Volume Title: The Detroit Prototype of the NBER Urban Simulation Model

Volume Author/Editor: Gregory K. Ingram, John F. Kain, and J. Royce Ginn

Volume Publisher: NBER

Volume ISBN: 0-870-14258-5

Volume URL: http://www.nber.org/books/ingr72-1

Publication Date: 1972

Chapter Title: Modeling the Housing Market: Demand Allocation, Filtering, Supply, and Market-clearing Submodels

Chapter Author: Gregory K. Ingram, John F. Kain, J. Royce Ginn

Chapter URL: http://www.nber.org/chapters/c3504

Chapter pages in book: (p. 103 - 128)

7

MODELING THE HOUSING MARKET: DEMAND ALLOCATION, FILTERING, SUPPLY, AND MARKET-CLEARING SUBMODELS

THE PREVIOUS CHAPTER contains descriptions of three of the major computer subroutines or submodels used to model a number of important demographic processes central to urban growth and development as well as a number of subsidiary computer operations. This chapter contains similar descriptions of the four major computer submodels used to represent demand and supply forces in urban housing markets: the demand allocation, filtering, supply, and market-clearing submodels. As in the previous chapter the discussion of these submodels follows the order in which they appear in the program and in program operations.

The Demand Allocation Submodel

The demand allocation submodel has the function of allocating the housing demanders described by workplace and household class, $RMOVE(H, J)$, among the twenty-seven housing submarkets or housing types used in the Detroit Prototype. The assignment of households to dwelling units is not a simple one-to-one matching of the seventy-two household classes to the twenty-seven housing types. Instead, it is probabilistic. Members of a household class can reside in more than one housing type, and the demand allocation submodel generates a percentage distribution of each household class over the

housing types. The percentage of a household class that chooses a housing type depends upon two main factors: the nature of the household class itself and variations in workplace-specific relative gross prices of the twenty-seven housing types.

The differences in the proportion of each household class choosing each housing type reflect the differences in taste and income among household classes. For example, large families strongly prefer large, low-density, private structural types such as single-family houses. Similarly, households with high incomes or education prefer and can afford high-quality dwelling units.

Workplace-specific variations in relative gross housing prices, the second major determinant of dwelling unit choice, arise from the interaction of spatial variations in demand for each housing type with spatial variations in its supply. Spatial variations in demand reflect the unequal demographic distribution of employment within metropolitan areas and differences in the characteristics of labor forces among workplaces. Spatial variations in supply reflect differences in the geographic distribution of the housing stock by type as well as some other contemporaneous factors.

Since relative gross housing prices vary among workplaces within the metropolitan area, the effect of prices on housing consumption patterns can be ascertained by analyzing differences in housing choices made by members of the same household class employed at different workplaces. Some of the empirical evidence of these gross price effects has been presented previously.[1] A detailed discussion of the econometric estimation of the parameters of the demand allocation submodel for the Detroit Prototype from Detroit and San Francisco data is presented in Chapter 8 and in Appendix B. At this point it is sufficient to observe that these analyses indicate that housing choices can be predicted more accurately by using information about both household characteristics and gross housing prices rather than information about household characteristics alone.

This result follows from the fact that households have demand curves for different housing types. The parameters of these demand curves are determined by the household's tastes and income, i.e.,

1. See Chapter 4, above, "Some Problems of Causality," and Chapter 3, "The Demand Sector."

its household class. And since the demand curves are downward sloping, an increase in the price of one housing type relative to the prices of other housing types will reduce the probability of a household's consuming the more expensive type. Therefore, information on both household class and gross prices is needed to predict the proportion of households choosing each housing type.

The NBER Urban Simulation Model adds two variations to the standard demand analysis. First, it recognizes that market prices for housing units of the same type vary in a complex manner among neighborhoods and communities in the same metropolitan area. And second, it assumes that gross prices, the market price of housing plus work-trip cost, are the relevant prices for demand determination.

In order to represent demand determination within the framework suggested above, the demand allocation submodel first transforms the expected market prices for the current time period—$P(K, I)$—into an array of gross housing prices which vary by workplace. This transformation first requires the calculation of interzonal travel costs.

Transportation costs between any two zones include the out-of-pocket costs of the trip, e.g., passenger fare or vehicle operating costs, plus the value of the time required to make the trip. For each of the four income classes, interzonal travel time is valued at four-tenths of the wage rate, with the wage rate reflecting the average income within each income class. An average speed of twenty-four miles per hour is used to determine mileages from interzonal travel times, and vehicle operating costs are assumed to be four cents per mile. For owner-occupants, travel costs are capitalized with an annual multiplier of 10.0 to form travel costs that are commensurate with housing prices. Equation 7.1 defines the travel cost array.

$$TCOST(I, J, HY, M) = OPC(I, J, M) + 0.4 * WAGE(HY) * HRS(I, J, M); \qquad (7.1)$$

where:

$TCOST(I, J, HY, M)$ = the travel cost from residence zone I to workplace J for income class HY and mode M;

$OPC(I, J, M)$ = out-of-pocket costs for mode M;

$WAGE(HY)$ = implicit wage rate of income class HY;

$HRS(I, J, M)$ = interzonal travel time for mode M.

