

Trade Shocks and Labor Adjustment: A Structural Empirical Approach*

Erhan Artuç, Koç University, Istanbul, Turkey;

Shubham Chaudhuri, World Bank, Washington, D.C.;

John McLaren, University of Virginia, Charlottesville, VA

October 3, 2007

Abstract

The welfare effects of trade shocks depend crucially on the nature and magnitude of the costs workers face in moving between sectors. The existing trade literature does *not* directly address this, assuming perfect mobility or complete immobility, or adopting reduced-form approaches to estimation. We present a model of dynamic labor adjustment that *does*, and is, moreover, consistent with a key empirical fact: that intersectoral gross flows greatly exceed net flows. Using an Euler-type equilibrium condition, we estimate the *mean* and the *variance* of workers' switching costs from the U.S. March Current Population Surveys. We estimate high values of both parameters, implying both slow adjustment of the economy, and sharp movements in wages, in response to a trade shock. Simulations of a trade liberalization indicate that despite the high estimated adjustment cost, in terms of lifetime welfare, the liberalization is Pareto-improving. The explanation for this surprising finding—which would be missed by a reduced-form approach—is that the high *variance* to costs ensures high rates of gross flow; this helps spread the liberalization's benefits around.

*We are grateful to seminar participants at the University of Virginia, Syracuse University, Wilfred Laurier University, and the North American Summer Meetings of the Econometric Society in June 2007. This work has been supported by National Science Foundation grant 0080731 and by the Bankard Fund for Political Economy at the University of Virginia.

1 Introduction.

Perhaps the most urgent question facing trade economists is the effect of liberalization and other trade shocks on the welfare of workers. This question has generated a large body of research, but a feature shared by most of the extant trade literature on this is a reliance on static models, in which workers are assumed to be either instantly costlessly mobile, or perfectly immobile (we will discuss important exceptions below). This prevents the trade literature from even addressing, let alone answering, some central questions: What are the costs faced by workers who wish to move to a new industry in response to import competition? How long will the labor market take to adjust, and find its new steady state? Will that steady state feature a lasting differential impact on workers in the import-afflicted sector, or will arbitrage equalize worker returns in the long run? What are the lifetime welfare effects on workers in different industries, taking into account moving costs and transitional dynamics?

This paper offers an approach to answering these questions. Within the context of a standard trade model, we specify a dynamic equilibrium model of costly labor adjustment, a model fully studied in Cameron, Chaudhuri and McLaren (2007). We then show how the structural moving-cost parameters of this model can be estimated, using Euler-equation-type techniques borrowed from macroeconomics. Estimating these parameters on data from the US Current Population Surveys (CPS), we then use these parameters to simulate stylized trade shocks and show their dynamic equilibrium impact.

A large number of studies in the trade economics field have attempted to measure the effects of trade shocks on wages. Some test labor-market predictions of the Heckscher-Ohlin model, as Lawrence and Slaughter (1993). Others regress changes in wages sector-by-sector on changes in import prices (as Revenga (1992)), trade policy (as Pavcnik, Goldberg and Attanasio (2004)) or import penetration (as Kletzer (2002)). Slaughter (1998) provides an overview.

Much has been learned from this literature, but it suffers from three weaknesses that this study addresses directly. First, these all suffer from the Lucas critique; for example, the change in wages observed when a tariff falls may be different depending on whether or not it was anticipated, or whether or not it is expected to be part of a continuing reduction. Reduced-form regressions cannot accommodate such dis-

tinctions. Second, in a dynamic environment, wage changes at a given moment are insufficient for identifying the lifetime effect on a workers' utility, which is what really matters for welfare analysis. Third, perhaps most importantly, these reduced-form studies take no account of the constant inter-industry gross flows of workers observed in the data. In the data, such gross flows are large, and have a large effect on welfare calculations. Indeed, we will see that the conclusions one would draw from standard reduced-form regression results can be *reversed* when gross flows are accounted for.

A small number of studies in the trade literature do study the empirics of dynamic labor market adjustment, but focus on employer-side adjustment. Utar (2007) estimates a dynamic model of firm adjustment to trade shocks with heterogeneous firms. Robertson and Dutkowsky (2002) use an Euler equation approach to estimate employers' labor adjustment costs in Mexico with a focus on international policy but employs a model that rules out gross flows in excess of net flows, thus ruling out an important feature of the data that is central to our approach.

On the other hand, a number of labor economists have developed highly sophisticated structural empirical models that allow them to estimate the impact of policy changes on labor adjustment in a manner similar in some respects to what we are doing here. Examples include Lee (2005) and Keane and Wolpin (1997), who focus on occupational choices of workers rather than inter-industry reallocation, and Kennan and Walker (2003), who study movement of workers across US regions. There are four key differences between those studies and our approach. First, with our emphasis on *intersectoral* reallocation we are tailoring our model to the analysis of trade policy, which cannot be addressed by those other studies. Second, with our Euler-equation approach, which appears not to have been used before in the analysis of workers' mobility choices, we do not need to make any strong assumptions about what workers know about the future (in particular, they do not need to know the future course of aggregate events with certainty, which is assumed in Keane and Wolpin (1997) and Lee (2005), for example). We assume that workers have rational expectations about the future, but we need to make no assumption regarding *how much* information they have about the future. Third, our estimation method is simple and computationally cheap, allowing its application potentially to a very wide range of data sets. The most closely related paper to ours is Artuç (2006), which does esti-

mate a general-equilibrium structural model of worker response to trade shocks, but focuses on *intergenerational* distributional issues and does not use an Euler-equation approach.¹

In our approach, we present a dynamic rational-expectations model² in which each worker can choose to move from her current industry to another one in each period, but must pay a cost to do so. The cost has a common component, which does not vary across time or workers; and a time-varying idiosyncratic component, which can be negative, reflecting non-pecuniary motives that workers often have for changing jobs (such as tedium, a need to relocate for family reasons, and the like). We derive an equilibrium condition, which is a kind of Euler equation, estimate its parameters using the Current Population Survey (CPS), and simulate a trade liberalization to illustrate their implications.

The element of idiosyncratic shocks is crucial to a realistic treatment of worker mobility, for two reasons. First, gross flows are an order of magnitude larger than net flows, implying large numbers of workers moving in opposite directions at the same time. Second, Bowlus and Newmann (2006) show that a significant fraction of workers who change jobs *voluntarily* move to jobs which pay less than the job the worker left behind. Approximately 40% of voluntary job changes have this feature, not very different from the 50% that would be expected if wage differences had *no* effect on mobility decisions. Both of these observations suggest a central role for idiosyncratic shocks in worker mobility. We quantify this in our estimates, and show that it is very important for evaluating the welfare effects of trade policy. In particular, the presence of these shocks imply that option value is an important element of each worker's utility calculation, which, although it can have a decisive effect on the welfare

¹Another related paper, Kambourov (2006), calibrates a model of labor reallocation, which is costly because of sector-specific human capital and firing costs, and applies it to trade reform. It turns out that firing costs have a large negative effect on the gains from trade reform. Unlike our paper, Kambourov's model does not provide workers with idiosyncratic shocks, so it cannot generate gross flows in excess of net flows. Given the importance of gross flows in the data, this is a significant feature of our approach.

²The model we use is presented in full in Cameron, Chaudhuri and McLaren (2007). It is a full-employment model with moving costs for workers. An alternative approach would be to focus on search frictions, as in Hosios (1990), Davidson, Martin and Matusz (1999) and Davidson and Matusz (2001).

effects of trade reform, to our knowledge has never before been introduced into the literature on trade and labor.

The estimates we obtain show very high average moving costs, and a very high standard deviation of moving costs, both estimated to be several times average annual wages for moving from one broadly aggregated sector of the economy to another. These surprisingly high estimated costs are actually in line with related findings by other authors using different techniques; for example, Kennan and Walker's (2003) estimates of costs of moving between US regions, and Artuç's (2006) estimates of intersectoral moving costs.³ In addition, as we will see, simulations based on these patterns produce realistic aggregate behavior. The message conveyed by these findings is that US workers change industry a great deal, but those movements do not respond much to movements in intersectoral wage differentials. Thus, non-pecuniary motives such as are captured by our idiosyncratic shocks must be driving a large portion of our workers' movements. This is important for the analysis of trade liberalization, as our simulations reveal. First, it suggests sluggish adjustment of the labor market to a trade shock, with the economy requiring several years to approach the new steady state. Second, as a corollary, it implies a large drop in wages in the import-competing sector that is hit by the liberalization; indeed, the wages in that sector never fully recover. Third, surprisingly, because of the high levels of mobility due to idiosyncratic shocks, *workers in the import-competing sector benefit from the liberalization*. Their welfare rises because the high volatility of their idiosyncratic shocks combined with rising real wages in other sectors implies that their option value is enhanced by the liberalization, and this effect overwhelms the direct loss from the lower wages in their own sector.⁴ This shows the utility of a dynamic structural approach; a reduced-form wage equation would have, in this case, produced exactly

³It should be noted that this is so even though Artuç (2006) uses a different data set, namely the NLSY; this paper uses the CPS.

⁴This is closely related to the empirical findings of Magee, Davidson and Matusz (2005). They find, for low-turnover industries, that political action committees are much more likely to donate to pro-trade politicians if they represent an export sector than an import-competing sector; but for high-turnover sectors the difference between export and import-competing industries essentially disappears. They rationalize this using a search model of labor adjustment as in Hosios (1990) and Davidson, Martin and Matusz (1999), but the underlying reason is similar: With a high degree of labor flows, workers do not identify closely with the industry in which they are currently located.

the wrong welfare conclusion by showing wage losses for the import-competing sector without identifying the countervailing option-value effect. Indeed, in our simulation section we will show how some of our findings are superficially similar to results from the static, reduced-form regressions of Revenga (1992), but with the opposite welfare implications.

In the following sections we present the model, deriving its estimating equation and explaining the identification strategy intuitively; then examine the data and its measurement issues; then present our estimates and interpret them. The next section deals with a number of measurement and specification issues, and a final section studies simulations based on our estimated parameters.

2 The model.

Consider an economy in which production may occur in any of N industries. We construct a dynamic rational expectations model of labor mobility across these industries, in which our goal is to derive an equilibrium condition that will allow estimation of moving cost parameters.⁵

2.1 Basic setup

Assume that in each industry i there are a large number of competitive employers, and that the value of their aggregate output in any period t is given by $x_t^i = X^i(L_t^i, s_t)$, where L_t^i denotes the labor used in industry i in period t , and s_t is a state variable that could capture the effects of policy (such as trade protection, which might raise the price of the output), technology shocks, and the like. Assume that X^i is strictly increasing, continuously differentiable and concave in its first argument. Its first derivative with respect to labor is then a continuous, decreasing function of labor, holding s_t constant; this is, then, the demand curve for labor in the industry. Assume

⁵In principle, the model can accommodate geographic as well as inter-industry mobility. Instead of N industries, we could have N industry-region cells, for example; all of the logic below would carry through without amendment. In practice, we have limited the discussion to inter-industry mobility because we have not found enough inter-regional mobility in the data to identify the parameters of interest.

that s follows a first-order Markov process on some state space S .

The economy's workers form a continuum of measure \bar{L} . Each worker at any moment is located in one of the N industries. Denote the number of workers in industry i at the beginning of period t by L_t^i . If a worker, say, $l \in [0, \bar{L}]$, is in industry i at the beginning of t , she will produce in that industry, collect the market wage w_t^i for that industry, and then may move to any other industry. In order for the labor market to clear, we must have $w_t^i = \frac{\partial X^i(L_t^i, s_t)}{\partial L_t^i}$ at all times.

If worker l moves from industry i to industry j , she incurs a cost $C^{ij} \geq 0$, which is the same for all workers and all periods, and is publicly known. In addition, if she is in industry i at the end of period t , she collects an idiosyncratic benefit $\varepsilon_{l,t}^i$ from being in that industry. These benefits are independently and identically distributed across individuals, industries, and dates, with density function $f : \mathfrak{R} \rightarrow \mathfrak{R}^+$ and cumulative distribution function $F : \mathfrak{R} \rightarrow [0, 1]$. Thus, the full cost for worker l of moving from i to j can be thought of as $\varepsilon_{l,t}^i - \varepsilon_{l,t}^j + C^{ij}$. The worker knows the values of the $\varepsilon_{l,t}^i$ for all i before making the period- t moving decision.⁶ We adopt the convention that $C^{ii} = 0$ for all i .

Note that the mean cost of moving from i to j is given by C^{ij} , but its variance and other moments are determined by f . It should be emphasized that these higher moments are important both for estimation and for policy analysis, as will be discussed below.

All agents have rational expectations and a common constant discount factor $\beta < 1$, and are risk neutral.

An equilibrium then takes the form of a decision rule by which, in each period, each worker will decide whether to stay in her industry or move to another, based on the current allocation vector L of labor across industries, the current aggregate state s , and that worker's own vector ε of shocks. In the aggregate, this decision rule will generate a law of motion for the evolution of the labor allocation, and hence (by the labor market clearing condition just mentioned) for the wage in each industry. Given this behaviour for wages, the decision rule must be optimal for each worker, in the sense of maximizing her expected present discounted value of wages plus idiosyncratic

⁶It is useful to think of the timeline as follows: The worker observes s_t at the beginning of the period, produces output and receives the wage, then learns the vector $\varepsilon_{l,t}$ and decides whether or not to move. At the end of the period, she enjoys $\varepsilon_{l,t}^j$ in whichever sector j she has landed.

benefits, net of moving costs.

2.2 The key equilibrium condition.

Suppose that we have somehow computed the maximized value to each worker of being in industry i when the labor allocation is L and the state is s . Let $U^i(L, s, \varepsilon)$ denote this value, which, of course, depends on the worker's realized idiosyncratic shocks. Denote by $V^i(L, s)$ the average of $U^i(L, s, \varepsilon)$ across all workers, or in other words, the expectation of $U^i(L, s, \varepsilon)$ with respect to the vector ε . Thus, $V^i(L, s)$ can also be interpreted as the expected value of being in industry i , conditional on L and s , but before the worker learns her value of ε .

Assuming optimizing behavior, i.e., that a worker in industry i will choose to remain at or move to the industry j that offers her the greatest expected benefits, net of moving costs, we can write:⁷

$$\begin{aligned} U^i(L_t, s_t, \varepsilon_t) &= w_t^i + \max_j \{ \varepsilon_t^j - C^{ij} + \beta E_t[V^j(L_{t+1}, s_{t+1})] \} \\ &= w_t^i + \beta E_t[V^i(L_{t+1}, s_{t+1})] + \max_j \{ \varepsilon_t^j + \bar{\varepsilon}_t^{ij} \} \end{aligned} \quad (1)$$

where:

$$\bar{\varepsilon}_t^{ij} \equiv \beta E_t[V^j(L_{t+1}, s_{t+1}) - V^i(L_{t+1}, s_{t+1})] - C^{ij} \quad (2)$$

Note that L_{t+1} is the next-period allocation of labor, derived from L_t and the decision rule, and s_{t+1} is the next-period value of the state, which is a random variable whose distribution is determined by s_t . The expectations in (1) and (2) are taken with respect to s_{t+1} , conditional on all information available at time t .

