ORGANIC ALGORITHM, ORGANIC URBANISM

By Travis Jared Marmarellis Bunt and Mathew Staudt, with Tat Lam and Timmie King Hong Tsang

Re-processing Beijing's hutong villages.

Published July 4, 2015

Qianmen alleyway in 2013. Photo courtesy of URB.

In early 2014, the URBANUS Research Bureau (URB)[1] was commissioned to study the historic Qianmen[2] hutong district in Beijing, south of Tiananmen Square. The dense nineteen-hectare site was originally constructed in the fifteenth century, and is home to more than five hundred residential buildings ranging from tiny informal dwellings to large courtyard villas in various stages of disrepair.

The Qianmen regeneration project was initiated by a state-owned enterprise with a dual mandate to achieve "social gain" and profitability.[3] In practice, though, the first priority of the venture was to ensure that the government's initial capital grant would fund regeneration of the entire hutong neighborhood over a twenty-year time horizon. Despite the client's expressed social objective, their initial preferred strategy echoed the tabula rasa approach common in contemporary Chinese land development. Large-scale demolition and fast-track construction would be phased so the first wave of luxury villa construction, built with all available capital, would sell quickly enough and at a high enough price that it could fund subsequent phases of ground-up construction. While relatively common, this type of operation is not only indiscriminately destructive, it carries considerable financial risk as well.

The disappearing hutong. New development closes in on Qianmen. 2001–2011, Google Earth images.

At URB, in collaboration with URBANUS's Shenzhen-based urban design team, the primary objective was to develop alternative strategies for redevelopment that would limit destruction of the existing neighborhood fabric and building stock. We initiated a parallel venture at the Hong Kong–based URBANUS Labs (URL)[4] to develop an analytic tool that could evaluate the implications of various program, phasing, lot-combination, and sizing strategies, and create an alternative "bottom-up" master plan. Our team established a sequence of analytic steps—a parametric workflow—that could evaluate potential correlations among geographic space, architectural volume, time, and money into an

integrated projective modeling tool: "Parse-L."

The first Parse-L prototype translates the design criteria of its users into a master plan for long-range development and a financial pro forma. The first step in its creation was for an URB field team in Beijing to input survey data, which included historic value, current use and ownership status, into Geographic Information System (GIS) software. This information was then compiled in 3-D modeling software to create an initial condition, or "year zero" base map, of the hutong. We next determined a series of design drivers—rule sets, preferences, and hierarchies through which the user defines design priorities—that would direct the area's growth. All inputs were then analyzed—to envision potential physical changes (building area, lot size, program), financial changes (value-added factors, gentrification factors) and a series of "actions" (buy, sell, build, renovate, rent, etc.) as outputs on a spreadsheet. These outputs were then recompiled to feed a twenty-year financial pro forma. Through this approach, the financial implications of different programmatic and phasing strategies could be projected. We could also apply different financial assumptions within the pro forma for a nuanced evaluation of urban redevelopment over time and under varied situations.

In addition to its underlying analytic and predictive potential, Parse-L's greatest innovative capacity lies in its ability to prioritize investment criteria (or "bang for the buck") while replicating distributed, organic development.[5] The two capacities are functionally interrelated within the model, where targeted investments could identify potential to initiate construction in particular areas.

Parse-L's workflow requires its users—designers, architects, or planners—to critically assess and quantify their presumed value structures in the early phases of a project. In this manner, Parse-L enables these users to apply parametric tools in accordance with their critical judgment to real-world scenarios that call for sensitive, studied solutions. Unlike many existing architectural parametric tools that breed a multitude of geometrically formal iterations, Parse-L models and projects abstract urban conditions to inform high-resolution planning strategies. It purposefully foregrounds the critical inputs from the design team and the generative potential of their strategic assumptions, as opposed to carrying out mechanical, undifferentiated applications. By responding directly to each parcel's unique attributes and projecting potential strategies in near real-time, Parse-L allows for a value-adjusted rendering of "highest and best use"[6] and an approximation of the long-term financial impact of opposing strategies.

Parse-L tests the potential for custom digital design tools to reconcile multiple design objectives and propagate unique design moves across diverse scenarios over time. The designer can introduce large-scale rules such as program distribution within a described zone, and introduce conditional rules that determine acceptable trajectories for a given lot based on its unique attributes. The role of the model itself is to reconcile these design drivers and produce a set of plausible solutions. In this sense, the model is designed to resolve and iterate this tension between small- and large-scale drivers.