Because income level affects the value of travel time, the array of travel costs is classified by residence zone (I), workplace zone (J), and income class (HY) for each of two modes. For each pair of workplace-residence zones the travel cost by the least costly mode is used; so the Detroit Prototype incorporates an elementary representation of modal split. By adding the travel cost array to the array of expected prices in the current period, a complete array of gross house prices is formed as shown in equation 7.2.

$$RES(I, J, K, HY) = TCOST(I, J, HY, MIN) + P(K, I); \quad (7.2)$$

where:

$RES(I, J, K, HY)$ = array of gross price surfaces over residence zones I, for each workplace J, housing type K, and income class HY;

$TCOST(I, J, HY, MIN)$ = travel cost for the cheapest mode for trips from residence zone to work zone by income class;

$P(K, I)$ = array of expected prices by housing type and residence zone.

It is unwieldy to use price surfaces in demand equations. Therefore, each of the surfaces is summarized by taking a weighted average of its surface points in order to form an expected gross price for each housing type. To create these expected gross housing prices, the residential zone values of the $RES(I, J, K, HY)$ surfaces are weighted by the proportion of available units of each type which are in the residence zone, and by the proportion of work trips by income class between each workplace and residence zone. The weights used for a given K, J, and HY, are generated by equation 7.3:

$$WT(I, J, K, HY) = \frac{AVAIL(K, I) * TRIP(I, J, HY)}{\sum_I [AVAIL(K, I) * TRIP(I, J, HY)]}; \quad (7.3)$$

where:

$WT(I, J, K, HY)$ = weight applied to gross price surfaces by residence zone I, workplace J, housing type K, and income class HY;

$AVAIL(K, I)$ = number of units available for occupancy by type and location;

$TRIP(I, J, HY)$ = work trips made by income class.

These variables take into account both the stock characteristics and the spatial range within which a household is likely to consider units. Equation 7.4 defines the gross prices used in the Detroit Prototype.

$$R(J, K, HY) = \sum_I [WT(I, J, K, HY) * RES(I, J, HY)]; \qquad (7.4)$$

where $R(J, K, HY)$ = expected gross housing price by workplace, housing type, and income class.

Since the demand allocation submodel assigns households to housing types by means of demand equations whose independent variables are relative gross housing prices, the matrix of expected gross prices is transformed into a matrix of relative gross prices by dividing each house-type value by a numeraire value. House-type 10 is used as the numeraire because it is the most plentiful unit; so the operation is

$$PCT(H, J, K) = A(H, K) + B1 * REL(J, HY, 1)$$
$$+ B2 * REL(J, HY, 2) + \ldots; \qquad (7.5)$$

where:

$PCT(H, J, K)$ = the proportion of housing demanders of class H at workplace J that chooses housing type K;

$A, B1, B2, \ldots$ = estimated parameters of the demand equation;

$REL(J, HY, 1)$ = the expected gross price of unit 1 divided by the expected gross price of unit 10, e.g., $R(J, 1, HY)/R(J, 10, HY)$, and so on for each housing type.

This approach can get out of hand very quickly, however, and it is necessary to make some simplifying assumptions. Since there are

27 housing types, each equation contains 27 relative price coefficients and a constant. If full interactions are permitted, i.e., if the price coefficients are allowed to vary by all household characteristics as well as housing type, there are 27 equations for each of the 72 household classes. Data requirements render this formulation infeasible, since it requires estimation of $27 \times 27 \times 72$ ($= 52{,}488$) coefficients. To make the demand estimation feasible, it is assumed that of the several household characteristics only income affects the relative gross price coefficients. The equation intercepts then represent the effect of all other household characteristics on the probability of choosing each housing type. By limiting the interaction effects to income class, only 2,916 ($27 \times 27 \times 4$) relative price coefficients, and 1,944 (27×72) constant terms have to be estimated. But, as is discussed in Chapter 8, the estimation from available data of even this number of coefficients, approximately 5,000, is a formidable undertaking.

Moving households are first allocated to housing submarkets by equation 7.6:

$$XMOV(J, K, H) = PCT(H, J, K) * RMOVE(H, J); \qquad (7.6)$$

where $XMOV(J, K, H)$ = number of housing demanders by workplace J and household class H who choose housing type K.
The housing demanders, $XMOV(J, K, H)$, are then summed by workplace, housing type, and income class, and stored in a matrix, $AMOV(J, K, HY)$, to conserve computer storage and running time. In a final bookkeeping operation the demand allocation submodel sums the matrix of assigned movers, $AMOV(J, K, HY)$, over workplaces and income classes to form a vector of demands for each housing type, $DMND(K)$. This demand for each housing type is augmented by a normal vacancy rate to form a total expected demand vector for each housing type, $DEMAND(K)$, as shown in equation 7.7. This vector is later used as input to the supply submodel.

$$DEMAND(K) = DMND(K) * VRATE(K); \qquad (7.7)$$

where:

$DEMAND(K)$ = total expected demand for each housing type K;

$DMND(K)$ = demand by households for each housing type
$$= \Sigma_{J,HY}AMOV(J, K, HY);$$
$VRATE(K) = 1 +$ normal vacancy rate for each housing type.