Taking the expectation of (1) with respect to the ε vector then yields:

$$V^i(L_t, s_t) = w_t^i + \beta E_t[V^i(L_{t+1}, s_{t+1})] + \Omega(\bar{\varepsilon}_t^i), \quad (3)$$

where $\bar{\varepsilon}_t^i = (\bar{\varepsilon}_t^{i1}, \dots, \bar{\varepsilon}_t^{iN})$ and:

$$\Omega(\bar{\varepsilon}_t^i) = \sum_{j=1}^N \int_{-\infty}^{\infty} (\varepsilon^j + \bar{\varepsilon}_t^{ij}) f(\varepsilon^j) \prod_{k \neq j} F(\varepsilon^j + \bar{\varepsilon}_t^{ij} - \bar{\varepsilon}_t^{ik}) d\varepsilon^j. \quad (4)$$

The average value to being in industry i can therefore be decomposed into three terms: (1) the wage, w_t^i , that a industry- i worker receives; (2) the base value of staying on in

⁷From here on, we drop the worker-specific subscript, l .

industry i , i.e., $\beta E_t[V^i(L_{t+1}, s_{t+1})]$; and (3) the additional value, $\Omega(\bar{\varepsilon}_t^i)$, derived from having the option to move to another industry should prospects there look better (and which is simply equal to the expectation of $\max_j\{\varepsilon^j + \bar{\varepsilon}_t^{ij}\}$ with respect to the ε vector). We will call this the ‘option value’ associated with being in that industry at that time.

Using (3), we can rewrite (2) as:

$$\begin{aligned}
C^{ij} + \bar{\varepsilon}_t^{ij} &= \beta E_t[V^j(L_{t+1}, s_{t+1}) - V^i(L_{t+1}, s_{t+1})] \\
&= \beta E_t[w_{t+1}^j - w_{t+1}^i + \beta E_{t+1}[V^j(L_{t+2}, s_{t+2}) - V^i(L_{t+2}, s_{t+2})] \\
&\quad + \Omega(\bar{\varepsilon}_{t+1}^j) - \Omega(\bar{\varepsilon}_{t+1}^i)], \text{ or} \\
C^{ij} + \bar{\varepsilon}_t^{ij} &= \beta E_t[w_{t+1}^j - w_{t+1}^i + C^{ij} + \bar{\varepsilon}_{t+1}^{ij} + \Omega(\bar{\varepsilon}_{t+1}^j) - \Omega(\bar{\varepsilon}_{t+1}^i)] \tag{5}
\end{aligned}$$

Note that $\bar{\varepsilon}_t^{ij}$ is the value of $\varepsilon^i - \varepsilon^j$ at which a worker in industry i is indifferent between moving to industry j and staying in i . Condition (5) thus has the simple, common-sense interpretation that for the *marginal* mover from i to j , the cost (including the idiosyncratic component) of moving is equal to the expected future benefit of being in j instead of i at time $t + 1$. This expected future benefit has three components. The first is the wage differential. The second is the revealed expected value to being in industry j instead of i at time $t + 2$, as revealed by the cost borne by the marginal mover from i to j at time $t + 1$, or $C^{ij} + \bar{\varepsilon}_{t+1}^{ij}$. The last component is the difference in option values associated with being in each industry. Thus, if I contemplate being in j instead of i next period, I take into account the expected difference in wages; then the difference in the expected values of continuing in each industry afterward; and finally, the differences in the values of the option to leave each industry if conditions call for it.

Put differently, condition (5) is an Euler equation. Given appropriate choice of functional forms, this can be implemented to estimate the moving-cost parameters. We turn to that task next.

2.3 The estimating equation.

Let m_t^{ij} be the fraction of the labor force in industry i at time t that chooses to move to industry j , i.e., the *gross flow* from i to j . With the assumption of a continuum of

workers and i.i.d idiosyncratic components to moving costs, this gross flow is simply the probability that industry j is the best for a randomly selected i -worker. Now, make the following functional form assumption. Assume that the idiosyncratic shocks follow an extreme-value distribution with parameters $(-\gamma\nu, \nu)$:

$$\begin{aligned} f(\varepsilon) &= \frac{e^{-\varepsilon/\nu-\gamma}}{\nu} \exp\{-e^{-\varepsilon/\nu-\gamma}\} \\ F(\varepsilon) &= \exp\{-e^{-\varepsilon/\nu-\gamma}\}, \end{aligned}$$

implying:

$$\begin{aligned} E(\varepsilon) &= 0, \text{ and} \\ \text{Var}(\varepsilon) &= \frac{\pi^2\nu^2}{6}. \end{aligned}$$

(For further properties of the extreme-value distribution, see Patel, Kapadia, and Owen (1976).)

Note that while we make the natural assumption that the ε 's be mean-zero, we do not impose any restrictions on the variance. The variance is proportional to the square of ν , which is a free parameter to be estimated, and crucial for all of the policy and welfare analysis.

By assuming that the ε_t^i are generated from an extreme-value distribution we are able to obtain a particularly simple expression for the conditional moment restriction, which we then plan to estimate using aggregate data. Specifically, it is shown in the Appendix that, with this assumption:

$$\bar{\varepsilon}_t^{ij} \equiv \beta E_t[V_{t+1}^j - V_{t+1}^i] - C^{ij} = \nu[\ln m_t^{ij} - \ln m_t^{ii}] \quad (6)$$

and:

$$\Omega(\bar{\varepsilon}_t^i) = -\nu \ln m_t^{ii} \quad (7)$$

Both these expressions make intuitive sense. The first says that the greater the expected net (of moving costs) benefits of moving to j , the larger should be the observed ratio of movers (from i to j) to stayers. Moreover, holding constant the

(average) expected net benefits of moving, the higher the variance of the idiosyncratic cost shocks, the lower the compensating migratory flows.

The second expression says that the greater the probability of remaining in industry i , the lower the value of having the option to move from industry i .⁸ Moreover, as the variance of the idiosyncratic component of moving costs increases, so too does the value of having the option to move. This also makes good sense.

Substituting from (6) and (7) into (5) and rearranging, we get the following conditional moment condition:

$$E_t \left[\frac{\beta}{\nu} (w_{t+1}^j - w_{t+1}^i) + \beta (\ln m_{t+1}^{ij} - \ln m_{t+1}^{jj}) - \frac{(1-\beta)}{\nu} C^{ij} - (\ln m_t^{ij} - \ln m_t^{ii}) \right] = 0. \quad (8)$$

This condition can be interpreted as a linear regression:

$$(\ln m_t^{ij} - \ln m_t^{ii}) = -\frac{(1-\beta)}{\nu} C^{ij} + \frac{\beta}{\nu} (w_{t+1}^j - w_{t+1}^i) + \beta (\ln m_{t+1}^{ij} - \ln m_{t+1}^{jj}) + \mu_{t+1}, \quad (9)$$

where μ_{t+1} is news revealed at time $t+1$, so that $E_t \mu_{t+1} \equiv 0$. In other words, the parameters of interest, C^{ij} , β and ν , can then be estimated by regressing current flows (as measured by $(\ln m_t^{ij} - \ln m_t^{ii})$) on future flows (as measured by $(\ln m_{t+1}^{ij} - \ln m_{t+1}^{jj})$) and the future wage differential with an intercept. Of course, the disturbance term, μ_{t+1} , will in general be correlated with the regressors, requiring instrumental variables. The theory implies that past values of the flows and wages will be valid instruments, and the optimal weighting scheme can be derived as in the Generalized Method of Moment (GMM) (Hansen (1982)). Note that while our choice of f obviously determined the form of the estimating equation, under the GMM estimation procedure, we do not need to make any additional assumptions about the process governing the state variables, s_t .

2.4 Identification

It may be helpful to review how the model provides a strategy for identifying the parameters of interest to us. Roughly, the logic of the model tells us that the *level* of gross flows in the data helps us pin down the *ratio* of average moving costs to the

⁸Note that $0 < m_t^{ii} < 1$, so $\Omega(\bar{\varepsilon}_t^i) = -\nu \ln m_t^{ii} > 0$.

variance of moving costs (that is, the ratios of the C^{ij} 's to ν), and the *responsiveness* of labor flows to anticipated wage differentials pins down the *level* of ν . Essentially, both the overall level of gross flows and their responsiveness to wages together pin down the values of the parameters. To see how, first note that worker flows are given by the following:

$$m^{ij} = \frac{\exp(\bar{\varepsilon}^{ij}/\nu)}{\sum_{k=1}^n \exp(\bar{\varepsilon}^{ik}/\nu)}$$

This is derived from the properties of the extreme-value distribution, and is essentially the same as the outcome of the familiar extreme-value multinomial choice problem (a detailed derivation is presented in the appendix). Now consider a simplified version of the model in which labor demand in each industry is identical and non-stochastic, and $C^{ij} \equiv C \forall i \neq j$. In the steady state of such a model, $L^i = L^j$ and $V^i = V^j \forall i, j$. Therefore, $\bar{\varepsilon}^{ij} = -C \forall i \neq j$, and:

$$\begin{aligned} m^{ij} &= \frac{\exp(-C/\nu)}{1+(n-1)\exp(-C/\nu)} \\ &= \frac{1}{\exp(C/\nu)+(n-1)} \end{aligned} \tag{10}$$

$\forall i \neq j$.

Thus, the level of steady-state gross flows is a decreasing function of C/ν . This is easy to understand, as a rise in C raises costs of changing industries, discouraging mobility, and a rise in ν fattens the tails of the idiosyncratic shocks, increasing the probability that a given worker has an idiosyncratic moving cost below the threshold required to move. (Or, viewed differently, a rise in ν raises the importance of non-pecuniary factors in mobility decisions, making workers more likely to change industries for non-pecuniary reasons.)

Thus, in this simplified model, observing what fraction of workers change their industry per period allows us to pin down the ratio C/ν . Note that in our estimation equation (9) this ratio is proportional to the intercept, so that a general increase in gross flows in the data (for given β) will result in lower values for the C^{ij}/ν ratios. This can be illustrated with Figure 1. A high value of observed flows would imply

a ray in C, ν space with a low slope, such as OA , while a lower value of gross flows would imply a point on a ray with a higher slope, such as OB . Now, what identifies the point upon that ray that the true parameter values must occupy?

Note from (9) that the coefficient multiplying the next-period wage differential is β/ν . A straightforward interpretation of this is that the coefficient β/ν measures the degree to which the future wage differential predicts the current rate of gross flow, $(\ln m_t^{ij} - \ln m_t^{ii})$. Thus, holding β constant, if future wage differentials are a good predictor for current labor flows, then we will obtain a low estimate for ν . This can be understood in two ways. First, realize that a high value of ν means that idiosyncratic and non-pecuniary factors are dominant in workers' mobility decisions, so that workers do not pay much attention to wages when making those decisions. Thus, a high value of ν implies that wages will be relatively irrelevant as a determinant of labor flows. A second interpretation is in terms of elasticities of labor supply: If we think of a labor supply model in which workers have individual disutilities to work and will join the labor force only if the wage exceeds the disutility, then a high variance of that disutility in the population of potential workers implies a vertical labor-supply curve and a low elasticity of supply, so that the wage has a small effect on the amount of aggregate labor supplied. This is analogous to the effect observed in our model, but in a setting of dynamic, intersectoral labor supply: A high idiosyncratic variance implies a low elasticity of response to wages.

Thus, roughly, the overall level of gross flows pins down the C/ν ratio, and the level of responsiveness of labor flows to future wage differentials pins down the level of C and ν . For a given level of flows, if wages do not matter much for explaining variation in flows over time, a high value of both C and ν will be implied.

A note on measurement error may be appropriate here, as well. If systematic errors in coding of workers' industry are present so that spurious industry mobility occurs in the data, that will both put the parameters on a lower ray (by putting excess mobility into the data) and put them on a higher point along that ray (by making wages appear less relevant to mobility, since coding errors are likely uncorrelated with anticipated wages). Thus, coding errors can in principal result in over- or underestimates of C , but will definitely provide an overestimate of ν and an underestimate of the ratio C/ν .

3 Data.

Our estimation strategy hinges on observing aggregate gross flows across industries. Since there are no published data on gross flows, we construct gross flow measures from individual-level data. For this purpose, we use the US Census Bureau’s March Current Population Surveys (CPS). Each year, the March CPS provides information on the individual’s industry, occupation, and employment status at the time of the March interview, as well as the industry, occupation, and employment status in which the individual spent the most time during the previous calendar year (i.e., January to December). We use this information to construct rates of flow, m_{t-1}^{ij} for each date t . We also obtain industry wages w_t^i as the average wage reported in the CPS samples for industry i at date t . These are deflated by the CPI, and normalized so that over the whole sample the average annualized wage is equal to unity. We restrict the sample to males aged 25 to 64 currently working full time who worked at least 26 weeks in the previous year and whose most recent weekly income was between \$50 and \$5,000.

Table 1: Descriptive Statistics: Gross Flows, 1975-2000.

	<i>Agric/Min</i>	<i>Const</i>	<i>Manuf</i>	<i>Trans/Util</i>	<i>Trade</i>	<i>Service</i>
<i>Agric/Min</i>	0.9292 (0.0146)	0.0126 (0.0040)	0.0142 (0.0046)	0.0075 (0.0032)	0.0160 (0.0063)	0.0206 (0.0057)
<i>Const</i>	0.0056 (0.0028)	0.9432 (0.0108)	0.0139 (0.0029)	0.0063 (0.0023)	0.0119 (0.0027)	0.0191 (0.0040)
<i>Manuf</i>	0.0020 (0.0008)	0.0041 (0.0008)	0.9708 (0.0035)	0.0031 (0.0010)	0.0080 (0.0012)	0.0120 (0.0021)
<i>Trans/Util</i>	0.0025 (0.0011)	0.0044 (0.0018)	0.0068 (0.0016)	0.9643 (0.0050)	0.0081 (0.0023)	0.0138 (0.0033)
<i>Trade</i>	0.0030 (0.0011)	0.0061 (0.0015)	0.0135 (0.0033)	0.0055 (0.0017)	0.9469 (0.0073)	0.0250 (0.0036)
<i>Service</i>	0.0018 (0.0008)	0.0043 (0.0011)	0.0079 (0.0013)	0.0037 (0.0008)	0.0103 (0.0014)	0.9720 (0.0033)

(Origin sector is listed by row, destination sector by column. Each cell of table contains mean flow rate followed by standard deviation in parentheses.)

If we have n industries, then there are n^2 rates of gross flow to keep track of each period (or $n(n - 1)$ if one excludes the fraction of workers in each industry who do not move). Thus, the number of directions for gross flows proliferates rapidly as the number of industries increases, leading in finite samples to zero observations and observations with very small numbers of individuals. As a result, we need to aggregate industries, and we aggregate to the following six: 1. Agriculture and Mining; 2. Construction; 3. Manufacturing; 4. Transportation, Communication, and Utilities; 5. Trade; and 6. All Other Services including government. As a result of this aggregation, the sample size for each regression is 720, since we have 26 years minus 2 to allow for lags, and 6 times 5 directions of flows.

Table 2: Descriptive Statistics: Wages, 1975-2000.

	<i>Mean Wage (in 2000\$)</i>	<i>Standard deviation of wage (in 2000\$)</i>	<i>Mean wage, normalized</i>	<i>Standard deviation of wage (nor- malized)</i>	<i>Sample size</i>
<i>Agric/Min</i>	34,739	24,978	0.8374	0.6021	20,952
<i>Const</i>	38,432	21,623	0.9265	0.5213	44,943
<i>Manuf</i>	42,655	21,706	1.0283	0.5233	140,339
<i>Trans/Util</i>	43,608	20,552	1.0512	0.4954	55,699
<i>Trade</i>	37,024	23,288	0.8925	0.5614	83,833
<i>Service</i>	43,617	26,810	1.0514	0.6463	173,012

An additional issue with the CPS is imputed data. In the CPS interviews, if an answer to a particular question is not received or is inconsistent with other answers, a variety of complex procedures are followed to impute the missing or inconsistent information (see Current Population Survey (2002), chapter 9, for a lengthy summary). As Kambourov and Manovskii (2004) point out, the imputation procedures changed in 1976 and 1989, and at those dates, rates of gross flow across industries and occupations in the publicly released CPS data changed dramatically. In particular, apparent rates of gross flow dropped dramatically with the 1976 change in

imputation procedures, and they increased dramatically with the 1989 change. Rates of gross flow are central to our estimation strategy, so we need to obtain the most reliable measures for such flows possible, and if imputation procedures introduce spurious flows we need to find a way to cleanse the data of these effects. From 1989 on, an indicator variable is recorded in the data to indicate if a portion of a given data record has been imputed. We follow Moscarini and Vella (2003), and perform the following two steps to minimize the imputation problem: (i) We drop data prior to 1976 (for which Moscarini and Vella argue that the imputation procedures were very crude and introduced a great deal of spurious gross flows, and no indicator exists in the data to identify which records are affected by imputation); and (ii) We drop any individual subsequent to 1988 whose data are partially imputed. In principle, this could create a selection bias, but since the sample means for the individuals who have been dropped are very similar to those for the rest of the sample (except for gross flow rates, which are much higher for the dropped workers), it does not appear to be a problem in this case.