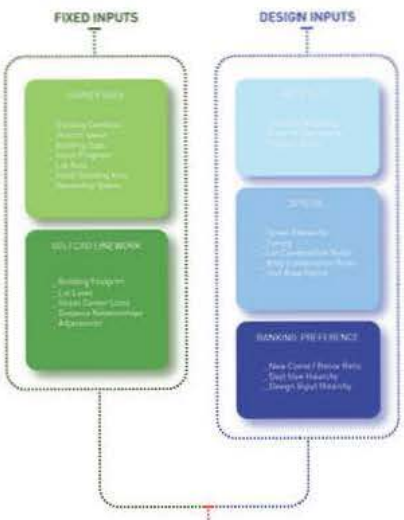

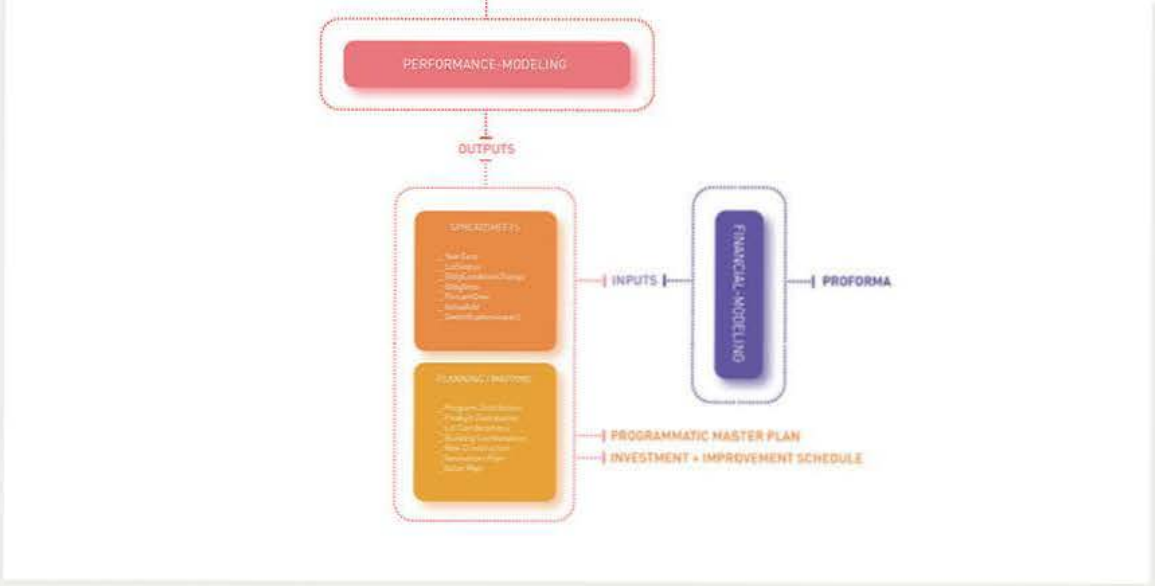

Parse-L workflow diagram. Bunt/Staudt, 2014.

The Parse-L process translates lot-based data from an "initial condition" to a "target condition" over simulated time. The initial condition is defined by any and all geo-data collected from the area. The target condition is then derived from reconciling user-defined design drivers. Therefore the model has two input types: fixed (to determine initial conditions) and variable (to determine target conditions).

Qianmen Building Stock, documented in field survey reports. Photo courtesy of URB.

FIXED INPUTS AND INITIAL CONDITIONS

During its multi-year land acquisition process, the client organization collected precise information about official historic designations, occupancy and property rights for each lot in the hutong. To supplement this data, the URB field team in Beijing surveyed the physical condition, architectural typology and historic features of each lot in Beijing. Data from these two sources comprised the basis of the model's "fixed inputs."