The Filtering Submodel

The filtering submodel represents quality improvements and
declines in the housing stock that result from wear and tear, the
aging of components, and maintenance, renovation, and repair
decisions by property owners. The theoretical bases for this portion
of the computer program are discussed in detail in Chapter 3. As we
note at that point, we planned initially to filter the entire housing
stock at the beginning of each simulation period. Filtering of the
entire housing stock, however, created bookkeeping problems within
the model that could not be reconciled with the amount of
information maintained in the Detroit Prototype. We were
confronted, therefore, with the need either to increase substantially
the amount of information stored within the model, and thereby
enormously increase its running time and cost, or to redesign the
filtering submodel. We chose the latter course.

The bookkeeping problems were overcome by having the submodel
operate only on the stock of available units rather than on the entire
housing stock. Since the stock of available units includes vacant units
from the previous period as well as units vacated by households
during the current period, it comprises approximately one-quarter
of the standing stock of dwelling units in each period.

In the Detroit Prototype, dwelling units of each structural type
are classified into three quality groups: sound, deteriorating, or
dilapidated. Quality change is represented in the model by altering
the quality classification of dwelling units in each residence zone.
The movement of dwelling units from one quality class to another
in the submodel is a function of two factors: the price difference
between units, which is due to differences in quality, and the cost
of upgrading a unit from the lower quality class to the higher quality
class.

The over-all model maintains an array of expected prices, $P(K, I)$,

which is classified by dwelling unit type, K, and residence zone, I. Each dwelling unit type is described by a vector of characteristics such as structural type, lot size, unit size, and unit condition. For example, one housing type in the model is a single-family house in sound condition with three bedrooms on a quarter-acre lot. Since a unit can have three quality classes, a second type of dwelling unit would be identical to the preceding type in every respect but quality. For example, it could be a single-family house with three bedrooms on a quarter-acre lot, but in deteriorating condition.

Because these two units are defined as separate housing types in the model, the expected price difference between them in a given residence zone can be calculated from the $P(K, I)$ matrix. This price difference, called the quality premium, is depicted by equation 7.8.

$$PQ(KS, Q1, Q2, I) = P(KS, Q1, I) - P(KS, Q2, I); \qquad (7.8)$$

where:

$PQ(KS, Q1, Q2, I) = $ the quality premium between quality levels $Q1$ and $Q2$ for structural type KS in zone I;

$P(KS, Q1, I) = $ the expected price for a unit of structural type KS, and quality level $Q1$ in zone I;

$P(KS, Q2, I) = $ expected price of a similar unit except of quality level $Q2$

The cost of upgrading a unit from a lower to a higher quality classification, $COSTF(KS, Q1, Q2)$, is supplied exogenously to the filtering submodel. These upgrading costs, shown in Table 7.1, are based upon the differences in the construction costs of dwelling unit types differing only in quality.[2] It will be noted from Table 7.1 that the costs of upgrading a unit are constant throughout the metropolitan area, and that lot size does not alter this cost for single-family units.

In equation 7.9, the quality premium for a structural type is calculated by the filtering submodel and divided by the correct upgrading cost to form a profitability measure for the activity of improving units of the given structural type in a residence zone.

2. The derivation of the construction costs is discussed in the Supply Submodel section below. The construction cost differences are increased 30 per cent to reflect the higher costs of working with an existing structure.

$$RATIO(KS, Q1, Q2, I) = PQ(KS, Q1, Q2, I)/COSTF(KS, Q1, Q2);$$
(7.9)

where:

$RATIO(KS, Q1, Q2, I) =$ a profitability measure of transforming a unit of structural type KS from quality level $Q2$ to quality level $Q1$ in zone I;

$COSTF(KS, Q1, Q2) =$ cost of upgrading a unit of structural type KS from quality level $Q2$ to quality level $Q1$.

If the ratio is greater than unity, upgrading of dwelling units will be expected to increase profits, and some owners will upgrade their units. If the ratio is less than unity, however, upgrading of units will not be profitable, and the filtering submodel will lower the quality of some dwelling units in the zone.

The direction and extent of quality change is determined by the interaction of the quality premium–upgrading cost ratio with the response function termed *FILTER*, shown in Figure 7.1. This response function operates on the profitability ratio and alters the quality levels of available units, as summarized in equation 7.10.

$$AVAILF(K, I) = FILTER [RATIO(KS, Q1, Q2, I)] * AVAIL(K, I);$$
(7.10)

Table 7.1
Filtering Costs by Activity[a]

Activity[b]	Structural Group[c]								
	1	2	3	4	5	6	7	8	9
1	1,250	1,250	2,180	2,180	3,960	3,960	14,300	8,580	79,560
2	2,510	2,510	3,050	3,050	7,150	7,150	7,180	4,970	39,520
3	3,760	3,760	5,240	5,240	11,020	11,020	21,480	13,550	119,100

a. Costs in total dollars per structure (model costs are dollars per month, equivalent to 0.01 × total cost).
b. Activity 1 = upgrade from level 2 to level 1.
  Activity 2 = upgrade from level 3 to level 2.
  Activity 3 = upgrade from level 3 to level 1.
c. The structural groups are the nine unit types remaining when quality is removed as a classifying dimension of housing units.

