A Alternate Assumptions on Unobserved Components

In Figure 8, we present alternate indices with varying assumptions on unobserved product quality and outside options. Here we explain why our preferred price indices assume that the unobserved quality of the outside option is changing over time. In the first quarter of 2000, only 2.1 percent of the patients in our sample were treated with a drug treatment that is included in our outside option – a drug regimen that was approved for colorectal cancer but had a very small market share throughout the sample period, where the efficacy and side effect data were unavailable, or where the regimen was not approved by the FDA for colorectal cancer (i.e., an off-label treatment). By the fourth quarter of 2003, on the other hand, the market share of the outside option had increased to 8.4 percent, with 2.8 percent receiving a regimen involving a drug that was approved for colorectal cancer (e.g., oxaliplatin, irinotecan, or capecitabine plus other component drugs), and the remaining 5.6 percent involving drugs approved for other types of cancer. Two years later, in the third quarter of 2005, the cumulative market share of the outside option had more than doubled to 18.3 percent, with almost 10 percent of this accounted for by a regimen involving bevacizumab (Avastin) plus other component drugs. Bevacizumab is expensive and is included in regimens with the longest median survival (see Figure 1). Physicians appear, therefore, to be shifting over time to regimens that use drugs approved for colorectal cancer combined with drugs that were not included in the trials that were submitted for FDA approval. This suggests that physicians attached substantial value to the treatments included in the outside option, which they relied on increasingly over time.

Panel (a) of Figure 8 depicts our preferred specification, as in Figure 6. Panel (c) depicts the three indices assuming that the quality of the outside option changes over time but $\Delta\xi_{jt}$ remains fixed at the first-period values (i.e. assuming $\Delta\xi_{jt}$ reflects a change in tastes, rather than product quality). Here, nearly all variation between the indices have been eliminated. All of the indices appear identical to the indices in Panel (a) pre-1997 (when there is only one available product), but thereafter exhibit a steady decline. By the end of the sample period, consumers experienced welfare gains of about 10 log points (or about a 40 percent decline in quality-adjusted prices). The implication is that the combination of time dummies and residuals explain a lot of the variation across indices. At first glance this might be seem counterintuitive, but it conforms to expectations when considering the implications of (Nevo, 2003). In particular, large swings both in aggregate and product-specific market shares in the sample imply that there are certain behaviors that are difficult to explain through observable characteristics alone, and they are captured by $\Delta\xi_{jt}$ and ξ_t.

Both plots above are constructed based on a scenario where the coefficients on the time indicators are interpreted as measuring how the quality of the outside option changes over time. We believe, as mentioned above, that given the prevalence of off-label use of bevacizumab in the later years of the sample, this interpretation is the most consistent with the data. However, in Panel (b) of Figure 6 we plot the log price indices for the three price indices under a scenario where the quality of the outside option is assumed to be fixed over time at zero (i.e., the coefficients on the time indicators capture changes in the mean quality of the inside goods), and $\Delta\xi_{jt}$ is assumed to reflect changes in unobserved product quality over time. As market share shifts to the outside option after 2000, all three price indices increase by more than those in Panel (a) and Panel (c).

Assuming time dummies reflect changing product quality appears to result in a significant deterioration of the mean quality of inside goods towards the latter part of the sample (when a larger share of patients switch to the outside option), driving all the indices upward. By the end of the sample period, the price indices are 18 to 33 log points (139 percent to 395 percent) higher than in 1993, implying substantial reductions in welfare.

Finally, in Panel (d) of Figure 6 we depict the three indices under the assumption that the quality of the outside option remains fixed (as in Panel (b)), but that $\Delta\xi_{jt}$ also remains fixed as product j's first period value. The pattern of these indices is similar to those in Panel (b), although the three indices do not rise as much by end of the sample period, implying greater welfare gains (or smaller losses).

In Figure 9, we plot the indices from the models with physician and patient heterogeneity under each of these assumptions. Again, panel (a) reflects our preferred specification, as in Figure 7. The indices in the other three panels have a similar pattern where the sickest patients experience welfare improvements due to innovation (and the prices of those products), and healthier patients never fare as well as the sicker patients. In Panel (c) and Panel (d), the welfare differences are fairly small, whereas they are larger in Panel (a) and Panel (b). Looking across the four panels, it is apparent that consumer heterogeneity is critically important when estimating welfare gains from introductions of new regimens, and therefore the implied quality-adjusted prices due to innovation. The two subpopulations experience very different welfare effects from the innovation, such that an aggregate index would incorrectly conclude that all consumers experience the same outcomes.