In incorporating this data into Parse-L, special focus was given to assign attribute types that could be directly compared through a discrete ranking method. The historic significance of each lot was scored on a scale of zero to three: from lots with no historic relevance[7] to ones with highly historic or officially protected buildings.[8] Similarly, the physical condition of structures on each lot was scored from zero to three: from vacant or ruined sites to excellent building conditions requiring no significant repair.[9]

The ranked survey data was next consolidated and mapped[10] to a spatial database recording information for each lot.[11] Data was organized within six parallel object classes:

lots, zones, streets, buildings, units, and courtyards.[12] Just as they have within the fabric of Qianmen, individual objects within these classes maintain spatial relationships with other objects both within their class (i.e., from building to building) and with those in other classes (i.e., from a unit within a building to a location on a certain street). When the model is initialized, Parse-L automatically computes and stores this relational data by constructing what is called a "mesh network topology"; each object recognizes its relative place within the urban fabric and cooperates in sending, receiving, and storing data. The mesh connections represent a computed relationship between any two elements.[13] In this way, Parse-L builds spatial intelligence into the model.

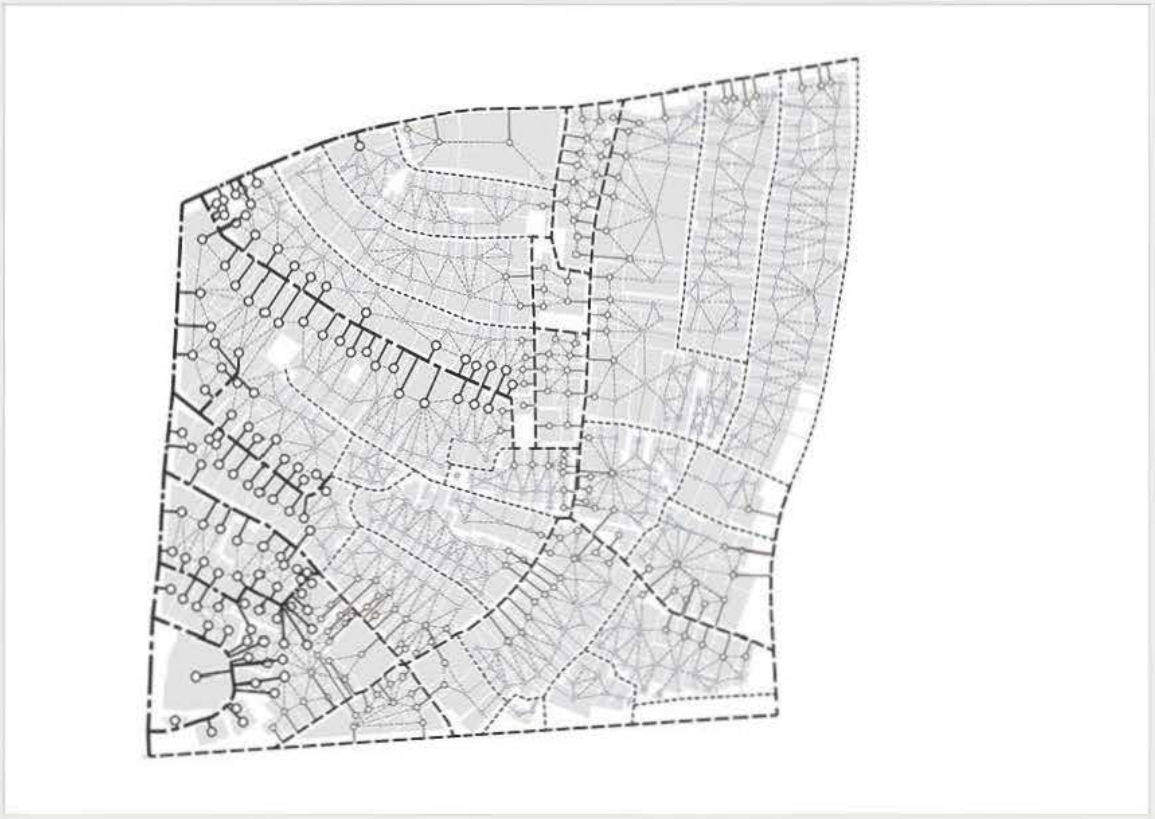

The model automatically builds and stores spatial relationship data using a mesh network topology. Bunt/Staudt, 2015.

The mesh is periodically updated during simulation. By entirely automating these calculations, this tool can be easily adapted for other urban environments. Fundamental changes to the urban fabric (such as new roads or other infrastructure) can be applied during simulation, and all relational data would be recalculated on the fly. We've utilized these spatial relationship networks in several ways, as explained in the following paragraphs.