Descriptive statistics for the resulting data are shown in Tables 1 and 2. Sample sizes added up across years range from 20,952 for Agriculture/Mining to 140,339 for Manufacturing and 173,012 for Service. Table 1 summarizes gross flows. Each cell of the table shows the average fraction of workers in the row sector who moved to the column sector in any given period; for example, on average, 0.56% of Construction workers in any year moved to Agriculture/Mining. The main diagonal shows the average fraction who did not change sector of employment (that is, m_t^{ii}), so one minus this value is a simple measure of the rate of gross flow. The value on the diagonal varies from 0.9292 for Agriculture/Mining to 0.9720 for Services, implying a rate of gross flow that varies across sectors from 2.8% to 7.1%. Table 2 shows descriptive statistics for wages. Normalized wages (that is, normalized to have a unit mean) averaged across time range from 0.8374 for Agriculture/Mining to 1.0514 for Services.

4 Results.

Before showing estimations, we should point out that we do not attempt to estimate β . This model is not designed to estimate rates of time preference, and although it could be done in principle, in practice it turns out that that one parameter is very poorly identified. Since it is not a parameter of interest for us, and since it is the one parameter for which we have strong information *a priori*, we simply impose $\beta = 0.97$ in all that follows.

Table 3 shows the results from the basic regression. For the simplest implementation of the model, we impose $C^{ij} \equiv C \forall i \neq j$, so that the mean moving cost for any transition from one industry to any other is the same. We will explore specifications that allow the C^{ij} 's to vary shortly. Throughout the table, the data are from 1976 to 2001, and the t-statistics are reported in parentheses.

Panel I shows the results for the full sample with no instruments, which (recalling (8)) amounts to regressing current flows $\ln m_t^{ij} - \ln m_t^{ii}$ on future flows $\ln m_{t+1}^{ij} - \ln m_{t+1}^{ii}$ and the future wage differential $w_{t+1}^j - w_{t+1}^i$ by OLS. Of course, this is likely to be biased, as the residual contains the shock revealed at time $t + 1$, which is likely to be correlated with date- $t + 1$ wages. For this reason, we use as instrumental variables the values of the gross flows and wages lagged twice, which must be uncorrelated with any new information revealed at time $t + 1$. The estimates using the instrumental values are reported in Panel II. Henceforth, unless otherwise stated, all estimates use this instrumental-variables approach.

For the basic specification, estimation with and without instrumental variables produces extremely high estimates of both C and ν , with both parameters highly significant. The instrumental-variables estimate of C in Panel II amounts to approximately thirteen times average annual wage earnings (given our normalization of average wages to unity). The value of ν of 2.897 implies a variance of the idiosyncratic shock equal to 13.8, or a standard deviation of 3.7; of course, the standard deviation of the idiosyncratic moving cost is twice that (since it is the difference between two idiosyncratic shocks). In other words, the mean moving cost between two industries is thirteen times the average wage, but its standard deviation is about seven times the average wage. We will argue in the following section that these estimates are likely to be biased upward and we will present lower estimates following corrections

Table 3: The Basic Regression.

<i>I. Full sample: OLS</i>	
ν	C
4.466** (1.829)	22.065 (1.780)**
<i>II. Full sample with instruments.</i>	
ν	C
2.897*** (2.667)	13.210*** (2.558)
<i>III. Younger workers.</i>	
ν	C
2.385*** (2.346)	10.312** (2.228)
<i>IV. Older workers.</i>	
ν	C
4.220 (1.153)	21.508 (1.149)
<i>V. No college education.</i>	
ν	C
5.665 (0.946)	25.658 (0.922)
<i>VI. Some college education.</i>	
ν	C
3.339** (2.000)	15.539** (1.785)

(T-statistics are in parentheses. One-sided significance: 1-percent***, 5-percent**, 10-percent*.)

for the bias, but these strikingly high figures do convey an important message that is robust to all corrections: *Labor movements in response to a differential in wages are very sluggish. The labor market acts as if it is very costly to change sectors, but at the same time a significant number of workers does so anyway, not in response to differences in wages, but because of unobserved and possibly non-pecuniary factors that are at least as important as wages in workers' decisions.* Later, in the simulations, we will see that the aggregate labor market behavior implied by our estimates is quite realistic, and fits well with some reduced-form regression results in the literature.

Panels III and IV show the results when the estimation is restricted to workers under the age of 45 and workers 45 years old or older, respectively. Once again, the coefficients are highly significant. The difference in results is that the mean and variance of moving costs are substantially higher for older workers (the value of C is about twice as high for the older workers), although it must be noted that the parameters are very imprecisely estimated for the older workers.⁹ This does not reflect a substantially reduced mobility *per se* for older workers (the ratio C/ν is about 5.1 for older workers and 4.3 for younger workers), but rather a much lower responsiveness of mobility to wage differentials.

Panels V and VI show the results when the estimation is restricted to workers with no college education and those with at least one year of college, respectively.¹⁰ Again, the coefficients are significant, although with a much lower level for those with no college. Again, the level of mobility *per se* is not very different between the two groups (the C/ν ratio is about 4.5 for both groups), but estimated mobility barriers are very much higher for those with no college.

5 Possible sources of bias.

There are two notable reasons the very high estimates we have obtained for C and ν may be the result of bias: Sampling error in industry wages and possible misinter-

⁹This may well reflect sample sizes, leading to more noise in the measured wages and gross flows for older workers. There are 518,778 worker-year observations in the sample, of whom 326,918 are young (63% of the total) and 191,860 are older (37% of the total).

¹⁰Workers with some college education comprise 149,329 worker years, or 29% of the total. Workers with no college comprise 369,449 worker years, or 71% of the total.

pretation of mobility rates in the CPS data due to timing issues. There is also the possibility that constraining all C^{ij} values to the the same is a misspecification that generates its own bias. We discuss these in turn.

5.1 Sampling error in wages.

We measure the industry wages w_t^i as the average wage in the industry in the CPS sample. If the sampling error is significant, the industry wage will be measured with noise, resulting in a classical errors-in-variables bias. Given the estimating equation (9), this will lower the estimated value $\frac{\beta}{\nu}$, thus raising the estimated value of ν and thus C . We investigate this possibility in two ways.

First, we re-do the estimation using time-averaged values of the variables. To the extent that the high estimates are driven by serially uncorrelated noise in the measured variables, this should reduce their level. We break the sample into consecutive, non-overlapping five-year segments. For each industry i , we average w_t^i over each segment, and for each i and j we average m_t^{ij} over the segment. The results are reported in Table 4.

Note that although the estimated moving costs are much smaller now, nonetheless C is estimated at eight and a half times average wages and the standard deviation of moving costs equal to eight times annual wages in the benchmark specification of panel II. (They are also much more precisely estimated, with significance at the 1% level for all parameters.) This specification is not useful for policy analysis, since the implied five-year period for each worker reallocation is unrealistically long, but it does make the point that only a portion of the explanation for the high moving costs could plausibly be due to sampling error in wages.

Second, we re-do the regression using wage data from the Bureau of Labor Statistics' Current Employment Surveys (CES) in place of the wage data we have constructed from the CPS. Since the CES is a broad employer-based survey with a large sample size, it is likely to have less of a problem with sampling error in the wages. The industry classifications for the two data sets are not exactly the same, but the nearest

Table 4: The Regression with Time-Averaging.

<i>I. Full sample: OLS</i>	
ν	C
3.587*** (5.924)	10.298*** (4.621)
<i>II. Full sample with instruments.</i>	
ν	C
3.338*** (7.932)	8.477*** (6.035)
<i>III. Younger workers.</i>	
ν	C
3.172*** (6.211)	7.513*** (4.668)
<i>IV. Older workers.</i>	
ν	C
3.680*** (4.156)	12.715*** (4.114)
<i>V. No college education.</i>	
ν	C
5.160*** (2.763)	12.733*** (2.626)
<i>VI. Some college education.</i>	
ν	C
3.177*** (4.948)	8.367*** (3.926)

(T-statistics are in parentheses. One-tailed significance: 1-percent***, 5-percent**, 10-percent*.)

match produces quite similar wage series,¹¹ and very similar regression results.¹²

We thus conclude that the high estimates of C and ν are *not* likely to be artifacts of sampling error in wages.

5.2 Timing and the misinterpretation of flow rates.

Kambourov and Manovskii (2004) point out a difficulty in interpreting flow rates that come out of the March CPS retrospective questions. Respondents are asked their industry and occupation in their longest-held job of the previous year. If the duration of jobs is distributed randomly and respondents remember correctly, on average they will be reporting their employment status as of the middle of the previous year, and thus mobility at a nine-month window (June to March) rather than a twelve-month window. However, if a respondent has had more than one job during that year and recalls the details of the later job more clearly, the later one might incorrectly be reported as the longest job. In this case, the respondent might be reporting details of his or her employment in October, for example, implying that what is being measured is mobility at a six-month window.

Therefore, although it appears superficially to be annual, the mobility measured by the March CPS is something less than annual. Kambourov and Manovskii (2004) point out that, consistent with this, occupational gross flow rates as measured by the March CPS tend to be smaller than those measured from other sources.

We can attempt to correct for this in the following way. Suppose that the gross flow rate we observe is the flow rate over some interval that is K months long, and denote the matrix of gross flow rates thus observed by \tilde{m} . We first convert this into a matrix of monthly gross flows, \hat{m} , by solving the equation $\hat{m}^K = \tilde{m}$, where \hat{m}^K denotes the matrix \hat{m} multiplied by itself K times.¹³ Without loss of clarity, we can

¹¹The correlation between the two wage series is 63% for Agriculture/mining; 91% for Construction; 44% for Manufacturing; 56% for Transportation/Utilities; 61% for Trade; and 55% for Government and other services.

¹²For example, for the OLS regression in Table 3, the point estimate and t-statistic for ν are 4.466 (1.829) for the CPS wage data and 4.237 (2.031) for the CES data respectively. The estimates for C are 22.065 (1.780) for the CPS wage data and 20.921 (2.021) for the CES data respectively.

¹³For example, in the two-industry case, if $K = 2$, the fraction of workers in industry 1 at the beginning of the two-month interval who are in industry 2 at the end of the two-month interval is

denote this matrix as $\tilde{m}^{1/K} = \hat{m}$. Suppose that within a year, the monthly flow rate matrix \hat{m} is constant. Then the year-by-year matrix of flow rates will be given by $m^{ANN} \equiv \hat{m}^{12} = \tilde{m}^{12/K}$, or the \hat{m} matrix multiplied by itself 12 times. We have data on gross flow rates from the National Longitudinal Survey of Youth (NLSY), which we can denote $m_t^{ij,NLSY}$ and which do not suffer from the timing problems just described for the March CPS. We choose K to minimize the following loss function:¹⁴

$$\sum_{i,j,t} ((\tilde{m}_t^{12/K})^{ij} - m_t^{ij,NLSY})^2 \quad (11)$$

for the portion of our sample restricted to younger workers. This results in a value of $K = 5$, implying that the March CPS measures mobility at a five-month horizon. We then replace our measured gross flows \tilde{m} with the annualized gross flows $m^{ANN} = \tilde{m}^{12/K}$ throughout, and perform the estimation again.

As expected, the annualized rates show higher gross flows overall. Table 5 provides a comparison of the original rates of gross flow (meaning $1 - \tilde{m}^{ii}$ for $i = 1, \dots, 6$) with the annualized rates of gross flow (meaning $1 - (m^{ANN})^{ii}$ for $i = 1, \dots, 6$).

Table 5: Rates of gross flow, original and annualized.

	<i>Raw data.</i>	<i>Annualized data.</i>
<i>Agric/Min</i>	0.071	0.161
<i>Const</i>	0.057	0.130
<i>Manuf</i>	0.029	0.068
<i>Trans/Util</i>	0.036	0.083
<i>Trade</i>	0.053	0.122
<i>Service</i>	0.028	0.065

The regression results are as shown in Table 6. The first two columns show results from annualized data as just described, and comparing them with the results in Table 3 shows that, as expected, the values for C and ν are lower for each version of the regression and for each subsample. The only exception is version IV, the regression with

equal to $\hat{m}^{12}\hat{m}^{22} + \hat{m}^{11}\hat{m}^{12}$, which is the product of the first row and the second column of the \hat{m} matrix.

¹⁴To clarify, $(\tilde{m}_t^{12/K})^{ij}$ is the ij element of the matrix $m_t^{ANN} = \tilde{m}_t^{12/K}$.

only older workers; that sample shows a slight increase for the annualized regression, but for neither the raw nor the annualized data are the coefficients significant. For the other cases, the drop in the estimates of C and ν is substantial; for example, the results from the benchmark regression (regression II) imply a mean moving cost of about six and a half times average annual wages (compared with thirteen in Table 3) and a standard deviation of moving costs of five times annual wages (compared with seven times for Table 3). The final two columns of Table 6 show the results of the regression when we apply time-averaging as in the previous subsection to the annualized data. Note the sharp reduction in estimates of C and ν , with the benchmark regression showing average moving costs at just below three times average annual wages. Note also that the coefficients for older workers are now significant, and that throughout Table 6, the patterns of the basic regression are preserved: Instrumental variables lower the estimated moving cost parameters, and moving costs are lower for younger workers and workers with some college (the exception being the last two rows of the last column).

Overall, we find strong indications that the timing problem due to the nature of CPS questions does bias our estimates for C and ν upward substantially, but correcting for this still leaves large values for moving costs, with C never falling substantially below three times average annual wages.

5.3 Mispecification of moving costs.

A last possible source of bias comes from the fact that we have imposed uniform moving costs for all sectors, so that $C^{ij} = C\forall i, j$. Degrees-of-freedom concerns prevent us from estimating the full set of C^{ij} parameters without restriction, but we have also estimated the model with a slightly richer specification allowing for sector-specific “entry costs.” In this approach, $C^{ij} = c^j$ for $i = 1, \dots, 6$. Table 7 shows the results of this regression with annualized data.

Compared with the benchmark regression (II) from Table 6, we find that most sectors exhibit lower entry costs (mostly between 4 and 5, compared to 6.6 for Table 6), but Sector 4, Transport, Communications and Utility exhibits substantially higher entry costs — more than eight times average annual wages. This reflects the fact that Sector 4’s wages are relatively high but few workers wind up in this sector. For

Table 6: The Regression with Annualized Flow Rates.