Figure 8: Estimated Price Indices Without Patient Heterogeneity

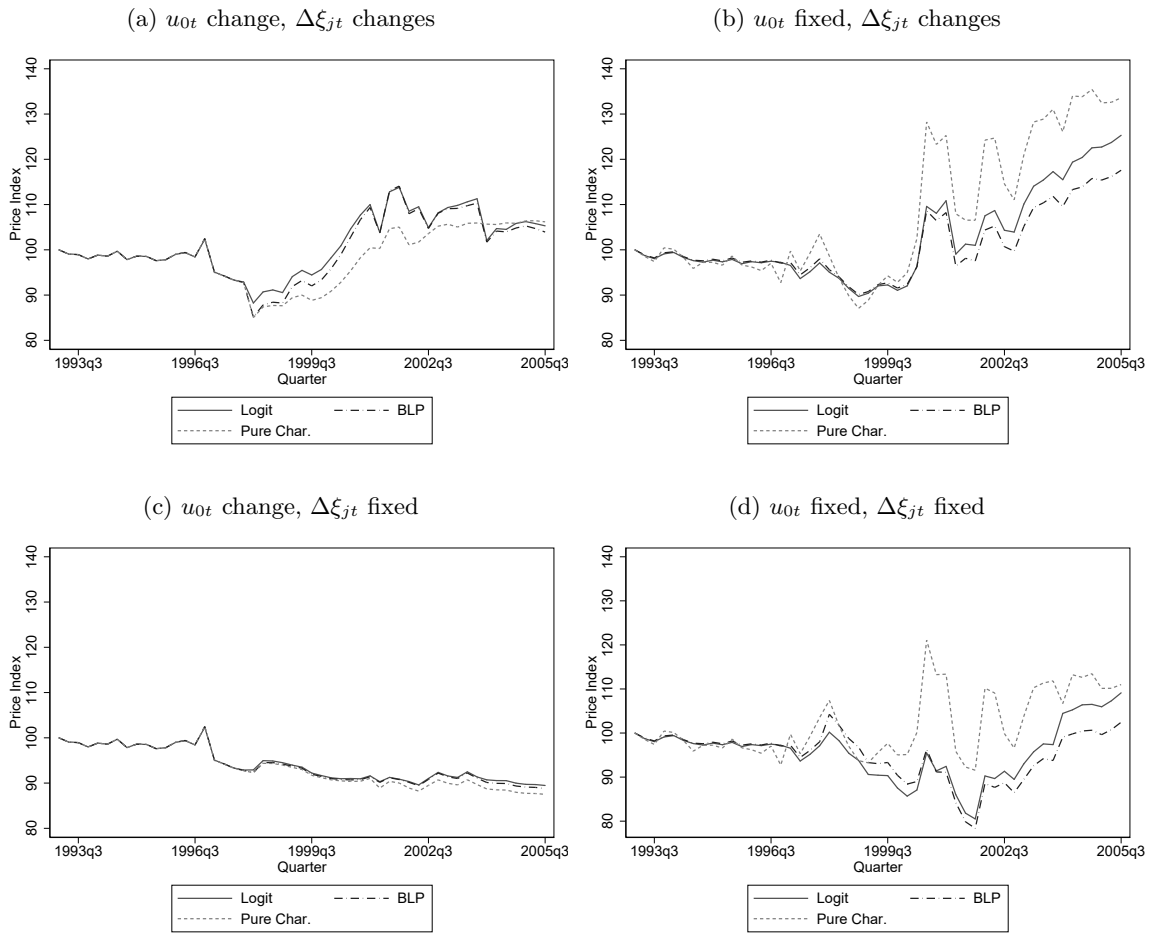

Notes: This figure plots the *aggregate* price indices from the estimated logit, BLP, and pure characteristics models. All indices are on a log scale (i.e. the dependent variable is the "log" price). Panel (a) plots the index from our preferred specification: assuming that the outside option (μ_{0t}) changes in quality each period and that the period-specific deviations from the mean in estimated regimen quality ($\Delta\xi_{jt}$) change each period. Panel (b) plots the indices assuming that μ_{0t} is normalized at 0 each period (and therefore the period-specific mean unobserved utility, ξ_t, changes), while $\Delta\xi_{jt}$ changes. Panel (c) plots indices where μ_{0t} changes each period, but $\Delta\xi_{jt}$ remains fixed at the regimens first-period value (its estimated value in the quarter the regimen first is introduced). Panel (d) plots indices assuming μ_{0t} is fixed at 0 each period and that $\Delta\xi_{jt}$ is also fixed at each regimen's first-period value.