The first layer of the mesh network simply determines lot-to-lot adjacency.[14] For Qianmen, the client's goal was to maximize the selection (or creation) of lots suitable for new luxury villas with an optimum parcel area of approximately five hundred square meters. The algorithm finds the greatest number of suitable lot combinations for this housing type while maintaining the existing urban fabric as much as possible. After culling all parcels not suitable for new ground-up construction (because they contain historic structures, buildings in relatively good condition and so on), the network of lot adjacencies is analyzed to find an optimal solution of combined lots.[15]

Another network layer identifies adjacencies between buildings and courtyards. As in the lot adjacency network, this data allows us to optimize potential unit combinations while respecting partially occupied buildings. By treating the courtyard as a discrete agent on the same level of data as a whole building, the courtyard's status as a central, spatially organizational tool in a traditional hutong can be fully incorporated. Within this structure, for example, two units that share adjacency through a courtyard can be combined even if they do not necessarily share the same lot.

Additional network layers, of spatial relationship such as streets and building stock, enable Parse-L begins to project potential ripple effects of new interventions. For example, our hypothesis was that localized and strategic urban improvement can accomplish the social goal of supporting cultural continuity on site while facilitating financial sustainability, as unimproved lots can gain value through proximity to those receiving direct capital investment. By investing judiciously in properties with the greatest catalytic potential, rising property values across a broader area can be achieved with an efficient capital allocation.

To measure this "gentrification" phenomenon, we created a method for calculating and projecting property values. Direct improvements to one parcel are assumed to have a halo effect in increasing surrounding lot values.[16] That is, if Lot A is redeveloped and sits in relatively close proximity to Lot B, B's value also increases. The closer B is to A, the more its property value is affected. This is not a simple geometric equation, of course. The measured distance between two parcels must take into account the topology of its urban area as well as its address—the primary face of a parcel's building on a street—and must be calculated according to the street network.[17] To achieve this, we constructed a lot-street network to calculate the walking distance between any two parcels using Dijkstra's shortest path algorithm.[18] During each year analyzed in the simulation, a property value adjustment factor (or gentrification effect) is calculated and stored for subsequent financial analysis.

Parse-L Lot Optimization Process." Animated GIF. Bunt/Staudt, 2014.

DESIGN INPUTS AND SIMULATION

The translation of each lot's "initial condition" to a "target condition" is directed by user-defined design drivers, which give the designer multiple points of entry to influence the model. The most significant of these is defining a "target product typology": a classification system defining all possible land uses and/or architectural types for any given lot. In other words, it explicitly defines all of the designer's intentions. We use the term "product" here as these definitions can identify more than land use—in the case of Qianmen, these categories also designated the type of construction (renovation or new), and were tiered based on the economic class of the intended occupant ("elite" or "moderate," in reference to single-family residences).

For each product defined, the designer declares acceptable target parameters for each

attribute. For example, the user can define minimum and maximum lot areas, historic significance levels or building conditions necessary for new ground-up construction to be initiated. In the case of Qianmen, these parameters were particularly crucial for the URBANUS team because they could demonstrate how major new construction can take place without sacrificing traditional urban form or significant, existing structures in the hutong.