Figure 7.1
The Filtering Function Used in the Model

Ratio = $\dfrac{\text{quality premium}}{\text{cost of improvement}}$

where:

$AVAILF(K, I)$ = the stock of available dwelling units after
some units have changed quality level;

$AVAIL(K, I)$ = number of available units supplied by the
vacancy submodel;

$FILTER$ = the filtering rate response function.

Since the over-all model describes housing types as an array of units indexed by type and residence zone, $AVAIL(K, I)$, the filtering submodel does not consider each dwelling unit individually. In any given zone dwelling units could be moving in both directions between two quality levels. The filtering submodel, however, estimates only the net effect of this movement because the over-all model deals with zonal aggregates rather than with individual units.

Because the filtering mechanism changes the characteristics of the supply of units available for occupancy each period, a limit of 10 per cent is imposed on the filtering rate in any period in the Detroit Prototype. This constraint, like many others in the model, is exogenously specified and can be modified for experimental purposes or if subsequent research suggests a more appropriate value. If the maximum filtering rate were raised to a much higher level, price instability and artificial price cycles in the model could result. The maximum filtering rate has been kept at a moderate level to prevent such cycles.

A second reason for limiting the filtering rate is that many owners of dwelling units may be responsive to profitable opportunities for increasing or decreasing the quality of their units only over fairly long periods of time. The maintenance decisions of owner-occupants in particular may not be very responsive to short-run potential gains or losses. Furthermore, the rate of filtering from higher to lower quality levels must be limited because disinvestment takes time. Quality decline from undermaintenance can occur fairly rapidly, but the extensive physical deterioration necessary for a unit to move from the higher to the lower quality levels of the model does not usually occur in one or two years.

It is very difficult to test the hypotheses in the filtering submodel and compare them with alternative formulations because time series observations of dwelling unit quality and value are virtually nonexistent. As the over-all model runs, this submodel must be calibrated on the consistency and likelihood of its results.

Finally, it should be noted that the filtering submodel is applied separately to each of the forty-four residence zones in the model. Because of differences in price levels across zones the filtering rates will differ in magnitude as well as sign, and the standing stock of dwelling units will adapt to new market conditions.

The Supply Submodel

Because of its visibility, residential development at the fringe of an urban area is often considered the sole component of housing supply activities, and several models which represent only fringe

development have been formulated.[3] However, housing supply activities occur throughout the metropolitan area if only because vacant land is available at locations other than the rural-urban fringe. Furthermore, some of the new construction involves a transformation of the existing stock, i.e., the demolition of existing units and their replacement by new and usually higher-density structures. And some structural alteration occurs which does not involve demolition but rather rehabilitation or the transformation of older single-family units into multiple-family structures.[4]

In order to simulate the workings of the housing market and the impact of public policies in a satisfactory manner, the alterations in the housing stock which occur each year must be represented. The filtering submodel represents some of these changes; and the supply submodel, the remainder. The function of the supply submodel is further to modify the supply of housing units by type and location in response to changes in demand and relative prices. Basic to the supply model is an excess demand vector by housing type that provides target demands for housing suppliers. This excess demand vector, shown in equation 7.11, is formed by subtracting the availabilities obtained in the filtering submodel from the demand forecast by housing type obtained in the demand allocation submodel.

$$XDMND(K) = DEMAND(K) - \sum_I AVAILF(K, I); \qquad (7.11)$$

where $XDMND(K) =$ the excess demand including normal vacancies for each housing type K, in the current period.

The supply submodel employs an input-output array which can be thought of as a production function or activities matrix. This array summarizes the set of efficient technologies and costs for transforming vacant land or existing structures into other structures. Structural alteration may include many different activities such as upgrading a unit in quality, increasing or decreasing the size of a single-family structure, partitioning a single-family unit for multiple-family use, or demolishing and reconstructing units on the site. These

3. See Harris, "Stochastic Process Model"; Morrill, "Expansion of Urban Fringe."
4. A brief but excellent discussion of the magnitude of these stock adjustment activities is presented in Frieden, "Housing and National Urban Goals."

transformation activities are assumed to involve the entire structure rather than just the dwelling unit; i.e., whole structures are used as inputs and produced as outputs.

The first operation in the supply submodel is the formation of prices per structure from dwelling unit prices. This computation, illustrated by equation 7.12, is done simply by multiplying the unit prices by the number of dwelling units per structure.

$$PSTRUT(K, I) = P(K, I) * AVGNO(K); \qquad (7.12)$$

where:

$PSTRUT(K, I)$ = expected structure price by housing type and zone;

$AVGNO(K)$ = average number of units per structure by housing type K.

The prices given by equation 7.12 are the expected prices of both the inputs and the outputs during the current time period. The expected price of vacant land in each zone is also an input price. Vacant land is represented as housing type 28 on the input side of the supply submodel, but is not a possible output.

The next step for the supply submodel is the calculation of the profitability of each possible transformation activity. Before this can be done, however, the price of the input structures for each transformation activity is adjusted to reflect the number of input structures required to produce an output structure by that activity. This adjustment is a function of lot size. For instance, if the input structure is a single-family house on a half-acre lot, and the output unit is a single-family house on a quarter-acre lot, two output structures can be produced from one input structure. Alternatively, when quarter-acre structures are transformed into half-acre structures, two input structures are required per output structure.