<i>Annualized flows.</i>		<i>Annualized with time-averaging.</i>	
<i>I. Full sample: OLS</i>			
ν	C	ν	C
2.178*** (3.755)	8.799*** (3.264)	1.857*** (9.982)	3.452*** (4.740)
<i>II. Full sample with instruments.</i>			
ν	C	ν	C
1.884*** (3.846)	6.565*** (3.381)	1.899*** (11.614)	2.837*** (4.300)
<i>III. Younger workers.</i>			
ν	C	ν	C
1.516*** (3.808)	4.988*** (3.011)	1.756*** (9.358)	2.296*** (3.313)
<i>IV. Older workers.</i>			
ν	C	ν	C
4.605 (0.504)	22.962 (0.502)	2.285*** (6.812)	5.640*** (4.893)
<i>V. No college education.</i>			
ν	C	ν	C
2.926* (1.405)	9.382 (1.261)	2.856*** (4.006)	4.037*** (2.351)
<i>VI. Some college education.</i>			
ν	C	ν	C
1.848** (1.938)	7.689* (1.499)	1.890*** (7.447)	2.969*** (2.984)

(T-statistics are in parentheses. One-tailed significance: 1-percent***, 5-percent**, 10-percent*.)

Table 7: Sector-specific entry costs.

	<i>Estimate.</i>	<i>t-statistic.</i>
ν	1.512***	(3.622)
C^1 (<i>Agriculture/Mining</i>)	4.124	(1.053)
C^2 (<i>Construction</i>)	4.899**	(1.735)
C^3 (<i>Manufacturing</i>)	4.994***	(3.165)
C^4 (<i>Transportation, Communication and Utilities</i>)	8.311***	(3.044)
C^5 (<i>Trade</i>)	3.703*	(1.437)
C^6 (<i>Government and Other Services</i>)	5.589***	(3.791)

(Full sample, with instruments. Gross flows are annualized as in Table 6. One-tail significance: 1-percent***, 5-percent**, 10-percent*.)

example, from Table 2, note that wages for both Services and Sector 4 are about five percent above the whole-sample average, but Sector 4 has fewer than a third as many workers. Alternatively, note from a comparison of the fourth and sixth columns of Table 1 that the rate of flow from each other sector into sector 4 is in each case around one third the rate of flow into sector 6, despite that fact that on average the wages in these two sectors are about identical. This indicates some implicit obstacle (or disutility) to entering sector 4 compared to other sectors, thus implying a high value of C^4 .

We can conclude that a portion of the reason for the high values estimated for C in the earlier regressions is the need to account for the unusually low flows into Transport, Communications and Utilities. However, even when this effect is separated out, most of the other sectoral moving costs are still high — at least four times average annual wages.

6 Simulation: A Sudden Trade Liberalization.

Now, we use the estimates to study the effect of a hypothetical trade shock through simulations. We assume that each of the six sectors has a Constant Elasticity of Substitution production function, with labor and unmodelled sector-specific capital

as inputs. Thus, for our purposes, the production function for sector i is given by:

$$y_t^i = \psi^i \left(\alpha^i (L_t^i)^{\rho^i} + (1 - \alpha^i) (K^i)^{\rho^i} \right)^{\frac{1}{\rho^i}}, \quad (12)$$

where y_t^i is the output for sector i in period t , K^i is sector- i 's capital stock, and $\alpha^i > 0$, $\rho^i < 1$, and $\psi^i > 0$ are parameters. Given the number of free parameters and our treatment of capital as fixed,¹⁵ we can without loss of generality set $K^i = 1 \forall i$. This implies that the wages are given by:

$$w_t^i = p_t^i \alpha^i \psi^i (L_t^i)^{\rho^i - 1} \left(\alpha^i (L_t^i)^{\rho^i} + (1 - \alpha^i) \right)^{\frac{1 - \rho^i}{\rho^i}}, \quad (13)$$

where p_t^i is the domestic price of the output of sector i . We set the values α^i , ρ^i , and ψ^i to minimize a loss function given our assumptions on prices (see below). Specifically, for any set of parameter values, we can compute the predicted wage for each sector and that sector's predicted share of GDP using (13) and (12) together with empirical employment levels for each sector and our assumptions about prices as described below. The loss function is then the sum across sectors of the square of each sector's predicted wage minus mean wage in the data, plus the square of the sector's predicted minus its actual share of GDP. In addition, we assume that all workers have identical Cobb-Douglas preferences, using consumption shares from the BLS consumer price index calculations for the consumption weights. The parameter values that result from this procedure are summarized in Table 8.

The moving-cost parameters used are found in our preferred specification, the annualized-flow-rate approach of the first two columns of Table 6, using the full sample with instruments.

Then, to provide a simple trade shock, we assume the following: (i) Units are chosen so that the domestic price of each good at date $t = -1$ is unity. (Given our available free parameters, this is without loss of generality.) (ii) There are no tariffs on any sector aside from manufacturing, at any date. (iii) The world price of manufacturing output is 0.7 at each date. The world price of all other tradeable goods

¹⁵We assume that capital is fixed in order to focus on the workers' problem and to keep the model manageable. Of course, capital should also be expected to adjust to trade liberalization, and that should also be expected to affect wages. We have experimented with simple simulations with perfect capital mobility, obtaining similar welfare results but sharper movements in wages. We defer a full treatment of this issue to future work.

Table 8: Parameters for Simulation.

	α^i	ρ^i	ψ^i	Consumer expenditure share.	Pre-liberalization domestic price.	World price.
<i>Agric/Min</i>	0.691	0.6828	0.6733	0.07	1	1
<i>Const</i>	0.6544	0.4924	0.7653	0.3	1	1*
<i>Manuf</i>	0.3224	0.3553	1.6965	0.3	1	0.7
<i>Trans/Util</i>	0.5721	0.5664	1.0393	0.08	1	1*
<i>Trade</i>	0.5714	0.445	0.9125	0	1	1*
<i>Service</i>	0.3418	0.5576	2.2135	0.25	1	1

(Note:* Under the second simulation specification, the sectors marked with an asterisk are non-traded, so they have no world price.)

is equal to unity at each date. (iv) There is initially a specific tariff on manufactures at the level 0.3 per unit, so that the domestic price of manufactures is equal to unity. (v) Initially, this tariff is expected to be permanent, and the economy is in the steady state with that expectation. (vi) At date $t = -1$, however, after that period's moving decisions have been made, the government announces that the tariff will be removed beginning period $t = 0$ (so that the domestic price of manufactures will fall from unity to 0.7 at that date), and that this liberalization will be permanent.

Thus, we simulate a sudden liberalization of the manufacturing sector. We compute the perfect-foresight path of adjustment following the liberalization announcement, until the economy has effectively reached the new steady state. This requires that each worker, taking the time path of wages in all sectors as given, optimally decides at each date whether or not to switch sectors, taking into account that worker's own idiosyncratic shocks. This induces a time-path for the allocation of workers, and therefore the time-path of wages, since the wage in each sector at each date is determined by market clearing from (13) given the number of workers currently in the sector. Of course, the time path of wages so generated must be the same as the time-path each worker expects. It is shown in Cameron, Chaudhuri and McLaren (2007)

that the equilibrium exists and is unique.¹⁶ The computation method is described at length in Artuç, Chaudhuri and McLaren (forthcoming).

We present two versions of the simulation. In the first, all goods are assumed to be traded, so all output prices are exogenous. In the second, some sectors are non-traded, and so their prices are determined as part of the equilibrium.

6.1 Specification I: All output is tradeable.

The results from the simulation with all goods tradeable can be seen in Figure 2, which plots the fraction of the labor force in each of the six sectors at each date, and Figure 3, which plots the time-path of wages. Figure 4 shows the average payoff V_t^i to being a worker in sector i at time t .

It is clear from Figure 2 that the employment share of manufacturing drops sharply as a result of the liberalization, from an old steady state value of 25% to a new steady state value of 16%, with corresponding modest gains to all other sectors. This transition is substantially complete within 8 years. The loss of manufacturing's share is of course not surprising given that manufacturing has lost its protection. It is also clear from Figure 3 that real wages in manufacturing fall as a result of the liberalization, from an old steady-state value of 1.06 to a new steady-state value of 1.03, and with corresponding modest gains to all other sectors due to the drop in consumer prices. It should be emphasized that in neither steady state are the wages equalized across sectors. This is a basic feature of the gross-flows model, absent in net-flows models (see Cameron, Chaudhuri and McLaren (2007) for an extended discussion).

Figure 3 shows, in addition, that each sector sees a non-monotonic path for real wages. The real wage in manufacturing overshoots its long-run value, with an initial drop of 22% and a new steady state just 2.45% below the original steady state. This overshooting occurs because after the sudden shock of the drop in domestic manufacturing prices, workers begin to move out of the sector, moving up and to the left along the sector's demand-for-labor curve and gradually bringing wages up.¹⁷

¹⁶Strictly speaking, the proof there applies to the case with all goods traded, but it can be extended mechanically to the case with non-traded goods under free trade.

¹⁷In principle, it is possible that this process could continue so that real wages in the liberalizing

Similar overshooting occurs in each of the other sectors, in the opposite direction, for parallel reasons.

Note that at each date following the liberalization announcement, the real wage in manufacturing is below what it was in the old steady state. It would be tempting to conclude that for this reason manufacturing workers must be worse off because of the liberalization. However, that is not true. As can be seen in Figure 4, which plots $V^i(L_t, s_t)$ from equation (3), *all* workers see a rise in their expected discounted lifetime utilities at the time of the announcement, *including* manufacturing workers.¹⁸ The reason is the presence of gross flows. Each manufacturing worker understands that, because of the liberalization, manufacturing wages are permanently lower but real wages in all other sectors are permanently higher. Further, there is in each period a positive probability that the manufacturing worker will choose to move to one of those other sectors and enjoy those higher wages. Taking into account these probabilities, the manufacturing worker considers himself/herself lucky to be hit with the liberalization.

Put differently, the liberalization lowers the wages in the manufacturing sector but *raises the option value* to workers in the sector by more than enough to compensate. Thus, in this case, despite the estimation of extremely high moving costs, the model predicts that even workers in import-competing sectors will welcome liberalization. This underlines the crucial importance of gross flows in welfare analysis.¹⁹

Finally, we can compute trade flows from the simulation. At each date with free trade, GDP can be computed from the labor allocation and production functions; from the utility function, we can compute consumption of each sector's output, and subtract the quantity produced to derive net imports. In the initial steady state with the tariff, calculation is slightly more complicated, as we need to add tariff revenue to income and compute consumption with domestic prices instead of world prices.

sector could rise past their original value, and wind up higher in the new steady-state than in the old, but that does not happen in this case. See Artuğ, Chaudhuri and McLaren (forthcoming) for examples.

¹⁸The rise in lifetime utility is between 4.5% and 5% in non-manufacturing sectors and 1.7% in manufacturing.

¹⁹We also have simulated exactly the same policy experiment with the estimates from the sector-specific "entry-cost" specification of Table 7. The results are qualitatively and quantitatively very similar, with rather higher wages overall for Transportation, Communications and Utilities.

Figure 5 shows the result for manufacturing output. At the date of liberalization, manufacturing consumption jumps up because of the abrupt drop in the domestic price of manufactures. Thereafter, it trends upward slightly because of increases in GDP as the economy reallocates its labor. Throughout, domestic production of manufactures falls, as workers leave the sector. Note that this implies that following the liberalization, manufactures imports continue to rise for several years, even as manufacturing wages rise. Thus, if one regressed manufacturing wages on import penetration and the date $t = -1$ was not part of the data, one would find a positive coefficient, while including the date $t = -1$ would change the sign to negative. This suggests that regressions that relate current manufacturing wages to current import penetration measures, such as are explored in Freeman and Katz (1991) and Kletzer (2002), need to be interpreted with great care.

Another point can be seen regarding the interpretation of reduced-form regressions. Revenga (1992), in her simplest specification, regresses changes in log industry wages and employment for the years 1981-5 on changes in log industry import prices for the same period, and finds an elasticity of 1.74 for employment and 0.40 for wages. An analogous wage ‘elasticity’ can be computed from our simulation, as

$$\left(\frac{w_3^{manuf} - w_{-1}^{manuf}}{w_{-1}^{manuf}} \right) \left(\frac{p_{-1}^{manuf}}{p_3^{manuf} - p_{-1}^{manuf}} \right) = -\frac{1}{0.3} \left(\frac{w_3^{manuf} - w_{-1}^{manuf}}{w_{-1}^{manuf}} \right),$$

where w_t^{manuf} and p_t^{manuf} denote the period- t manufacturing wage and domestic output price, respectively. The employment ‘elasticity’ is analogous. The employment ‘elasticity’ from our simulation is 0.88, and the wage ‘elasticity’ is 0.38. Thus, the orders of magnitude are similar to the Revenga elasticities and the signs match up – despite the tremendous differences in method. However, as pointed out above, the welfare implications of our dynamic model with gross flows are the opposite of the implications of Revenga’s static model without gross flows. In her interpretation, workers in a sector whose import price falls are hurt, while in our simulation they are not. The option value effects that are key to our analysis have no possible role in the static approach.

Thus, although our model generates aggregate behavior broadly similar to what is found in some reduced-form regression results, the welfare implications are extremely different.

6.2 Specification II: Non-traded sectors.

In our second simulation specification, Construction, Transportation/Utilities and Trade are taken to be non-traded.²⁰ Thus, their prices are endogenous, and adjust so that the quantity produced in each of those sectors at each date (as determined from the production function and the number of workers in the sector at that date) is equal to the quantity demanded, given GDP and tradeable-goods prices. The endogenous domestic prices are shown in Figure 11, which can be contrasted with the exogenous prices of Figure 6. Figure 7 shows the reallocation of labor. Compared with Figure 2, the pattern is similar, but the non-traded sectors expand less while the traded sectors expand more. This is because the suddenly less expensive manufactured goods cause consumer expenditure to switch toward manufacturing and away from non-traded goods, effectively shifting the demand curve for non-traded goods sharply downward at the date of the liberalization. This is reflected in the sudden drop in non-traded prices exhibited in the time-plot of domestic prices shown in Figure 11. As a result, the movement in wages in non-traded sectors is much less sharp in Figure 8 than in Figure 3. Figure 10 shows that, once again, liberalization increases the imports of manufactures, with an initial jump and gradual adjustment over the following several years.

The main point is unchanged. Real wages in manufacturing fall sharply and never recover, as shown in Figure 8, but workers in manufacturing benefit from the liberalization along with workers in all of the other sectors, as shown in Figure 9. Once again, the explanation is enhanced option value for workers in the liberalized sector.

7 Conclusion.

We have presented a dynamic, rational expectations model of labor adjustment to trade shocks, which is, through Euler-equation techniques, easy to estimate econo-

²⁰This division is, of course, to some degree arbitrary. It is difficult to argue that Services should be classified as non-traded, since services trade has occupied much attention and created much controversy at the WTO. On the other hand, the ‘Trade’ sector is, of course, mainly domestic wholesale and retail trade, and thus not *internationally* traded, which is the meaning of ‘non-traded’ here.

metrically, yielding structural parameters. It is then easy to simulate to study policy. Among our findings are the following.

(i) Since gross flows of workers across industries are substantial but do not respond much to intersectoral wage differences, both the mean and the standard deviation of workers' moving costs implied by the model are large – several times an average workers' annual earnings, in fact.

(ii) Because of this, the model predicts somewhat sluggish reallocation of workers following a trade liberalization. In our simulation of the elimination of a 30% tariff on manufacturing, 95% of the reallocation is completed in 8 years.

(iii) This implies sharp movement of wages in response to the liberalization, with the short-run response overshooting the long-run response by a wide margin.

(iv) Option value, not previously part of the discussion in analysis of trade policy, matters a great deal in evaluating the welfare effects of trade liberalization. In our simulation, the manufacturing wage falls both in the short run and in the long run, but manufacturing workers are better off than before the liberalization because of their enhanced option value. This echoes some findings by Magee, Davidson and Matusz (2005) on patterns in political contributions.