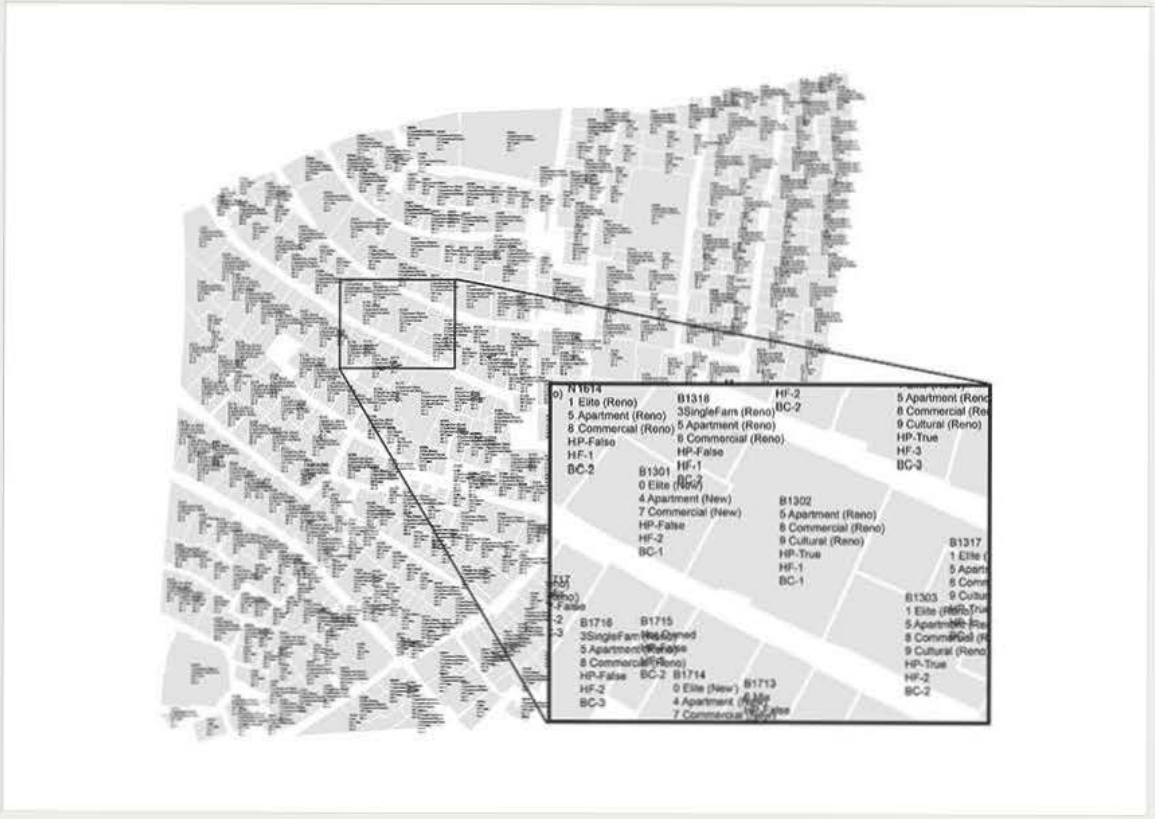

Based on the lot area, the buildings' physical condition, and historic significance, all potential products for each site are determined." Bunt/Staudt, 2014.

It is important to note that fixed and designer inputs are used in a deductive[19] manner, in which data is used to rule out unsuitable solutions. The designer can reduce the field of possible product options by inputting regional typologies—territorial overlays that define priorities within each zone—to the model. For example, if the designer wishes to ensure that an area has a fair amount of multi-family residential and community-based retail, he or she can input a zone boundary[20] and assign a high target percentage value to all multi-family residential products.