Figure 9: Estimated Price Indices with Patient and Physician Heterogeneity

(a) u_{0t} changes, $\Delta\xi_{jt}$ changes

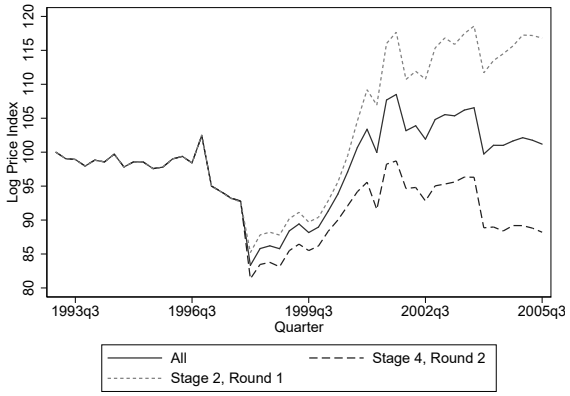

(b) u_{0t} fixed, $\Delta\xi_{jt}$ changes

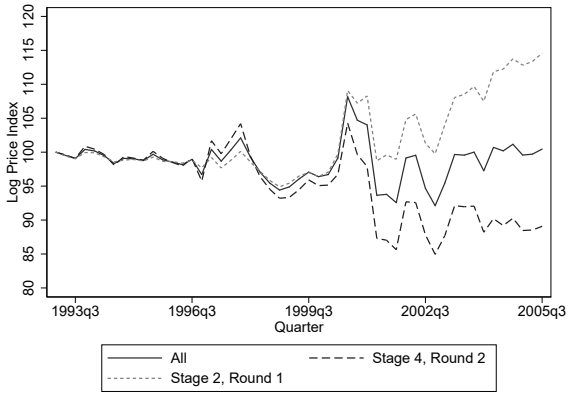

(c) u_{0t} changes, $\Delta\xi_{jt}$ fixed

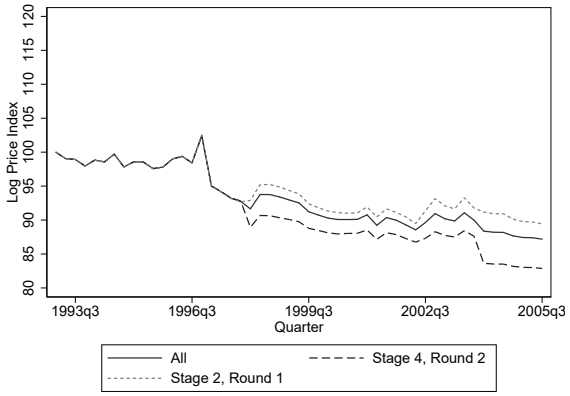

(d) u_{0t} fixed, $\Delta\xi_{jt}$ fixed

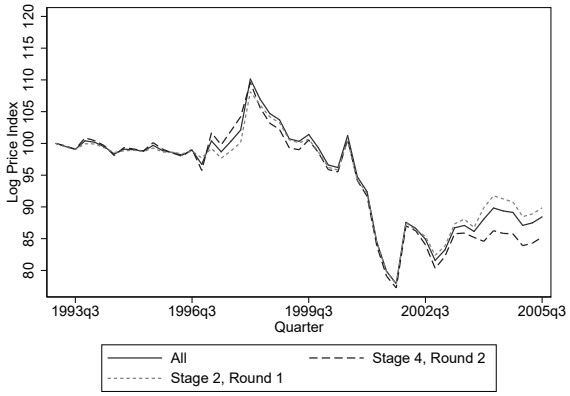

Notes: This figure plots the price indices from the estimated micro-BLP models, incorporating both observed and unobserved heterogeneity. Each panel presents indices for the entire population, as well as for the population of patients with Stage II cancer diagnoses undergoing their first-round of treatment and patients with Stage IV diagnoses undergoing their second-or-higher round of treatment. All indices are on a log scale (i.e. the dependent variable is the "log" price). Panel (a) plots the index from our preferred specification: assuming that the outside option (μ_{0t}) changes in quality each period and that the period-specific deviations from the mean in estimated regimen quality ($\Delta\xi_{jt}$) change each period. Panel (b) plots the indices assuming that μ_{0t} is normalized at 0 each period (and therefore the period-specific mean unobserved utility, ξ_t, changes), while $\Delta\xi_{jt}$ changes. Panel (c) plots indices where μ_{0t} changes each period, but $\Delta\xi_{jt}$ remains fixed at the regimens first-period value (its estimated value in the quarter the regimen first is introduced). Panel (d) plots indices assuming μ_{0t} is fixed at 0 each period and that $\Delta\xi_{jt}$ is also fixed at each regimen's first-period value.

B Estimating the Pure Characteristics Model

The pure characteristics model described in subsection 3.4 is notoriously difficult to estimate. The challenge stems from the absence of the idiosyncratic error term, ϵ_{ijt}, present in both the logit and random-coefficients models. In the previous models, the error term ensured that the market share function was a smooth function of the regimen characteristics, which allowed for easy integration. Without the error term, there is no longer a guaranteed contraction mapping that can be used to equate the predicted and observed market shares, and hence back out the mean utility, δ_j. We, therefore, resort to alternate methods of estimating the mean product quality as described by Berry and Pakes (2007). Our methods also follow closely that of Minjae Song (Song, 2007, 2008).

As in Berry and Pakes (2007) and Song (2008), the estimation procedure proceeds in three steps. The first step is to bring the mean product quality closer to the true product quality by use of a contraction mapping, as in BLP. Recall the utility function in BLP is:

$$u_{ijt} = -\alpha p_{jt} + \beta x_{jt} - \alpha^u v_i p_{jt} + \beta^u v_i x_{jt} + \xi_j + \Delta\xi_j + \epsilon_{ijt} \tag{12}$$

In principle, then, the utility function from the pure characteristics model is a limiting case of the utility function from the random coefficients model as ϵ_{ijt} approaches zero. We thus begin by implementing the BLP contracting mapping with a scaling factor applied to the error term that gradually proceeds to zero. The market share in this model collapses to:

$$S_{jt}(\delta_{jt}, \alpha^u, \beta^u) = \frac{1}{ns}\sum_{i=1}^{ns} \frac{exp[(\delta_{jt} - \alpha^u v_i p_{jt} + \beta^u v_i x_{jt})\mu]}{1 + \sum_{k=1}^{J} exp[(\delta_{kt} - \alpha^u v_i p_{kt} + \beta^u v_i x_{kt})\mu]} \tag{13}$$

where μ is the scaling factor and gradually grows larger in the estimation routine. In practice, as μ grows larger, the exponential function rapidly blows up, resulting in incalculable or missing mean values. We iterate on μ and proceed slowly until the point that we can no longer compute a δ_j. The closes δ_j we obtain using this scaled share function brings us closer to the true mean value from the pure characteristics model, and we hence use this value as the starting point for the next part of the procedure.

The second step is an element-by-element fixed-point homotopy method. This goal of this inversion is to find a mean utility value, δ_j, to satisfy the following equation:

$$|S_{jt}(\delta_j, \delta_{-j}, \theta) - s_{jt}| < tol \tag{14}$$

The difficult is that this inversion is not guaranteed to be a weak contraction mapping. Therefore, both Berry and Pakes (2007), and Song, combine this element-by-element inversion with a homotopy method using the following:

$$\delta'_j(t) = (1-t)\delta_{0j} + t\delta_j, j = 1, ..., J \tag{15}$$

where δ_{0j} is an initial guess for the mean value, δ_j is the current iteration of the mean value, and t is between 0 and 1. As Song points out, when $t = 1$, this collapses back to the strict element-

by-element inverse, which is not guaranteed to contract. However, when $t < 1$, there is a strict contraction mapping guaranteed, albeit the fixed point the model converges to may not necessarily be the true value of the pure characteristics model, δ_j. Therefore, much like the first step of the procedure in which the BLP contraction is implemented with a scaling factor to bring the value of δ_j closer to the true value, the homotopy method is implemented with a value of t that is strictly less than 1, but repeated while approaching one very slowly. Following Berry & Pakes, and Song, we begin with $t = 0.99$ and increase by 0.0025, with the element-by-element mapping repeating at least 50 times before altering the value of t.