(v) Although our model generates aggregate behavior broadly similar to what is found in some reduced-form regression results, the welfare implications are extremely different.

A Appendix: Derivation of Equilibrium Conditions with the Extreme Value Distribution.

A.1 Overview of the Derivation.

The cumulative distribution function for the extreme value distribution with zero mean is given by:

$$F(\varepsilon) = \exp(-\exp(-\varepsilon/\nu - \gamma)),$$

where $\gamma \cong 0.5772$ is Euler's constant. The associated density function is:

$$f(\varepsilon) = (1/\nu) \exp(-\varepsilon/\nu - \gamma - \exp(-\varepsilon/\nu - \gamma)).$$

In the following subsection we will derive equation (6), which relates gross flow rates to the value function. In the subsection after that we will derive the form for the option-value function reported in (7).

A.2 The m^{ij} function.

The gross flow of workers from i to j at date t , m_t^{ij} , is equal to the probability that a given i -worker will switch to j at date t , or the probability that, for an i -worker, utility $w_t^i + \varepsilon_t^j + \beta E_t[V^j(L_{t+1}, s_{t+1})] - C^{ij}$ will be higher for a move to j than for any of the other $n - 1$ options. In other words, from (2),

$$m_t^{ij} = \text{Prob}_{\varepsilon_t} [\bar{\varepsilon}_t^{ij} + \varepsilon_t^j \geq \bar{\varepsilon}_t^{ik} + \varepsilon_t^k \text{ for } k = 1, \dots, n].$$

Suppressing the time subscript, this can be written:

$$m^{ij} = \int_{-\infty}^{\infty} f(\varepsilon^j) \prod_{k \neq j} F(\varepsilon^j + \bar{\varepsilon}^{ij} - \bar{\varepsilon}^{ik}) d\varepsilon^j.$$

Define, for convenience: $x \equiv \varepsilon^j / \nu + \gamma$, $z^j \equiv \bar{\varepsilon}^{ij}$, $\bar{\varepsilon}^{ik} = z^k$, and $\lambda \equiv \log\left(\frac{\sum_{k=1}^n \exp(z^k / \nu)}{\exp(z^j / \nu)}\right)$.

Then the expression for gross flows can be rewritten:

$$\begin{aligned} m^{ij} &= \frac{1}{\nu} \int \exp(-\varepsilon^j / \nu - \gamma - \exp(-\varepsilon^j / \nu - \gamma)) \prod_{k \neq j} \exp(-\exp(-[\varepsilon^j + \bar{\varepsilon}^{ij} - \bar{\varepsilon}^{ik}] / \nu - \gamma)) d\varepsilon^j \\ &= \frac{1}{\nu} \int \exp(-\varepsilon^j / \nu - \gamma - \exp(-\varepsilon^j / \nu - \gamma)) \exp(-\sum_{k \neq j} \exp(-[\varepsilon^j + \bar{\varepsilon}^{ij} - \bar{\varepsilon}^{ik}] / \nu - \gamma)) d\varepsilon^j \\ &= \frac{1}{\nu} \int \exp(-\varepsilon^j / \nu - \gamma) \exp(-\sum_{k=1}^n \exp(-[\varepsilon^j + \bar{\varepsilon}^{ij} - \bar{\varepsilon}^{ik}] / \nu - \gamma)) d\varepsilon^j \\ &= \frac{1}{\nu} \int \exp [(-\varepsilon^j / \nu - \gamma) - \sum_{k=1}^n \exp(-[\varepsilon^j + \bar{\varepsilon}^{ij} - \bar{\varepsilon}^{ik}] / \nu - \gamma)] d\varepsilon^j \\ &= \frac{1}{\nu} \int \exp [(-\varepsilon^j / \nu - \gamma) - \exp((- \varepsilon^j / \nu - \gamma)) \sum_{k=1}^n \exp(-[z^j - z^k] / \nu)] d\varepsilon^j \\ &= \frac{1}{\nu} \int \exp [(-\varepsilon^j / \nu - \gamma) - \exp((- \varepsilon^j / \nu - \gamma)) (\sum_{k=1}^n \exp(z^k / \nu)) / \exp(z^j / \nu)] d\varepsilon^j \\ &= \int \exp(-x - \exp(-(x - \lambda))) dx. \end{aligned}$$

This again can be rewritten:

$$m^{ij} = \exp(-\lambda) \int \exp(-(x - \lambda) - \exp(-(x - \lambda))) dx.$$

Now set $y = x - \lambda$. Noting that the antiderivative of

$$\exp(-y - \exp(-y))$$

is

$$\exp(-\exp(-y)),$$

we can derive:

$$\begin{aligned} m^{ij} &= \exp(-\lambda) \int \exp(-y - \exp(-y)) dy \\ &= \exp(-\lambda) \\ &= \frac{\exp(z^j/\nu)}{\sum_{k=1}^n \exp(z^k/\nu)} \\ &= \frac{\exp(\bar{\varepsilon}^{ij}/\nu)}{\sum_{k=1}^n \exp(\bar{\varepsilon}^{ik}/\nu)}. \end{aligned}$$

Given that $\bar{\varepsilon}^{ii} \equiv 0$, this yields (6).

A.3 The Option-Value Function.

Define:

$$\begin{aligned} \Psi^{ij} &\equiv \int_{-\infty}^{\infty} (\varepsilon^j - C^{ij}) f(\varepsilon^j) \prod_{j \neq k} F(\varepsilon^j + \bar{\varepsilon}^{ij} - \bar{\varepsilon}^{ik}) d\varepsilon^j \\ &= \frac{1}{\nu} \int (\varepsilon^j - C^{ij}) \exp(-\varepsilon^j/\nu - \gamma - \exp(-\varepsilon^j/\nu - \gamma)) \prod_{k \neq j} \exp(-\exp(-[\varepsilon^j + \bar{\varepsilon}^{ij} - \bar{\varepsilon}^{ik}]/\nu - \gamma)) d\varepsilon^j \end{aligned}$$

Going through the steps of Subsection (A.2), we find:

$$\begin{aligned}
\Psi^{ij} &= \int (\nu(x - \gamma) - C^{ij}) \exp(-x - \exp(-(x - \lambda))) dx \\
&= (-C^{ij} - \nu\gamma) \exp(-\lambda) + \nu \int x \exp(-x - \exp(-(x - \lambda))) dx \\
&= (-C^{ij} - \nu\gamma) \exp(-\lambda) + \nu \exp(-\lambda) \int x \exp(-x + \lambda - \exp(-(x - \lambda))) dx
\end{aligned}$$

We know that $\exp(-\lambda) = m^{ij}$ from the previous derivation. Substituting this in:

$$\begin{aligned}
\Psi^{ij} &= (-C^{ij} - \nu\gamma)m^{ij} + \nu m^{ij} \int x \exp(-x + \lambda - \exp(-(x - \lambda))) dx \\
&= (-C^{ij} - \nu\gamma)m^{ij} + \nu m^{ij} \int x \exp(-x + \lambda - \exp(-(x - \lambda))) dx \\
&\quad + \nu m^{ij} \int \lambda \exp(-x + \lambda - \exp(-(x - \lambda))) dx \\
&\quad - \nu m^{ij} \int \lambda \exp(-x + \lambda - \exp(-(x - \lambda))) dx \\
&= (-C^{ij} - \nu\gamma)m^{ij} + \nu m^{ij} \int (x - \lambda) \exp(-x + \lambda - \exp(-(x - \lambda))) dx \\
&\quad + \nu m^{ij} \int \lambda \exp(-x + \lambda - \exp(-(x - \lambda))) dx \\
&= (-C^{ij} - \nu\gamma)m^{ij} + \nu m^{ij} \int y \exp(-y - \exp(-y)) dy + \nu \lambda m^{ij} \int \exp(-y - \exp(-y)) dy \\
&= (-C^{ij} - \nu\gamma)m^{ij} + \nu m^{ij} \int y \exp(-y - \exp(-y)) dy + \nu \lambda m^{ij}.
\end{aligned}$$

Noting that $\int y \exp(-y - \exp(-y)) dy = \gamma$ (Euler's constant) (Patel, Kapadia and Owen (1976, p. 35)), we can simplify:

$$\begin{aligned}
\Psi^{ij} &= (-C^{ij} - \nu\gamma)m^{ij} + \nu \lambda m^{ij} + \nu \gamma m^{ij} \\
&= -C^{ij} m^{ij} - \nu \log(m^{ij}) m^{ij} \\
&= m^{ij} (-C^{ij} - \nu \log(m^{ij})).
\end{aligned}$$

Adding this up across possible destinations j , note that the utility of a worker in i is equal to:

$$\begin{aligned}
V_{t+1}^i &= w_{t+1}^i + \sum_{j=1}^n (\Psi_t^{ij} + \beta m_t^{ij} V_{t+1}^j) \\
&= w_{t+1}^i + \sum_{j=1}^n [m_t^{ij} (-\nu \log(m_t^{ij}) - C^{ij} + \beta V_{t+1}^j)] \\
&= w_{t+1}^i + \sum_{j=1}^n [m_t^{ij} (-\nu \log(m_t^{ij}) - C^{ij} + \beta(V_{t+1}^j - V_{t+1}^i))] + \beta V_{t+1}^i \\
&= w_{t+1}^i + \sum_{j=1}^n [m_t^{ij} (\bar{\varepsilon}_t^{ij} - \nu \log(m_t^{ij}))] + \beta V_{t+1}^i.
\end{aligned}$$

Now, recall from Subsection (A.2) above that $\log(m^{ij}) = \bar{\varepsilon}_t^{ij}/\nu - \log(\sum_{k=1}^n \exp(\bar{\varepsilon}^{ik}/\nu))$. This yields:

$$\begin{aligned}
V_{t+1}^i &= w_{t+1}^i + \sum_{j=1}^n \left[m_t^{ij} \left(\nu \log \left(\sum_{k=1}^n \exp(\bar{\varepsilon}^{ik}/\nu) \right) \right) \right] + \beta V_{t+1}^i \\
&= w_{t+1}^i + \nu \log \left(\sum_{k=1}^n \exp(\bar{\varepsilon}^{ik}/\nu) \right) + \beta V_{t+1}^i.
\end{aligned}$$

This implies that the option value $\Omega(\bar{\varepsilon}^i)$ can be written as:

$$\Omega(\bar{\varepsilon}^i) = \nu \log \left(\sum_{k=1}^n \exp(\bar{\varepsilon}^{ik}/\nu) \right).$$

Alternatively, recalling that $\bar{\varepsilon}^{ii} = 0$, we have:

$$\begin{aligned}
\log(m^{ii}) &= 0 - \log \left(\sum_{k=1}^n \exp(\bar{\varepsilon}^{ik}/\nu) \right) \\
&= -\log \left(\sum_{k=1}^n \exp(\bar{\varepsilon}^{ik}/\nu) \right),
\end{aligned}$$

so in equilibrium

$$\Omega(\bar{\varepsilon}^i) = -\nu \log(m^{ii}).$$

This, then, is (7).

References

- [1] Artuç, Erhan (2006). “Intergenerational Effects of Trade Liberalization.” Mimeo: Koç University.
- [2] Artuç, Erhan, Shubham Chaudhuri, and John McLaren (forthcoming). “Delay and Dynamics in Labor Market Adjustment: Simulation Results.” Forthcoming, *Journal of International Economics*.
- [3] Bowlus, Audra J., and George R. Neumann,(2006). “The Job Ladder” In H. Bunzel, B.J. Christensen, G. Neumann, and J-M Robin (eds.), *Structural Models on Wage and Employment Dynamics*. New York, NY: Elsevier, pp. 217-235.
- [4] Cameron, Stephen, Shubham Chaudhuri, and John McLaren (2007). “Trade Shocks and Labor Adjustment: Theory.” NBER Working Paper #13463 (October).
- [5] Current Population Survey (2000). *Technical Paper 63: Design and Methodology*. Washington, D.C.
- [6] Davidson, Carl and Steven J. Matusz (2001). “On Adjustment Costs.” Michigan State University Working Paper.
- [7] Davidson, Carl, Lawrence Martin and Steven J. Matusz (1999). “Trade and Search Generated Unemployment.” *Journal of International Economics* 48:2 (August), pp. 271-99.
- [8] Davis, Steven J. and John Haltiwanger (1999). “Sectoral Job Creation and Destruction Responses to Oil Price Changes.” NBER Working Paper 7095 (April).
- [9] Freeman, Richard B., and Lawrence F. Katz (1991). “Industrial Wage and Employment Determination in an Open Economy,” in John M. Abowd and Richard Freeman, eds., *Immigration, Trade and the Labor Market* (Chicago: University of Chicago Press, 1991), pp. 235-59.
- [10] Hanson, Gordon H. and Matthew J. Slaughter (1999). “The Rybczynski Theorem, Factor-Price Equalization, and Immigration: Evidence from U.S. States.” NBER Working Paper 7074 (April).

- [11] Hosios, Arthur J. (1990). "Factor Market Search and the Structure of Simple General Equilibrium Models," *The Journal of Political Economy*, 98:2 (April), pp. 325-355.
- [12] Kambourov, Gueorgui (2006). "Labor Market Regulations and the Sectoral Reallocation of Workers: The Case of Trade Reforms."
- [13] Kambourov, Gueorgui and Iourii Manovskii (2004). "A Cautionary Note on Using (March) CPS Data to Study Worker Mobility." Working paper, University of Toronto.
- [14] Keane, Michael P., and Wolpin, Kenneth I. (1997). The Career Decisions of Young Men, *Journal of Political Economy*, vol. 105:3.
- [15] Kennan, John and James R. Walker (2003). "The Effect of Expected Income on Individual Migration Decisions." NBER Working Paper No. 9585.
- [16] Kletzer, Lori G. (2002). *Imports, Exports and Jobs: What Does Trade Mean for Employment and Job Loss?* Kalamazoo, Michigan: W. E. Upjohn Institute for Employment Research.
- [17] Lawrence, Robert Z. and Matthew J. Slaughter (1993). International Trade and American Wages in the 1980s: Giant Sucking Sound or Small Hiccup? *Brookings Papers on Economic Activity (Microeconomics 1993)*.
- [18] Lee, Donghoon (2005). An Estimable Dynamic General Equilibrium Model of Work, Schooling and Occupational Choice. *International Economic Review*, 46 (February), pp. 1-34.
- [19] Magee, Christopher S.P., Carl Davidson, and Steven J. Matusz (2005). "Trade, turnover, and tithing," *Journal of International Economics* 66, pp. 157-176
- [20] Moscarini, Giuseppe and Kaj Thomsson (2006). "Occupational and Job Mobility in the US." Working paper, Yale.
- [21] Moscarini, Giuseppe and Francis Vella (2003). "Aggregate Worker Reallocation and Occupational Mobility in the United States: 1976-2000." Mimeo, Yale University.

- [22] Patel, Jagdish, C. H. Kapadia, and D. B. Owen (1976). *Handbook of Statistical Distributions*, New York: Marcel Dekker, Inc.
- [23] Pavcnik, Nina, Orazio Attanasio and Pinelopi Goldberg (2004). "Trade Reforms and Income Inequality in Colombia." *Journal of Development Economics*, 74 (August), pp. 331-366.
- [24] Revenga, Ana L. (1992). "Exporting Jobs?: The Impact of Import Competition on Employment and Wages in U.S. Manufacturing," *The Quarterly Journal of Economics* 107:1. (February), pp. 255-284.
- [25] Robertson, Raymond and Donald H. Dutkowsky (2002). "Labor adjustment costs in a destination country: the case of Mexico." *Journal of Development Economics* 67:1 (February), pp. 29-54.
- [26] Slaughter, Matthew J. (1998). "International Trade and Labor-Market Outcomes," *Economic Journal*, 108:450 (September), pp. 1452-1462.
- [27] Utar, Hâle (2007). "Import Competition and Employment Dynamics." Working paper, University of Colorado at Boulder.
- [28] United States Bureau of the Census, (1992) "Geographic Mobility: March 1987 to March 1991." U.S. Government Printing Office, Washington D.C.