In order to handle these possibilities in a consistent way, the input price of the supply activity and its estimated cost are calculated so that whole structures are used as inputs and result as outputs. Therefore, when one structure is transformed into two structures, the input price is that of a single structure, the transformation cost includes the construction of two output structures, and the expected output price for the activity is the price of both output structures.

Conversely, when two structures are required as inputs for the production of one output structure, the input price should be at least twice the price of one input structure.

In fact, the assumption made in this latter case is that the input price is more than twice the single-structure price. When two structures are used as inputs to produce an output structure, the input structures must be next to one another. A larger parcel must be assembled from smaller ones, and a developer usually pays a premium for this. Therefore, in the Detroit Prototype the supply submodel increases the input price by an amount equal to the number of input structures required times 2.5 per cent. Thus if a supply activity requires three input structures to produce an output structure, the price of each of the input structures is 1.075 times its expected market price. The determination of input prices in the supply submodel is summarized as follows:

$$PINPUT(K, KO) = PSTRUT(K, I) * INPTNO(K, KO)$$
$$* AGLOM; \tag{7.13}$$

where:

$PINPUT(K, KO) =$ price of input K when housing type KO is output;

$INPTNO(K, KO) =$ number of structures of type K required to produce a structure of type KO;

$AGLOM = [1.0 + 0.025 * INPTNO(K, KO)]$ if $INPTNO$ exceeds 1.0, and 1.0 otherwise.

After calculating the input price, $PINPUT$, and the output price, $POUTPUT$, the gross profitability of transforming structures of type K into structures of type KO in zone I is

$$PROFIT(I, K, KO) = POUTPUT(I, K, KO) - [PINPUT(I, K, KO)$$
$$+ COST(K, KO)]; \tag{7.14}$$

where:

$PROFIT(I, K, KO) =$ the expected profit of producing structures of type KO from inputs of type K in zone I;

$POUTPUT(I, K, KO) =$ the total expected price of output structures produced by the activity;

$COST(K, KO)$ = exogenously estimated cost of trans-
forming K to KO.

The gross profit amount is then transformed to a gross profit rate,

$$RATE(I, K, KO) = PROFIT(I, K, KO)/[PINPUT(I, K, KO) \\ + COST(K, KO)]; \quad\quad (7.15)$$

where $RATE(I, K, KO)$ = gross profit rate of producing output KO from input K.

At least two features of the supply submodel as described so far should be noted. First, the transformation cost matrix is both independent of the residence zones and constant across them. And second, the transformation cost matrix includes pure rather than mixed transformations, i.e., transformation activities that combine two different structural types as inputs or produce two different structural types as outputs have not been included. Mixed activities could be replicated within the supply framework described thus far, but the process would generate many more possible supply activities. After weighing the computational costs of including these extra activities against their returns in model verisimilitude, we excluded them.

The algorithm used to assign levels to the transformation activities, *SUPPLY*, first ranks all feasible activities according to their profit rate. Then, starting with the most profitable activity first, it assigns activity levels that are consistent with available inputs, zoning, and expected demands. Since only feasible activities are ranked, considerably fewer than the possible total of 33,264 activities are considered by the algorithm. Each of the forty-four residential zones probably will not have structures available to be used as inputs in each of the twenty-seven structure-type categories. When a zone lacks an input structural type, no transformation activities using that input type in the zone are considered. Transformation activities which have profit rates less than zero are also eliminated.

The number of activities considered could have been reduced further by casting out those transformation activities which produced structural types prohibited by zoning laws in the relevant residence zone. This was not done in the Detroit Prototype, however, because we wished to see the extent to which zoning constitutes a constraint

on supply activity. By including activities prohibited by zoning laws, one can get a sense of the opportunity costs of zoning and the pressure for zoning changes from the profitability levels of the prohibited activities.

Experience with the model suggests that two to four thousand activities will remain after all eliminations. Of course, the actual number of feasible activities in any period is a function of that period's expected housing prices and of the distribution of the existing stock. These feasible activities are then ranked by their profit rates, and activity levels are assigned in accordance with the three constraints on supply transformations.

The availability constraint, $AVAILF(K, I)$, indicates the number of structures by residence zone and housing type that are available for participation in the market this period. These include units which have been standing vacant as well as units which households vacate during the market period. In the Detroit Prototype, the quantity of vacant land available in each zone during a period, $VLAND(I)$, is 10 per cent of the total vacant land remaining in each zone in the current period. This proportion is another exogenously specified parameter that is meant to reflect more complex behavior of the land market, and it can be modified easily in experimental simulation.

The zoning constraint, $ZONE(K, I)$, limits the number of output structures of each type which can be produced in a zone. If the structure is prohibited, the constraint would have a value of zero. If there are no zoning constraints, the constraint has a very large value. This constraint can handle intermediate zoning quantities as well. Thus a zone might allow up to ten apartment buildings to be present.

The demand constraint for each housing type, $XDMND(K)$, prevents the supply model from overbuilding units of a given structural type during each simulated period. Since an oversupply of units will cause prices to soften and force builders to carry their inventory of vacant units longer (a very costly arrangement for them), the supply submodel is sensitive to aggregate demand conditions. The demand constraint results from the assignment of households to submarkets that is carried out in the demand allocation submodel above.