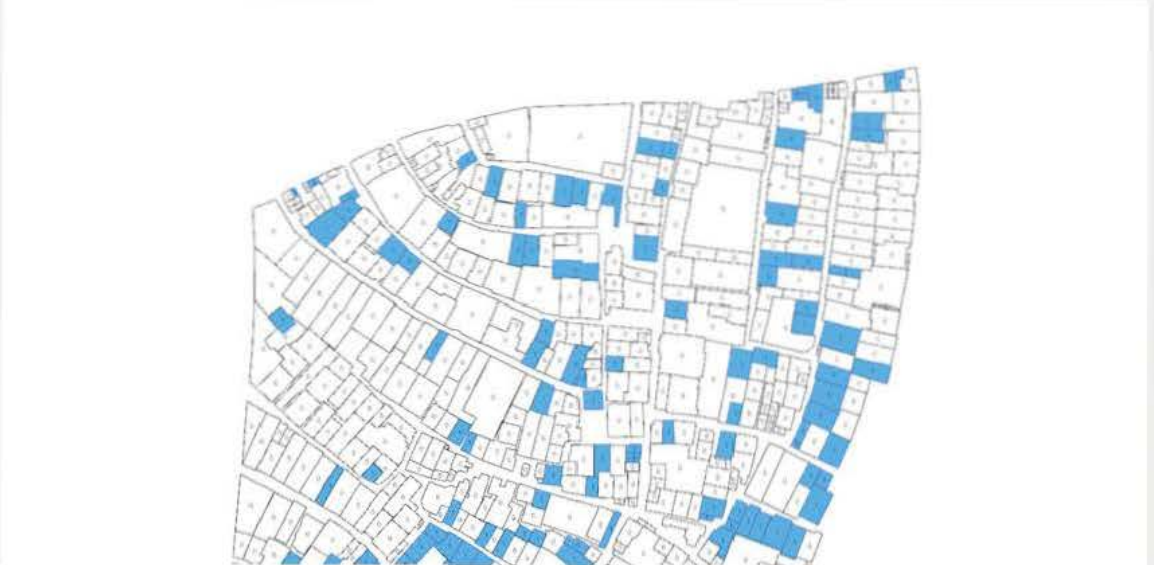

Parse-L Lot Determination in Rhino/Grasshopper. Animated GIF. Bunt/Staudt, 2014.

The designer also defines a series of global thresholds that govern decision making, such as the calculation of lot values or the phasing of a project. Specifically for Qianmen, the client's goal was to maximize the number of lots suitable in scale for new luxury villas. Within Parse-L, an optimization function reorganizes lot boundaries to maximize the number of lots within this target area range.[21] The designer determines how aggressive lot transformation is by parametrically limiting the lots considered.[22]

According to a user-defined phasing strategy, the model distributes all proposed interventions along a twenty-year timeline and simulates implementation of the design.[23] Upon completion, Parse-L outputs a database representing the proposed masterplan. In addition to outputting the proposed land use (or product) for each lot, the Parse-L masterplan is a lot-by-lot and year-by-year prescription of "actions" taken to execute the given design. Additionally, the model outputs relative financial data chronicling a measured anticipated impact of each action taken. This database, stored in the form of spreadsheets, is the master plan in its entirety. All subsequent steps taken are calculations and representations. For this, the data is passed to two secondary models: a visualizer to draw maps and a pro forma model for financial analysis. [24]

Parse-L 20-Year Action Strategy. Animated GIF. Bunt/Staudt, 2014.

Parse-L Lot-Based Master Plan. Bunt/Staudt, 2014.

RESPONSIVE FINANCIAL PRO FORMA

Output data from the simulation is populated on a series of spreadsheets within an automatically updating financial-projection model. The foundational sheet, documenting initial conditions, is produced in the first step of algorithmic processing described above, essentially transcribing both the survey data from GIS and the quantitative design inputs through to Excel to form a Year Zero foundation for financial modeling. The remaining sheets, produced at the end of the simulation, are formatted to give an instruction (or data point) for each of twenty years for each individual lot.[25]