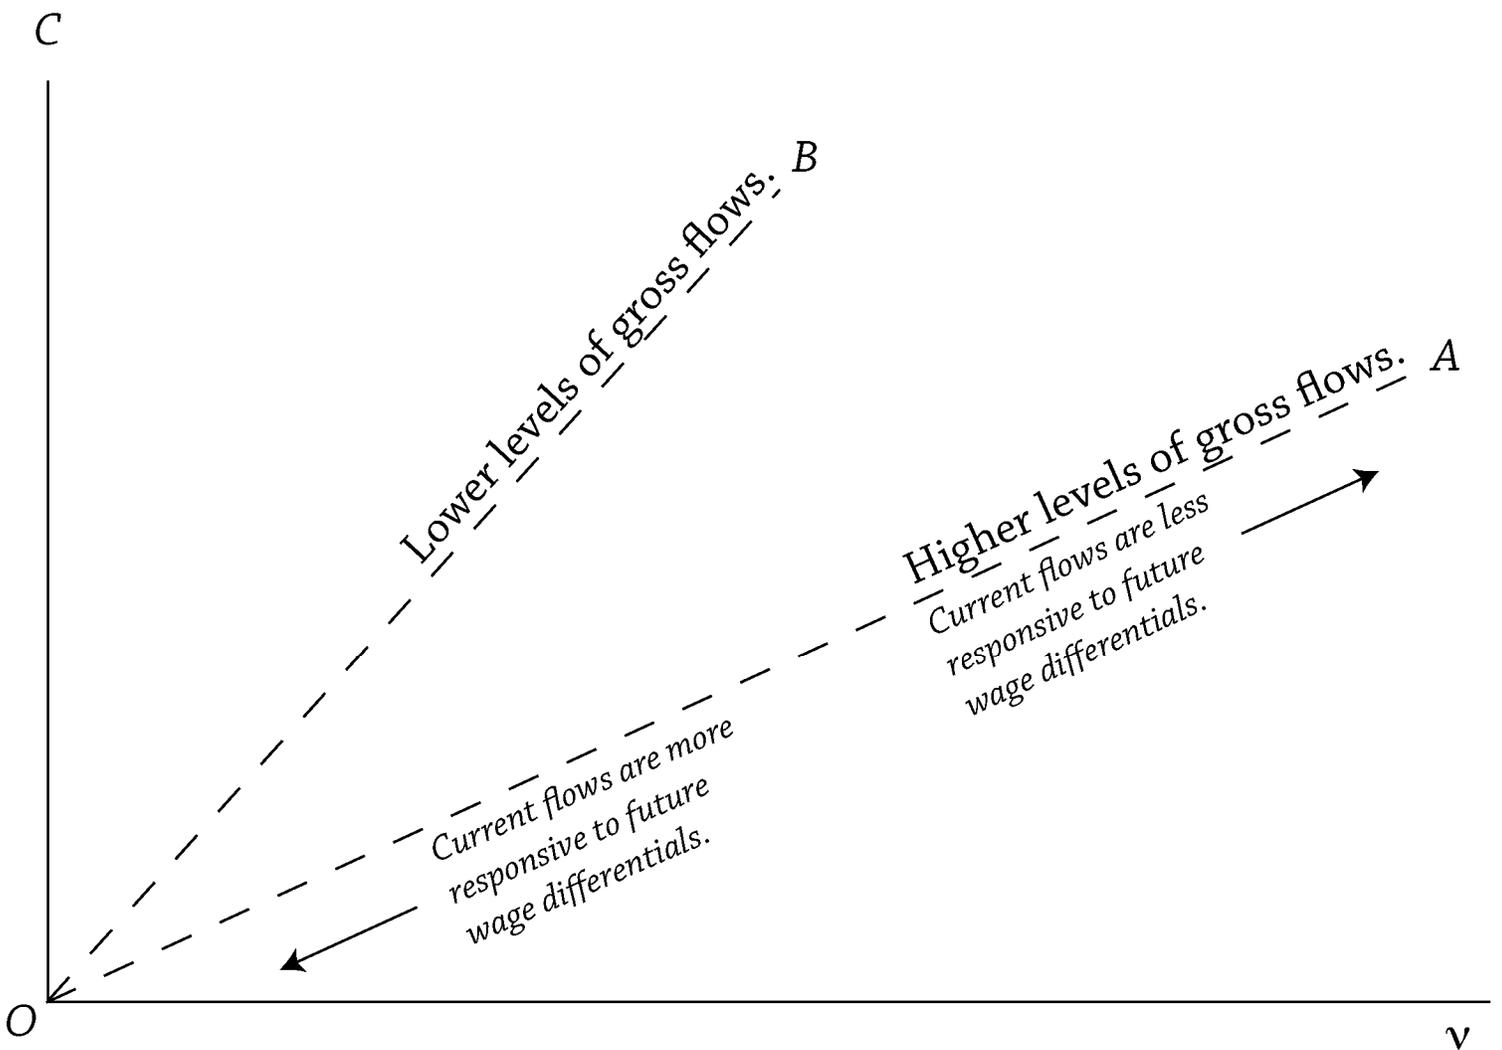


Figure 1: How the mean and variance of moving costs are identified.