The level assigned to each transformation activity, starting with the most profitable one, is the level of the smallest of the zoning,

availability, and demand constraints. After an activity level is assigned, the constraints are revised to reflect the assignment, and the next highest ranking activity is considered.

When the submodel has finished examining the list of profitable activities, there is no guarantee that the remaining excess demand for each housing type will be zero. Furthermore, it is possible that the excess demand for some housing types would actually increase because of the supply submodel. The submodel might, for instance, transform a certain housing type into other types and increase the "deficit" of that input housing type. In fact, however, this has happened rather infrequently in test runs of the simulation model. When it has happened, the unfilled demand has typically been in lower-quality dwelling units, a result not greatly at variance with observed market behavior. Furthermore, a small amount of unfilled excess demand typically will only reduce the vacancy rate within that housing type, since demands have been augmented by vacancy rates. Equation 7.16 summarizes the operations of the supply submodel in simplified terms.

$$AVAILS(K, I) = SUPPLY[AVAILF(K, I)]; \qquad (7.16)$$

subject to:
  a. Profit:
$$RATE(I, K, KO) > 0;$$
  b. Availability:
$$AVAILS(KO, I) \leq \sum_K [AVAILF(K, I)/INPTNO(K, KO)]$$
$$+ VLAND(I)/INPTNO(28, KO);$$
  c. Zoning:
$$AVAILS(K, I) - AVAILF(K, I) \leq ZONE(K, I);$$
  d. Forecast demand:
$$\sum_I AVAILS(K, I) - \sum_I AVAILF(K, I) \leq XDMND(K);$$

where

$AVAILS(K, I) =$ the number of units available for occupancy in the current period after new construction and transformations;

$SUPPLY =$ algorithm used to assign levels to transformation activities.

Data for the Supply Submodel

In its operation the supply submodel relies on a significant body of exogenous data, *COST(K, KO)*, the transformation cost matrix. The relative magnitude of these transformation costs is an important determinant of the construction and transformation activities carried out by the model. Furthermore, these costs are used as a basis for the costs in the filtering submodel and also play an important role in the price formation routines of the market-clearing submodel. In a sense, the transformation cost matrix is a numerical statement of the production functions used in the supply submodel. Because of its ubiquitous role, the derivation of the cost matrix is a matter of some importance.

The transformation costs used in the model are based on data given in *The Dow Building Cost Calculator*, a manual designed to help real estate agents, tax appraisers, and fire insurance agents determine the approximate replacement cost of buildings.[5] The publication contains the basic cubic foot costs of a wide range of building types as well as tables of multipliers for adjusting the basic costs for different cities in North America. The data are based on construction industry statistics compiled by McGraw-Hill and the F. W. Dodge Company (now a McGraw-Hill subsidiary). It is asserted that the cost figures generated by the metropolitan multipliers will apply over an area within a radius of twenty-five miles of the listed urban area.[6] In order to generate cost figures for this simulation model, the basic 1960 cost multiplier for Detroit was applied. The construction costs should therefore be compatible with the housing prices which were generated for the same year in Detroit.

Several problems were involved in estimating the construction and transformation costs used in the simulation model. Perhaps the main one was that the housing types used in the over-all simulation model were defined in a relatively unspecific way, whereas in order to estimate the construction cost of a structure, the structure had to be described in fairly great detail. Thus estimation of the construction cost matrix required specification of at least the size of each housing type.

5. *Dow Building Cost Calculator.*
6. Ibid., pp. c and d.

The second problem encountered in formulating transformation costs involved the definition of quality. In particular, one might wonder how to calculate the construction cost of a structure falling in the lower quality category. The *Building Cost Calculator* was of help with this problem since it defined each house described as being of good, average, or fair quality and gave information on the unit's interior finish and fixture grade.

The construction cost finally selected for each quality level represented an average of the costs given for several similar structures. The size of the structure also reflects an average for structures with the number of rooms or units used in the model. These basic construction costs for the twenty-seven structural types when vacant land is used as the input are shown in Table 7.2.

The basic construction costs must be expanded considerably in order to obtain the entire transformation cost matrix. This expansion was based upon the data in Table 7.2 as well as upon estimates of demolition costs given in the *Building Cost Calculator*.[7] For quality changes within a general structural type it was assumed that structures could be upgraded by an amount equal to 30 per cent more than the original construction cost difference between quality levels. Downgrading a structure in quality is assumed to be costless. Transformations other than quality change and new construction include increasing unit size via additions, making units smaller via demolition or subdivision, and changing structural types via demolition and reconstruction. Transformation activities for which single-family structures are both the input and output structures involve alteration to the standing structure. Transformations between single-family types and multiple-family types and between different multiple-family types involve demolition of the standing structure and reconstruction. Therefore, the transformation costs reflect specific ways of effecting transformations. The technological coefficients are fixed, as are their costs.