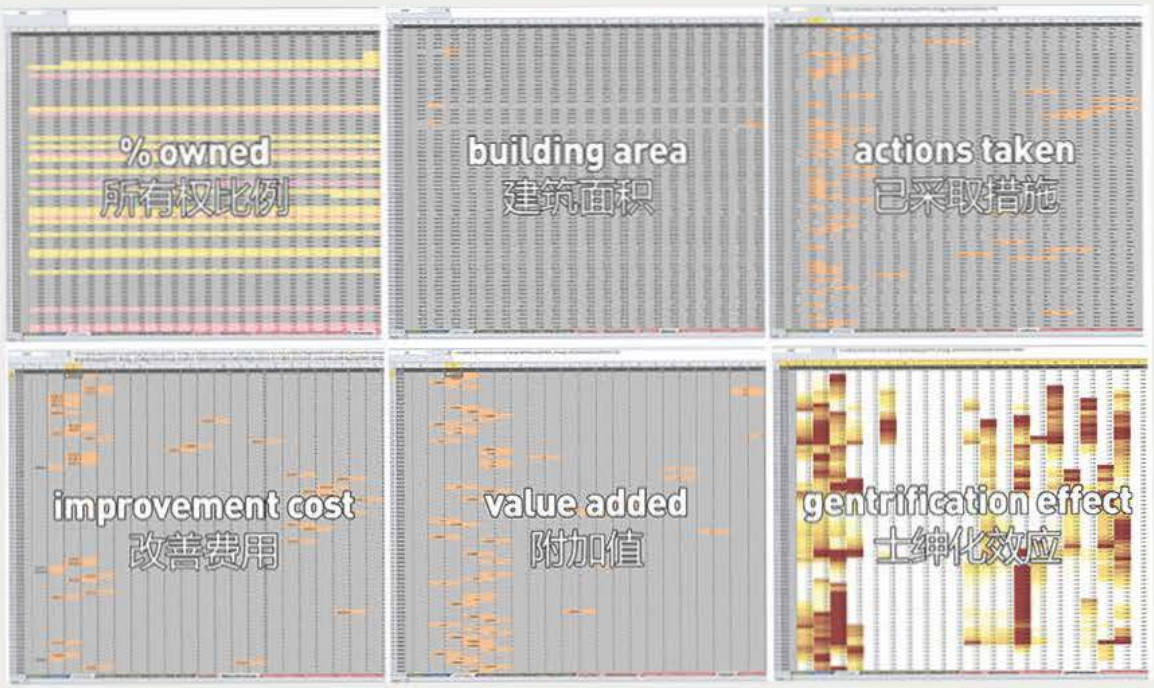

Parse-L Pro Forma inputs. For clarity, conditional formatting in Excel highlights cells that represent a change in condition on the prior year. Bunt/Staudt, 2014.

These auto-populated input sheets translate spatial actions from the simulation into quantifiable bits. The financial model is structured so that each sheet, with its matrix of lot-year data points, can function as an additional variable to be calculated in projecting future values. As such, the model is expandable indefinitely—new design input variables (or a new category of survey data) can be layered on, sheet by sheet, as the simulation adds complexity. In the current configuration, they track actions taken on an individual lot, changes to a lot's building condition, total built area, ownership percentage, value increase due to direct improvements, and underlying appreciation due to area improvements. These matrices are then cross-calculated to produce an estimated real estate value for each lot in any given year.

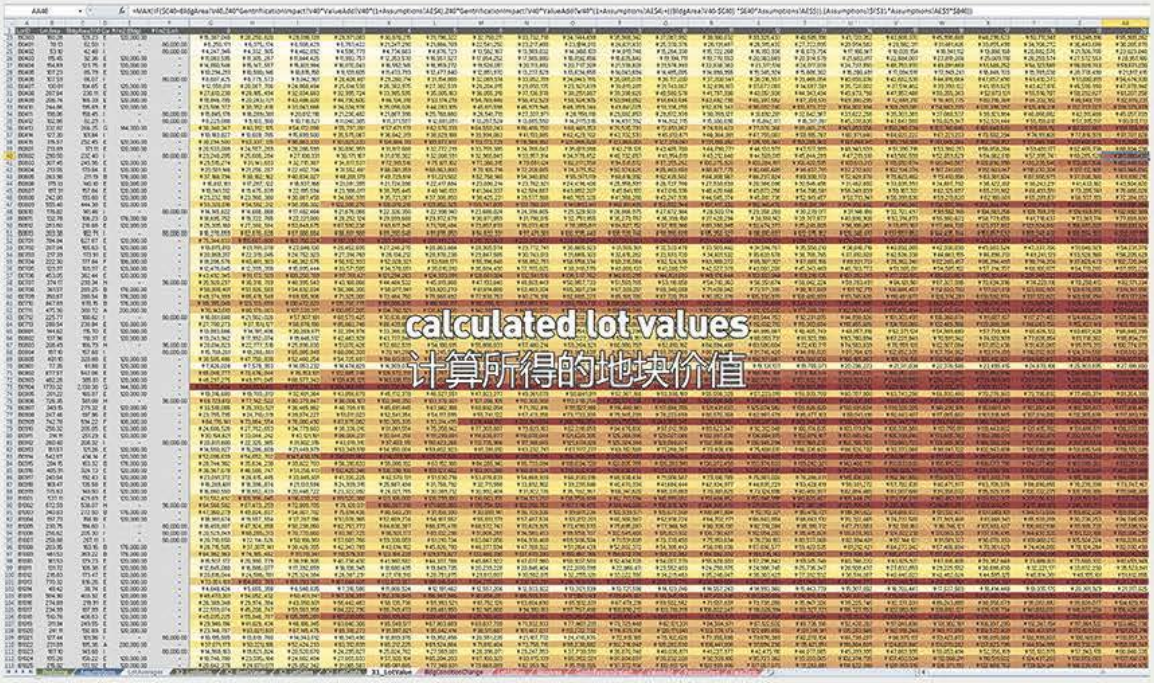

Parse-L Pro Forma, Calculated Lot Values. To assist the user in spotting trends—or potential errors—conditional formatting is used to "heat-map" the resultant values. Bunt/Staudt, 2014.

