Hedonic Price Function

Foundations • The Valuation Engineer

Bert Craytor• May 29, 2026

Version 0.1.0-draft

Formal definition. The hedonic price function is the equilibrium mapping from points in the characteristics space of a heterogeneous good to transaction prices, formed as the envelope of buyer bid functions and seller offer functions in a competitive market.

Intuitive framing. We now have two ingredients: a heterogeneous good (a bundle of characteristics) and a characteristics space (a vector space in which each bundle is a point). The hedonic price function combines them. It is the function $p(\mathbf{z})$ that assigns a price to each point $\mathbf{z}$ in the characteristics space, telling us what the market pays for a bundle with that particular combination of attributes.

The key insight from Rosen (1974) is that $p(\mathbf{z})$ is not chosen by any single buyer or seller. It emerges from market equilibrium. Each buyer has a *bid function* expressing the maximum they would pay for any bundle, given their preferences and income. Each seller has an *offer function* expressing the minimum they would accept, given their costs and alternatives. The hedonic price function is the envelope of these bids and offers — the boundary along which buyers and sellers actually transact.

For real estate, $p(\mathbf{z})$ is the answer to the question: if you specified a complete bundle of characteristics, what price would the market produce for it? The answer is not arbitrary. It reflects the aggregate preferences of all buyers in the market and the aggregate costs and alternatives of all sellers, mediated by the equilibrium matching of buyers to bundles.

The hedonic price function is the central theoretical object of hedonic analysis. Implicit prices are partial derivatives of it. Hedonic regression is an empirical estimate of it. The defensibility of any sales-comparison adjustment ultimately reduces to a claim about this function.

Where appraisers encounter it. Appraisers reason about $p(\mathbf{z})$ continually, again usually without naming it. Several practice activities are implicit manipulations of the hedonic price function:

Adjustments in the sales comparison approach. When an appraiser adjusts Comp B upward by $100,000 to account for the fact that the subject has a view and Comp B does not, the appraiser is asserting a claim about $p(\mathbf{z})$: namely, that holding all other characteristics fixed, the function increases by $100,000 as the view characteristic moves from 0 to 1. This is a statement about the local slope of p in the view dimension.

Reconciliation between approaches. The cost approach constructs a value indication from below, summing the component costs of a replicated bundle. The income approach derives value from cash flows, which in equilibrium reflect the price of bundles that produce similar cash flows. The sales comparison approach reads price directly from observed transactions of similar bundles. In a well-functioning market, all three approaches should converge on the same point on $p(\mathbf{z})$. When they diverge, the appraiser must judge which approach has measured the function most reliably for the subject's location in characteristics space.

Market analysis and time adjustments. The hedonic price function shifts over time as supply, demand, financing, and expectations change. A 6-percent annual appreciation rate is a claim about how $p(\mathbf{z})$ has shifted as a whole; a finding that view-premium properties have appreciated faster than no-view properties is a claim that the function has shifted unevenly across the characteristics space.

Why it matters for defensibility. Naming the hedonic price function explicitly clarifies several foundational defensibility questions:

Does the function exist? The Rosen framework assumes a competitive market with enough transactions that an equilibrium function is meaningful. In thin markets, in rapidly changing markets, or in submarkets with idiosyncratic buyer pools, the function may be poorly defined or unstable. An appraiser working in such a market should acknowledge that the inferential machinery is operating on a shaky theoretical foundation.

Is the function the same everywhere? Rosen's framework yields one function per market segment, not one function globally. Pacifica and San Francisco are distinct markets with distinct hedonic functions; single-family homes and condominiums are distinct submarkets even within a single city. An appraiser conflating segments is implicitly assuming a single function spans both, which is usually wrong.

What functional form does p take? The function need not be linear, log-linear, or any other specific form. Rosen showed only that p is the envelope of bids and offers; he did not specify its shape. Empirically, p exhibits nonlinearities (the marginal value of an additional bedroom is often nonconstant in GLA), interactions (view premium depends on lot orientation), and thresholds (a kitchen below some quality bar may attract a discount disproportionate to its condition score). Defensible estimation requires either a functional form supportable from the market or a flexible estimator capable of recovering the shape from data.

Worked appraisal example. For the Pacifica neighborhood under study, we will work with a simple additive functional form that the next entry will exercise more fully:

$$p(\mathbf{z}) \;=\; \beta_0 \;+\; \beta_{\text{GLA}}\, z_{\text{GLA}} \;+\; \beta_{\text{lot}}\, z_{\text{lot}} \;+\; \beta_{\text{view}}\, z_{\text{view}} \;+\; \beta_{\text{cond}}\, z_{\text{cond}}.$$

This form embeds two strong assumptions: that the contribution of each characteristic is linear in its level, and that the characteristics do not interact. Both assumptions are testable and frequently violated; we adopt them here as a starting point.

For the eight comps in this issue's running example, our analysis in the next entries will estimate values of β that yield a function something like:

Parameter	Estimated value
β_0 (intercept)	\$80,777
β_{GLA} (per sf)	\$575
β_{lot} (per sf)	\$22.70
β_{view} (yes/no)	\$86,179
β_{cond} (per step)	\$49,824

Plugged into the additive form, this function predicts a price for any point in the four-dimensional characteristics space. Substituting Comp A's coordinates yields a predicted price near its observed \$1,425,000; substituting the hypothetical subject at $(1700, 7500, 1, 2)$ yields a value indication for that property.

The point of the entry is not these specific numbers but the conceptual move: *the function exists, in principle, before we estimate it.* The function is a property of the market. Our estimation is an attempt to recover it. The recovery may be partial, biased, or noisy, and the next three entries are about how to do that recovery responsibly.

Cross-references: heterogeneous good; *characteristics space*; *implicit price*; *hedonic regression.*