Since these costs are exogenous inputs, it is an easy matter to alter them to represent new technologies or the introduction of cost-saving methods in the residential construction industry. A shifting of the entire array, e.g., increasing all construction costs

7. P. 14.15.

Table 7.2
Housing Types and Construction Costs

No.	Quality[a]	Lot Size per Structure (acres)	No. of Rooms per Unit	Cubic Ft. per Unit (thousands)	Units per Structure	Cost per Cubic Ft., 1960	Construc- tion Cost per Unit
1	1	.25	5	9	1	$0.993	$ 8,940
2	2	.25	5	9	1	0.886	7,980
3	3	.25	5	9	1	0.672	6,050
4	1	.50	5	9	1	0.993	8,940
5	2	.50	5	9	1	0.886	7,980
6	3	.50	5	9	1	0.672	6,050
7	1	.25	6–7	14	1	1.023	14,300
8	2	.25	6–7	14	1	0.901	12,620
9	3	.25	6–7	14	1	0.733	10,270
10	1	.50	6–7	14	1	1.023	14,300
11	2	.50	6–7	14	1	0.901	12,620
12	3	.50	6–7	14	1	0.733	10,270
13	1	.25	8	20	1	1.070	21,390
14	2	.25	8	20	1	0.817	18,340
15	3	.25	8	20	1	0.642	12,840
16	1	.50	8	20	1	1.070	21,390
17	2	.50	8	20	1	0.817	18,340
18	3	.50	8	20	1	0.642	12,840
19	1	.0625	5	15	1	1.040	15,590
20	2	.0625	5	15	1	0.856	12,840
21	3	.0625	5	15	1	0.764	11,460
22	1	.25	4	12	6	0.947	11,370
23	2	.25	4	12	6	0.856	10,270
24	3	.25	4	12	6	0.794	9,530
25	1	.25	3	10	15	1.380	13,750
26	2	.25	3	10	15	1.070	10,690
27	3	.25	3	10	15	0.920	9,170

a. High quality is denoted 1: medium, 2; and low, 3.

by a proportional amount, would have little effect since it would not change the rank ordering of the transformation activities by their profit rates, but would render some formerly profitable activities unprofitable. The net effect would be to shorten the list of feasible activities without altering the order in which they are considered. In order to generate more substantial changes in stock alteration

activities, one must alter the cost of producing some types of
structures relative to others. Such a change would produce a greater
variation in the activities actually performed.

The Market-clearing Submodel

When the market-clearing submodel is reached within the over-all
model, the demand allocation submodel has allocated the relocating
households to discrete housing types, and the filtering and supply
submodels have generated a revised stock of units available for
occupancy. The major function of the market-clearing submodel
is to locate current-period demanders of housing in the available
units. In addition, the market clearing submodel updates the trip
matrix and generates a matrix of prices which will be the expected
prices during the next model time period.

In order to locate the movers assigned to housing types—
$AMOV(J, K, HY)$—in the available units—$AVAILS(K, I)$—it is
assumed that each housing type forms a separate and independent
housing submarket. That is, within the market-clearing submodel a
household can choose only among dwelling units of the type it
selected in the demand allocation submodel. This restriction allows
the general location assignment problem to be partitioned by housing
type into twenty-seven smaller, more easily solved assignment
problems, each of which involves assigning households classified
by four income classes and nineteen work zones to available units
in the forty-four residence zones as shown in equation 7.17.

$$MIN \sum_{I,J,HY} TCOST(I, J, HY) * X(I, J, HY); \qquad (7.17)$$

for each separate K subject to:

$$\sum_{I} X(I, J, HY) = AMOV(J, K, HY);$$

$$\sum_{J,HY} X(I, J, HY) = AVAILS(K, I);$$

where $X(I, J, HY)$ = households of income class HY employed at
workplace J who locate in zone I, given they have chosen housing
type K.

In terms of substitution possibilities, the foregoing assumption means that elasticities of substitution between housing types are accounted for solely by the demand allocation submodel, wherein households are allocated to housing types on the basis of these elasticities and the expected housing prices of the current period. Subsequently, the market-clearing submodel prohibits substitutions between housing types and only allows substitution between locations within a housing type submarket. In this latter submodel, therefore, the dwelling units making up a submarket are perfect substitutes for one another in all unit attributes except residence zone location.

It should be recalled that specific households are not assigned to specific dwelling units within this submodel. Both the demand allocation submodel and the market-clearing submodels have probabilistic interpretations. They generate distributions of households first over housing types and then over residential space. The groups allocated in these submodels are households classified by household characteristics and workplace zone. However, in the demand allocation submodel the whole vector of household characteristics is used in conjunction with workplace to determine the housing types chosen by households, since all of the characteristics affect a household's tastes. On the other hand, in the market-clearing submodel household income class is the only household characteristic used in conjunction with workplace, since it is assumed that travel costs are independent of the remaining household characteristics.

The treatment of excess demand and excess supply within the market-clearing submodel has been discussed in detail in Chapter 4.[8] By adding pseudo-households or pseudo-units to make the total number of households equal the total number of available dwelling units in each submarket, the linear programming problems in equation 7.15 can be solved and the identity of vacant units or excess households established. Units which are in excess supply are treated by the model as being vacant. They are identified as such and will become available units during the next market period. Households that are not located in a dwelling unit of the type they chose in the

8. See in Chapter 4, above, "Market Clearing, Excess Demand, and Disappointed Expectations."

demand allocation submodel are also carried over to the next
market period when they will again participate in the housing
market.