The final step is to use the Newton-Rhapson search method. This occurs when the element-by-element inversion in Equation 15 is not satisfied after an iteration of the homotopy search, yet all the predicted market shares are non-zero. This implies the true value of δ_j has not been found yet, but since the current predicted market shares are non-zero, an alternate, more rapid search method for smooth functions can be implemented. While both Berry & Pakes, and Song had to rely on the Newton method in their simulations of the pure characteristics model, our implementation rarely relied on this method, with the homotopy method approaching the true δ_j and satisfying Equation 15 most of the time.

C Composition and Dosages

Table 6: Composition and Dosages of Chemotherapy Drugs

Regimen	1^{st} Drug	2^{nd} Drug	3^{rd} Drug	4^{th} Drug
5-FU + LV	425 mg of 5-Fu/m^2/day for days 1-5, every 4 weeks	20 mg of LV/m^2/day for days 1-5, every 4 weeks		
Irinotecan (Pfizer)	125 mg of irinotecan per week/m^2 for 4 weeks, every 6 weeks			
Irinotecan + 5-FU/LV	180 mg of irinotecan/m^2 on day 1, every 2 weeks	1,000 mg of 5-FU/m^2 on day 1 and 2, every 2 weeks	200 mg of LV/m^2 on day 1 and 2, every 2 weeks	
Capecitabine (Roche)	2,500 mg of capecitabine per m2/day for days 1-14, every 3 weeks			
Capecitabine + Irinotecan	70 mg of irinotecan/m^2/week, every 6 weeks	2,000 mg of capecitabine per m^2/day for days 1-14, every 3 weeks		
Oxaliplatin (Sanofi) + 5-FU/LV	85 mg of oxaliplatin per m^2 on day 1, every 2 weeks	1,000 mg of 5-FU/m^2 on day 1 and day 2, every 2 weeks	200 mg of LV/m^2 on day 1 and day 2, every 2 weeks	
Oxaliplatin + Capecitabine	130 mg of oxaliplatin per m^2 on day 1, every 3 weeks	1,700 mg of capecitabine per m^2/day for days 1-14, every 3 weeks		
Cetuximab (ImClone)	400 mg of cetuximab per m^2 on day 1; then 250 mg/m^2 once a week, every 6 weeks			
Cetuximab + Irinotecan	400 mg of cetuximab per m^2 on day 1; then 250 mg/m^2 once a week, every 6 weeks	125 mg of irinotecan per week/m^2 for 4 weeks, every 6 weeks		
Bevacizumab (Genentech) + Oxaliplatin + 5-FU/LV	5 mg of bevacizumab per kg, every 2 weeks	85 mg of oxaliplatin per m2 on day 1, every 2 weeks	1,000 mg of 5-FU/m^2 on day 1 and day 2, every 2 weeks	200 mg of LV/m^2 on day 1 and day 2, every 2 weeks
Bevacizumab + Irinotecan + 5-FU/LV	5 mg of bevacizumab per kg, every 2 weeks	180 mg of irinotecan per m^2 on day 1, every 2 weeks	1,000 mg of 5-FU/m^2 on day 1 and day 2, every 2 weeks	200 mg of LV/m^2 on day 1 and day 2, every 2 weeks
Bevacizumab + Oxaliplatin + Capecitabine	7.5 mg of bevacizumab per kg, every 3 weeks	130 mg of oxaliplatin per m^2 on day 1, every 3 weeks	1,700 mg of capecitabine/m^2/day for days 1-14, every 3 weeks	

mg=miligram of active ingredient; m^2=meter squared of a patient's surface area; kg=kilogram of a patient's weight.
Source: National Comprehensive Cancer Network, Colon Cancer, Version 2.2006; package inserts.