Figure 2: Simulated Trade Liberalization I - Labor Allocation

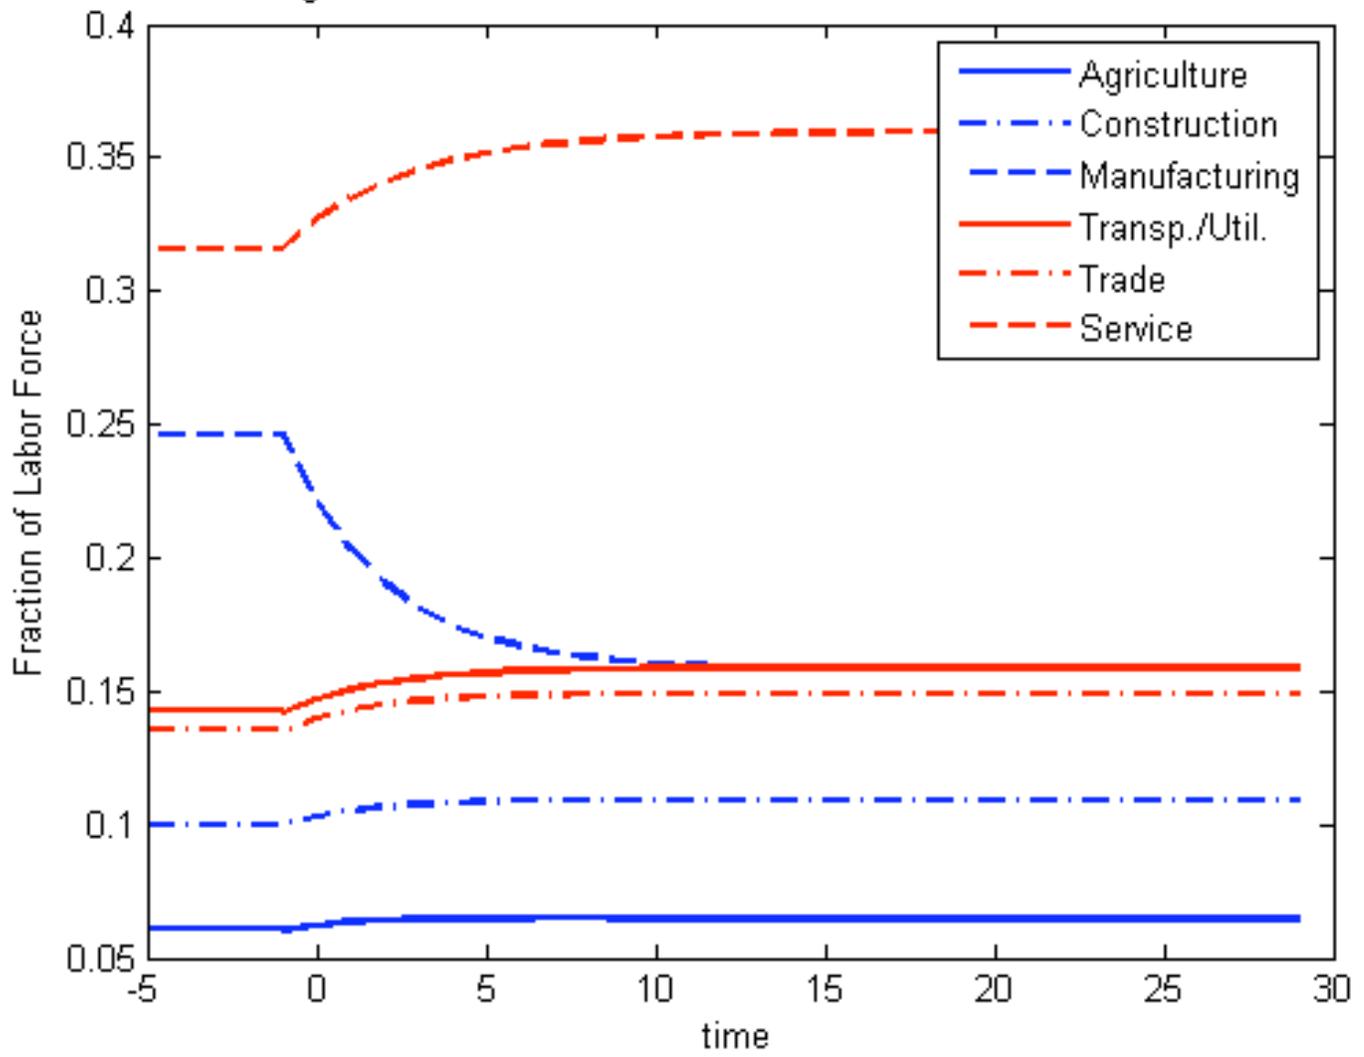


Figure 3: Simulated Trade Liberalization I - Wages

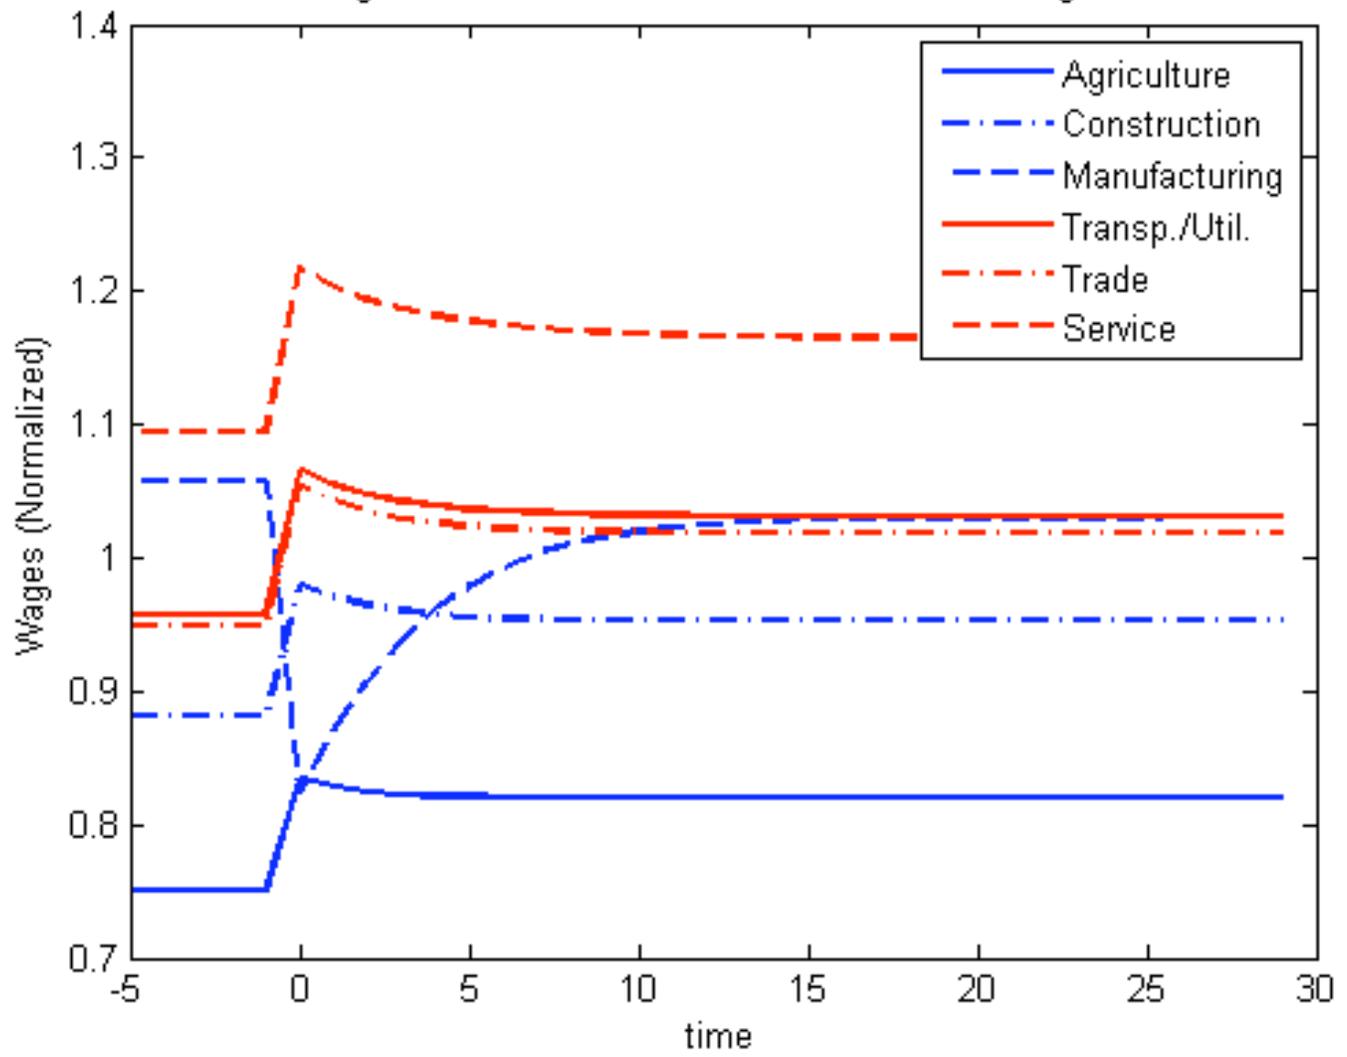


Figure 4: Simulated Trade Liberalization I - Values

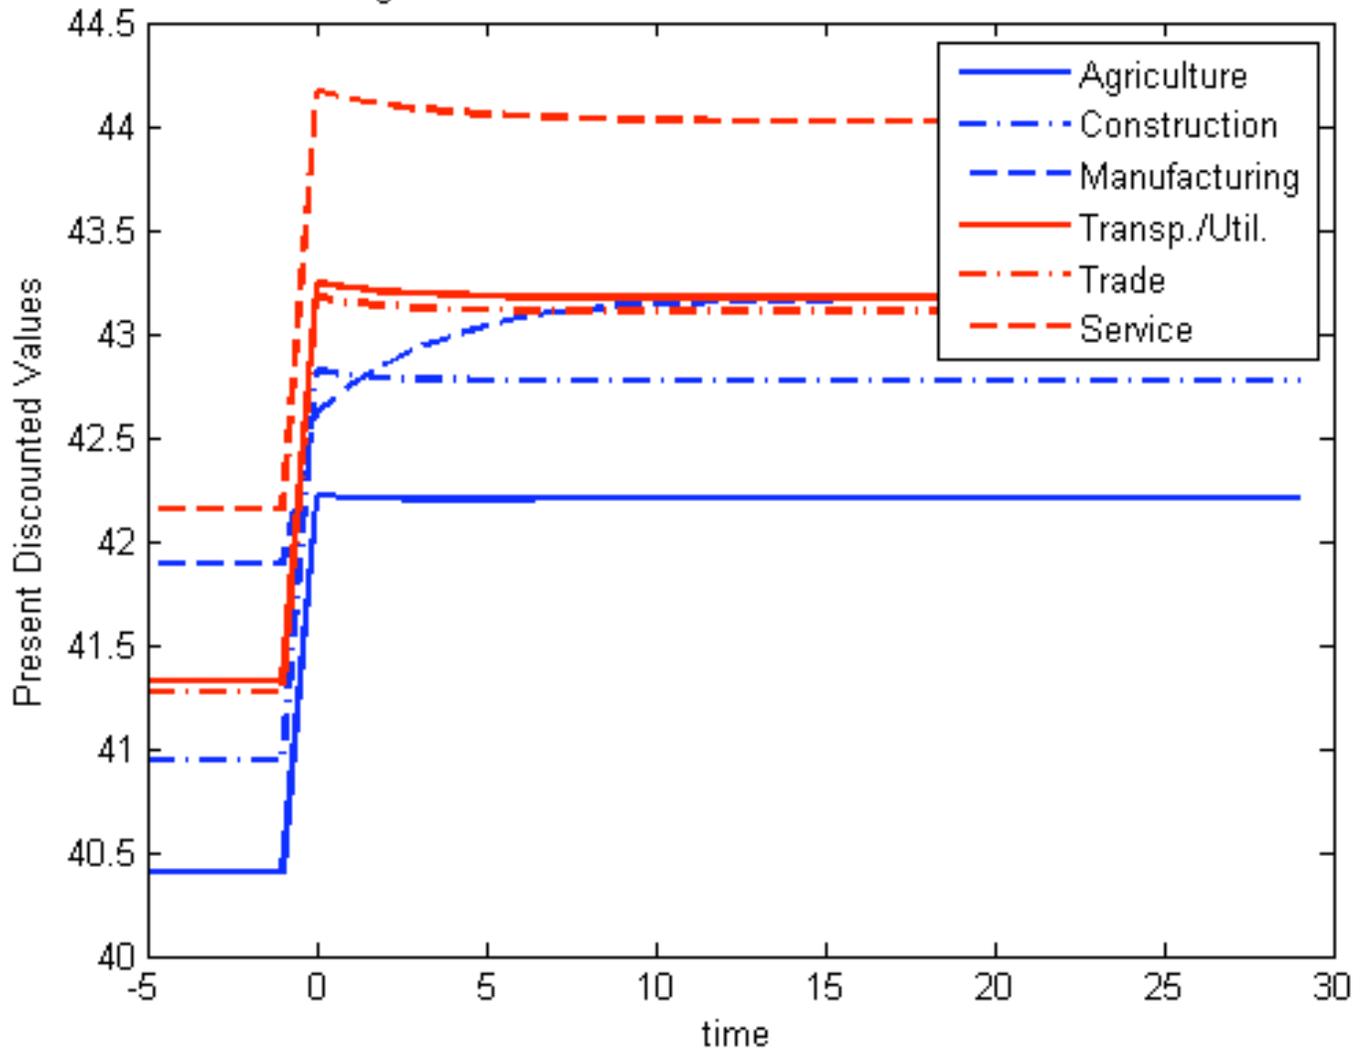


Figure 5: Simulated Trade Liberalization I - Trade of Manufacturing

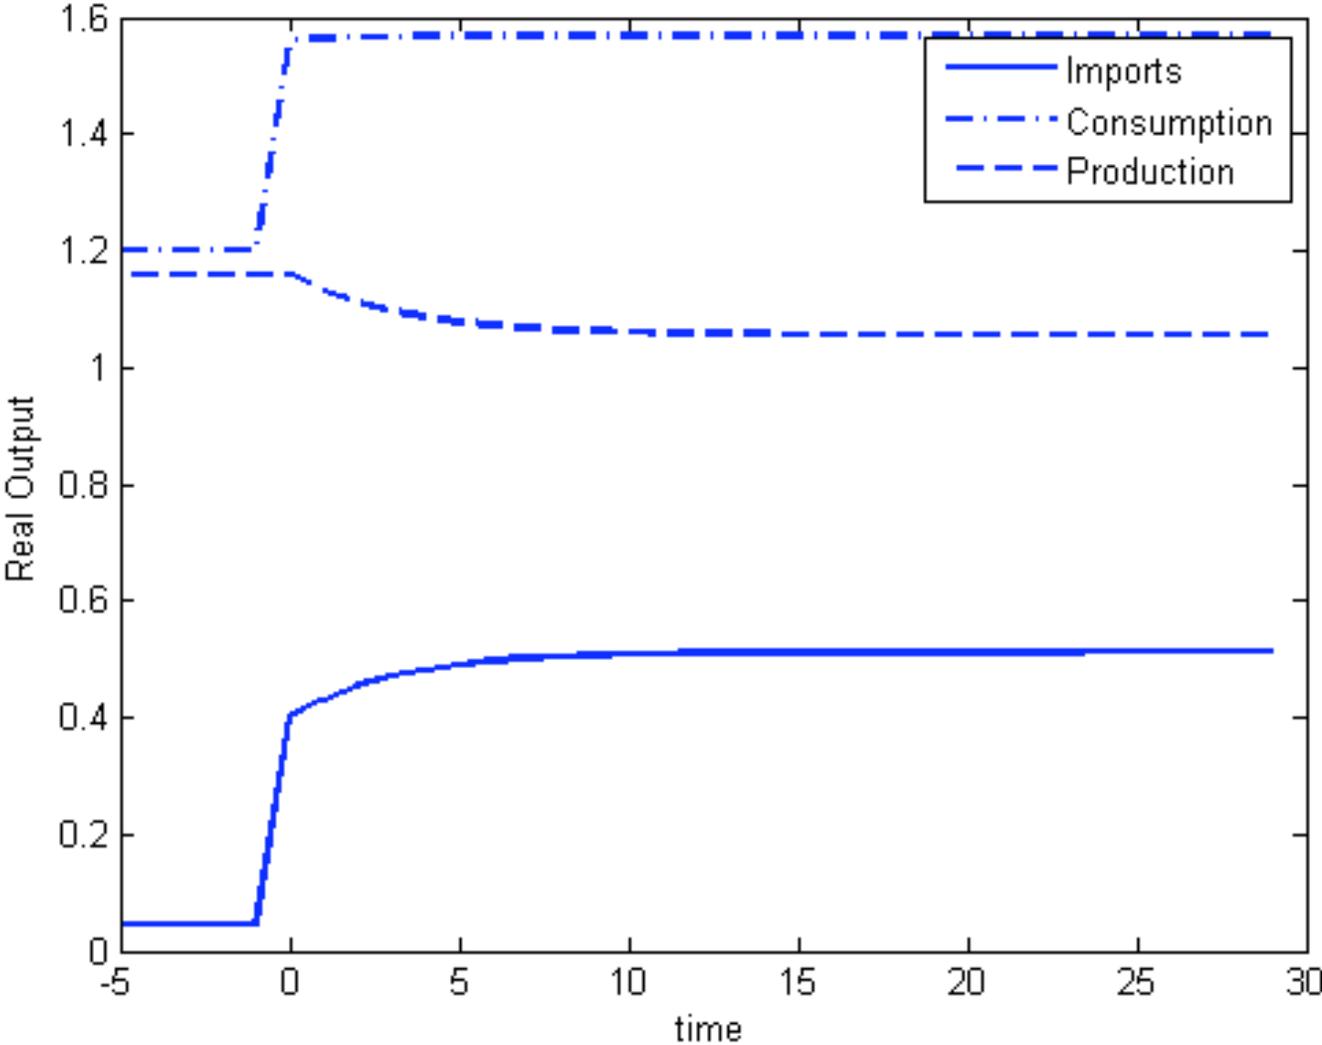


Figure 6: Simulated Trade Liberalization I - Prices

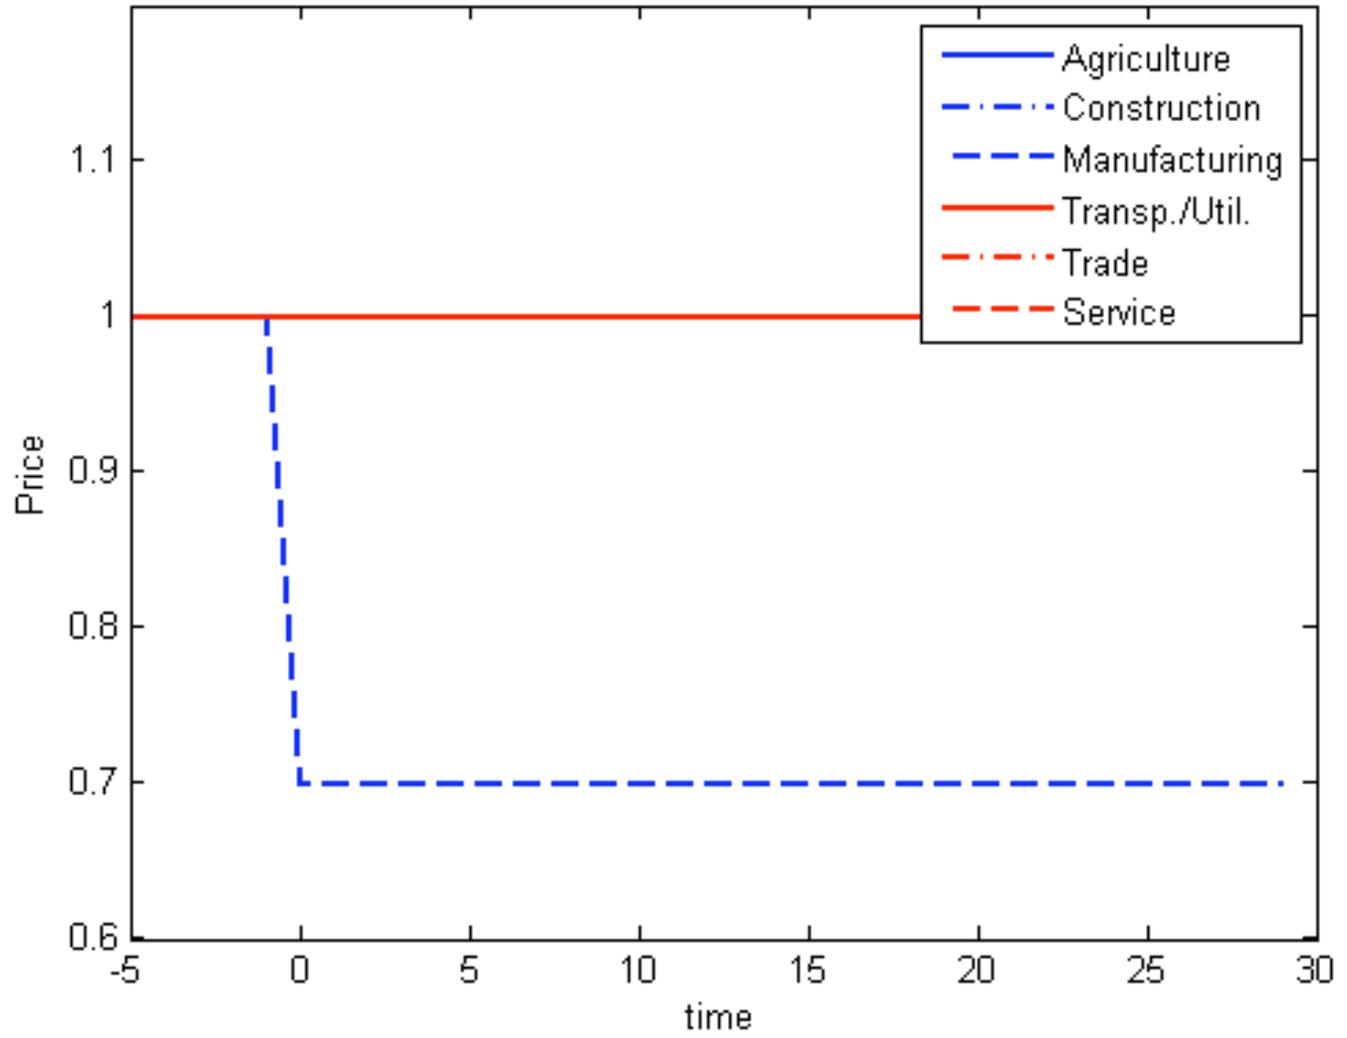


