Change." *Quarterly Journal of Economics*, 2005, 120(2), pp. 517–49.
- Katz, Lawrence F. and Autor, David H.** "Changes in the Wage Structure and Earnings Inequality," in Orley Ashenfelter and David Card, eds., *Handbook of labor economics*. Vol. 3A. Amsterdam: Elsevier Science, North-Holland, 1999, pp. 1463–555.
- Katz, Lawrence F. and Murphy, Kevin M.** "Changes in Relative Wages, 1963–1987: Supply and Demand Factors." *Quarterly Journal of Economics*, 1992, 107(1), pp. 35–78.
- Kennedy, Charles.** "Induced Bias in Innovation and the Theory of Distribution." *Economic Journal*, 1964, 74(295), pp. 541–47.
- King, Robert G. and Levine, Ross.** "Capital Fundamentalism, Economic Development, and Economic Growth." *Carnegie-Rochester Conference Series on Public Policy*, 1994, 40(0), pp. 259–92.
- Klenow, Peter J. and Rodríguez-Clare, Andrés.** "The Neoclassical Revival in Growth Economics: Has It Gone Too Far?" in Ben S. Bernanke and Julio J. Rotemberg, eds., *NBER macroeconomics annual 1997*. Cambridge, MA: MIT Press, 1997, pp. 73–103.
- Klenow, Peter J. and Rodríguez-Clare, Andrés.** "Externalities and Growth," in Philippe Aghion and Steven Durlauf, eds., *Handbook of economic growth*. Vol. IA. Amsterdam: Elsevier North-Holland, 2006.
- Krusell, Per; Ohanian, Lee E.; Ríos-Rull, José-Victor and Violante, Giovanni L.** "Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis." *Econometrica*, 2000, 68(5), pp. 1029–53.
- Mankiw, N. Gregory; Romer, David and Weil, David N.** "A Contribution to the Em-

- pirics of Economic Growth." *Quarterly Journal of Economics*, 1992, 107(2), pp. 407–37.
- Parente, Stephen L. and Prescott, Edward C.** *Barriers to riches. Walras–Pareto lectures*. Vol. 3. Cambridge, MA: MIT Press, 2000.
- Ruiz Arranz, Marta.** "Wage Inequality in the U.S.: Capital–Skill Complementarity vs. Skill-Biased Technological Change." Unpublished Paper, 2002.
- Samuelson, Paul A.** "A Theory of Induced Innovation along Kennedy–Weisacker Lines." *Review of Economics and Statistics*, 1965, 47(4), pp. 343–56.
- Samuelson, Paul A.** "Rejoinder: Agreements, Disagreements, Doubts, and the Case of Induced Harrod-Neutral Technical Change." *Review of Economics and Statistics*, 1966, 48(4), pp. 444–48.
- Summers, Robert and Heston, Alan.** "The Penn World Table (Mark 5): An Expanded Set of International Comparisons, 1950–1988." *Quarterly Journal of Economics*, 1991, 106(2), pp. 327–68.
- Trefler, Daniel.** "International Factor Price Differences: Leontief Was Right!" *Journal of Political Economy*, 1993, 101(6), pp. 961–87.