Price Formation

The over-all model employs an array of expected prices, $P(K, I)$,
indexed by housetype, K, and residence zone, I, to mediate between
the supply and demand sectors of the model. These expected prices
are altered each market period by the market-clearing submodel.
The manner in which prices are generated in the model has been
outlined in Chapter 3, and an example of this price generation will
now be described.[9]

Since the market-clearing submodel uses a linear programming
algorithm of the Hitchcock type to allocate households to residence
zones, dual variables or shadow prices are generated as part of the
assignment solution. The dual variable associated with a residence
zone denotes the change in total transport cost which would occur
if one unit of capacity were added to the residence zone and
households were reassigned. Thus assume there are three residence
zones, A, B, and C, in a given submarket with dual variables of
-50, -30, and 0, respectively. If a unit of capacity were added to
zone A, and households reassigned, total transport costs would fall
by 50. In forming prices the dual variables are altered by a sign
change to signify differences in travel savings instead of differences
in travel costs. Therefore, the travel saving associated with zone A
is 50, and the difference between zone A and B is still 20. The
travel-saving shadow price associated with each zone reflects its
marginal locational advantage, and these quantities are interpreted as
location rents.[10] Zone C earns no location rent and is the marginal
zone because it has excess capacity. Vacant units within the
submarket are located in zone C.

In forming prices it is assumed throughout the market-clearing
submodel that units within a submarket are perfect substitutes
except for location. Given this assumption and an instantaneous
adjustment of prices to their equilibrium within the period, the

9. See in Chapter 3, above, "The Price Formation Sector."
10. Stevens, "Linear Programming and Location Rent," p. 20.

differences in market prices observed between residence zones
would correspond to the differences in location rents between
zones. That is, the price difference between zones A and B in
one-period equilibrium would be 20, the difference in location rent.
Since the differences in price across zones are known from the
location rents generated, to specify the entire price surface it is
necessary only to specify the market price in one residence zone.
Generating the price surface thus amounts to specifying the height
of the location rent surface.

In the market-clearing submodel the price surface is formed by
adding to the location rent surface the cost of producing a unit of
the appropriate type in the marginal zone. In the three-zone example,
the price of units in zone C is set equal to the cost of producing a
unit in zone C. If this cost were 200, the unit prices in zones A, B,
and C would be 250, 230, and 200, respectively. The production cost
selected is the cost of the supply activity which is feasible and
least expensive; it could involve new construction or the conversion
or quality change of an existing unit. In the supply submodel it is
also possible to transform a unit into itself at zero cost, and this
null activity could be the least expensive way of producing a housing
type which is in excess supply. If this were the case, the cost of
producing the unit would be merely its expected price for the current
period. However, when the null activity is the least expensive one
in the marginal zone, the supply cost in the zone is set to nine-tenths
of the expected price of the unit in the period so that prices will fall
over time when a unit is in long-run excess supply. But when a unit
is not in long-run excess supply, the one-period equilibrium prices
are determined by supply costs.

Of course, prices in the housing market do not adjust
instantaneously to a short-run equilibrium, and the prices formed
from the location rents do not purport to be transaction prices in the
market in the current period or even to be the next period's expected
prices. Therefore, as shown in Figure 3.8, the "equilibrium" prices
generated for the current period are combined with the period's
expected prices in an adaptive expectations framework to form the
expected prices for the next period.[11] The adaptive expectations

11. See "The Price Formation Sector" in Chapter 3, above.

model is used because housing prices are sticky and change relatively
slowly over time. Furthermore, the location rents generated each
period are based upon incremental additions to zonal capacity in a
given period. They may vary a great deal from period to period
because the supply submodel can make relatively large changes in
a zone's capacity during any one market period. Over several
periods, however, these one-period equilibrium prices are expected
to produce reliable indications of price levels in the housing
market.

After assignments have been made to each submarket and the
location rents for each housing type have been calculated, the next
period's expected price of land in each residence zone is formed.
Within a given zone, each housing type will have a different location
rent. The one-period equilibrium land price in a zone is calculated
as an average of the location rents weighted by the respective stock
of units of each type in the zone, as follows:

$$PLAND(I) = \sum_K \left[\frac{LRENT(K, I) * STOCK(K, I)}{\sum_K STOCK(K, I)} \right]; \qquad (7.18)$$

where:

$PLAND(I)$ = the one-period equilibrium price of land in
zone I;

$LRENT(K, I)$ = the location rent of land in zone I under housing
type K;

$STOCK(K, I)$ = the number of units of type K in zone I.

The average location rent is then combined with the current period's
expected land price in the adaptive expectations framework to form
the next period's expected land price. In early runs of the model
the highest location rent occurring in each zone was used as the
basis for land prices, but this procedure had to be abandoned
because land prices rose so dramatically over time. The weighted
average of location rents has produced much more reasonable
levels of land prices in the model.

After prices are formed, the remaining operations of the market-
clearing submodel are essentially bookkeeping details. Since

assignments of households have been made to the standing stock, the occupied stock and the number of available units are modified to reflect this. The residence assignments also generate work trips for each income class, so the work-trip array is updated.

The market-clearing submodel is the last of the seven submodels encountered in a model time period. After it is executed the simulation model has all of the updated quantities necessary to begin the simulation of the next time period. If simulation continues, control returns to the employment location submodel, and the next time period is produced.