Figure 7: Simulated Trade Liberalization II - Labor Allocation

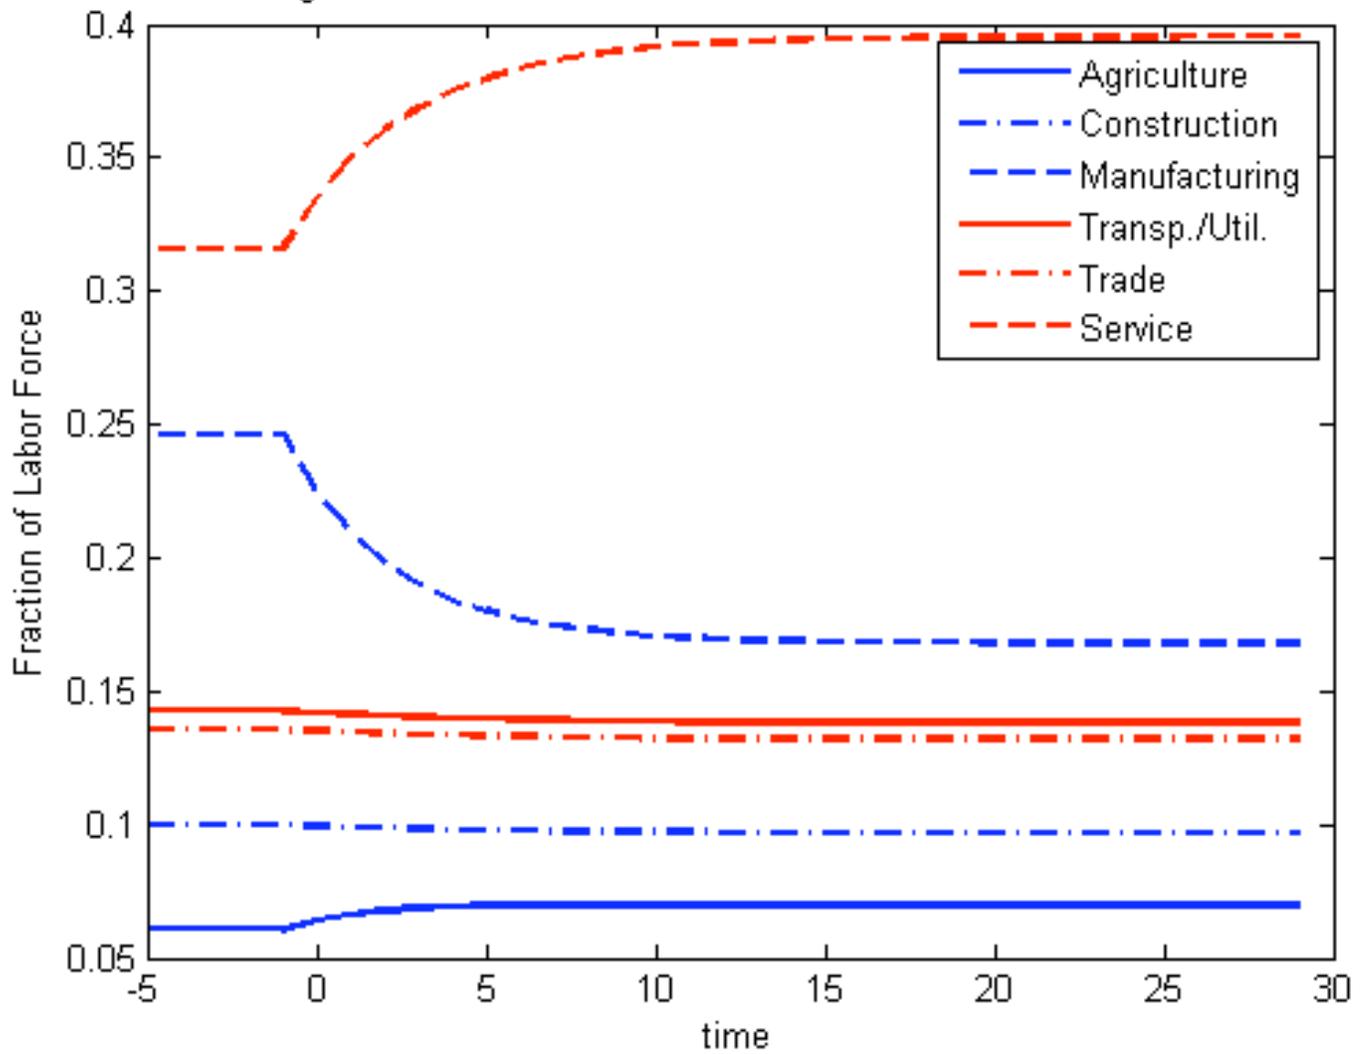


Figure 8: Simulated Trade Liberalization II - Wages

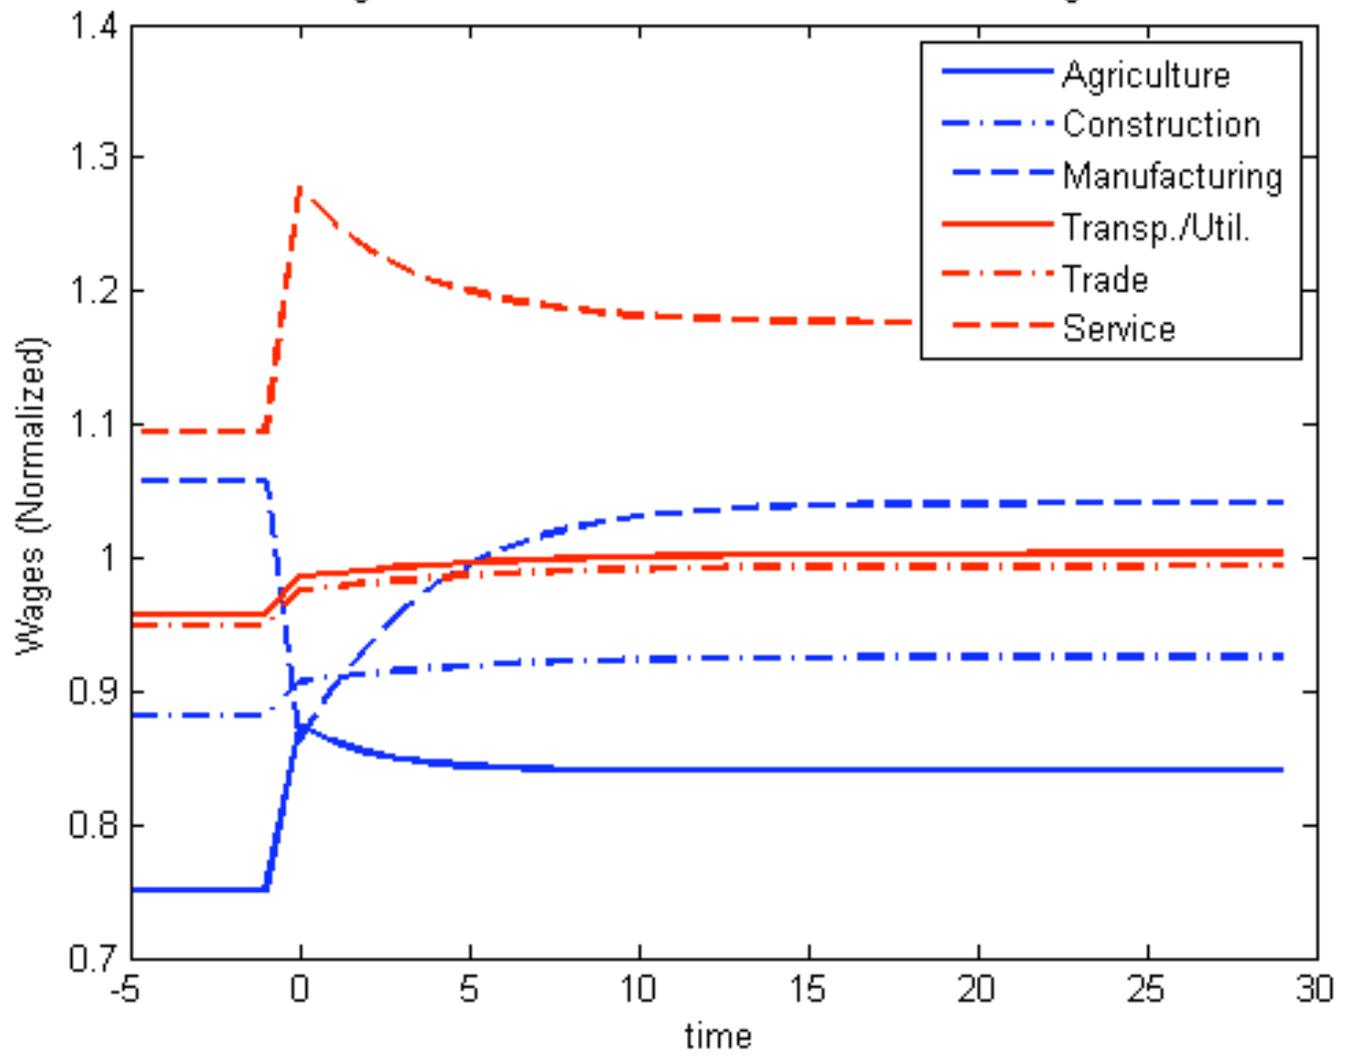


Figure 9: Simulated Trade Liberalization II - Values

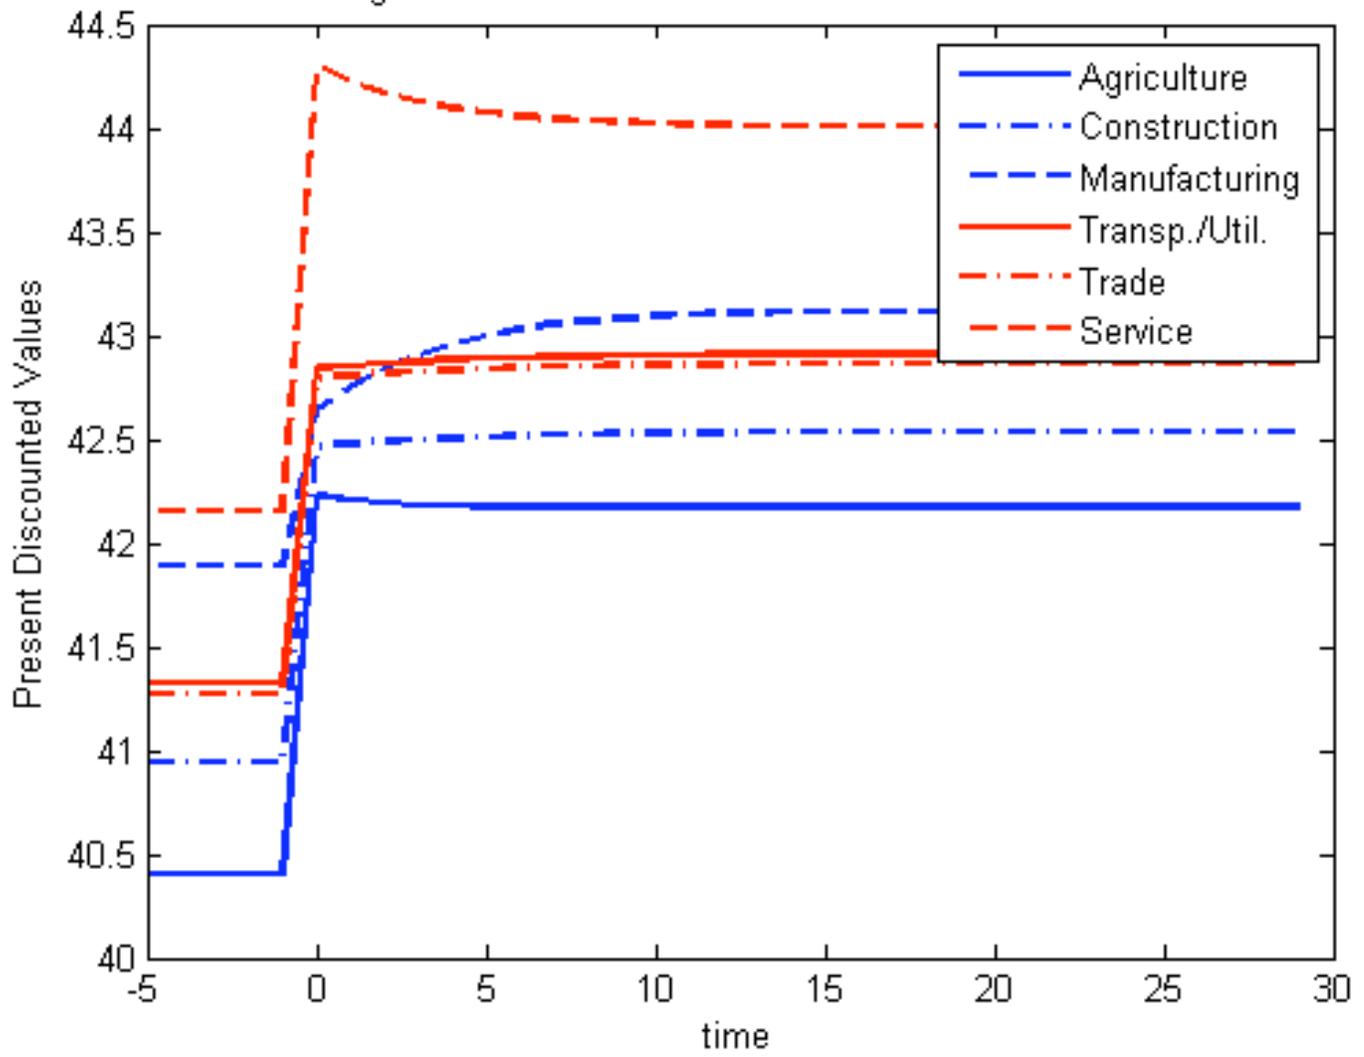


Figure 10: Simulated Trade Liberalization II - Trade of Manufacturing

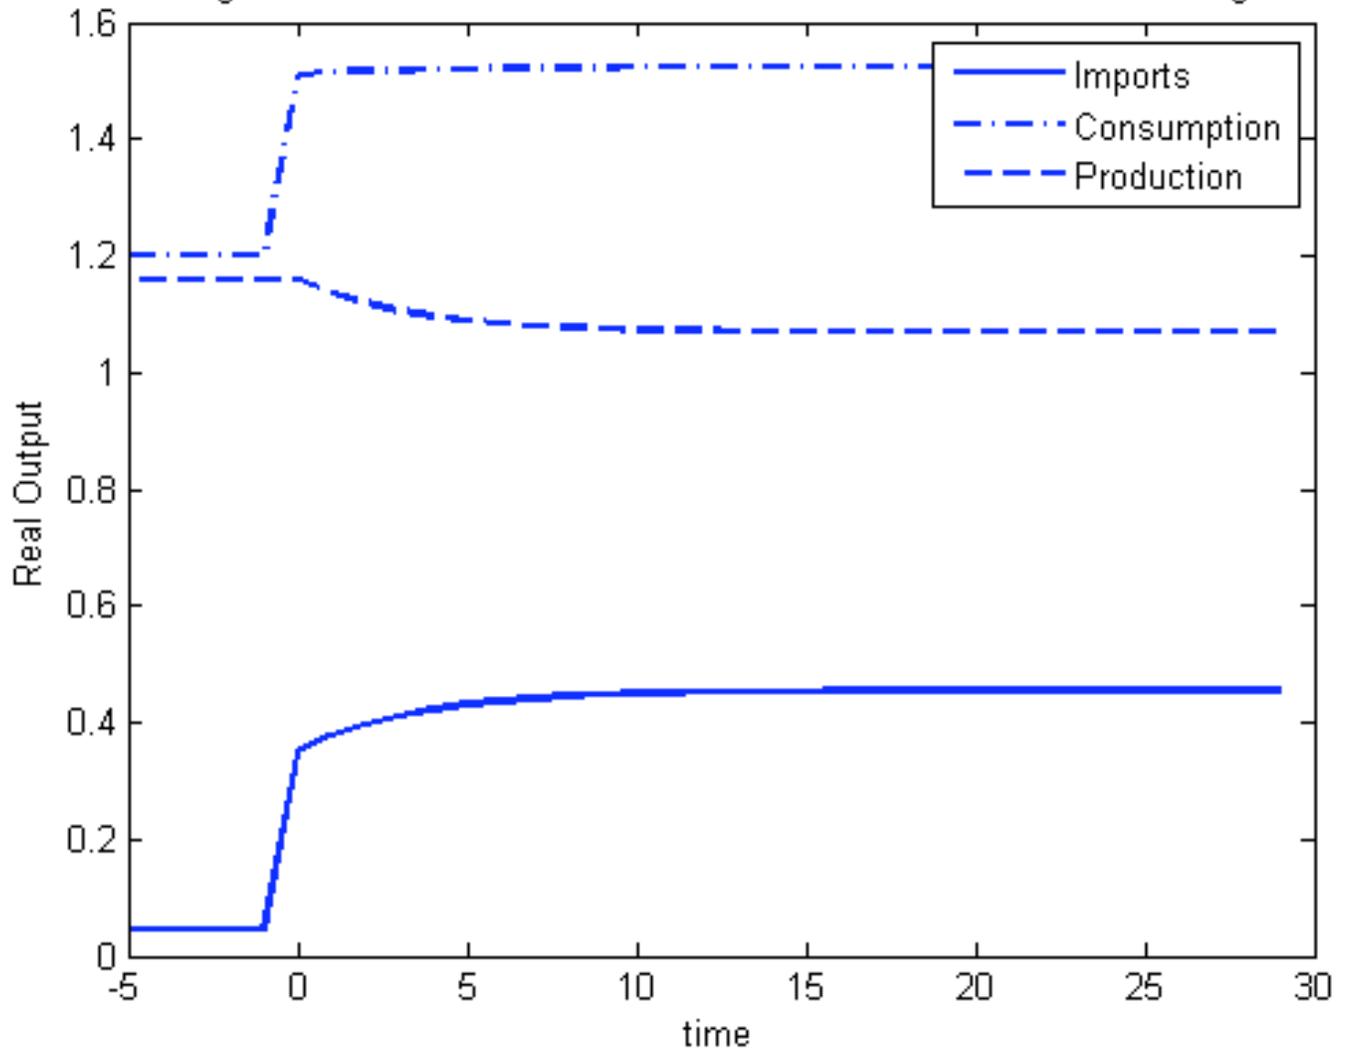


Figure 11: Simulated Trade Liberalization II - Prices