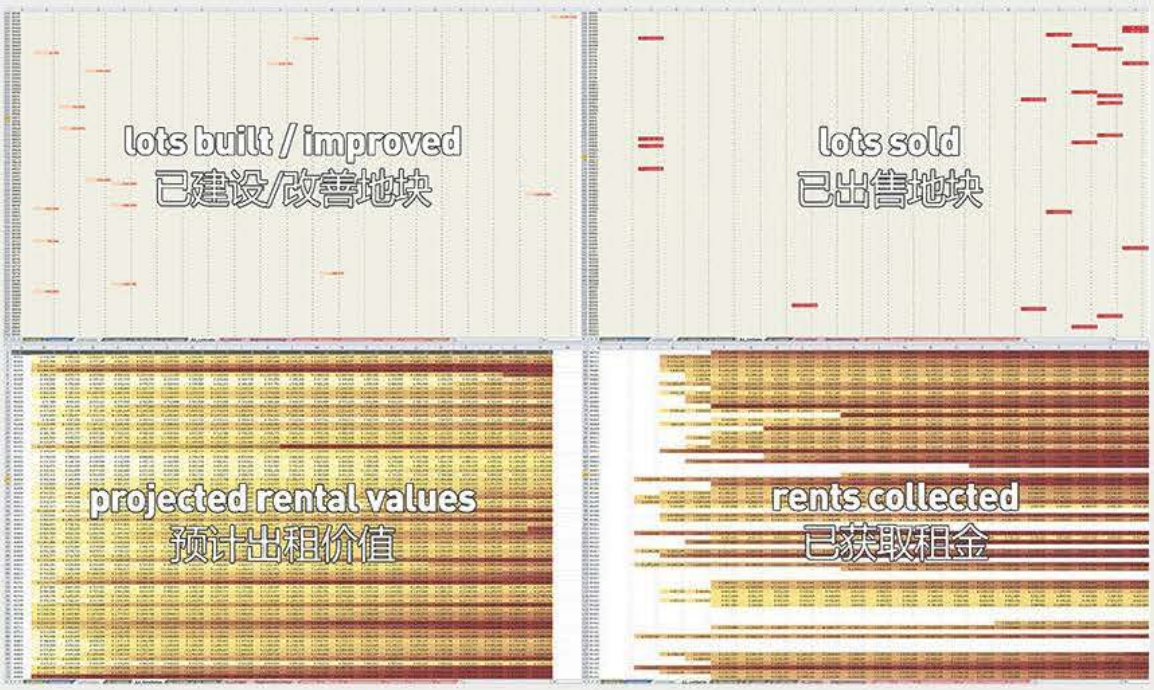

Parse-L Pro Forma, Secondary Calculations." Bunt/Staudt, 2014.

The final pro forma tallies all inflows and outflows (including land cost, infrastructure improvement cost, lot improvement costs, rental income, and sales income) from year zero to twenty, and calculates the net operating income (NOI), internal rate of return (IRR) and annualized return on investment (ROI) for the simulation. The NOI, IRR, and ROI values can then be quickly compared across different simulations to test the financial impact of differing inputs. These are "heat-mapped"[26] with conditional formatting to be understood graphically.

IMPLEMENTATION: TESTING THREE DESIGN APPROACHES

For Qianmen, we simulated three different schemes. The first followed localized criteria on a lot-by-lot basis. The second and third were site-wide "master-planned" alternatives set by the designers' preferences—one prioritized the assignment of program and phasing according to specific street corridors, or spines, and the other implemented strictly

defined neighborhood zoning. Parse-L allowed us to compare all three schemes in terms of their projected financial outcome.

Simulation A: Localized Scheme

The first simulation was intended to replicate an organic "gentrification" scenario, defined by a hierarchy of default values and planning policy nudges instead of any dominant, guiding hand of the designer. This was understood as a neutral setup in which no particular classes of objects took precedence. The existing conditions of the Qianmen site (such as lots and roads) were kept intact except for potential lot mergers and infrastructural upgrades necessary for basic dwelling needs. There were no defined interventions (such as a retail-oriented street or a luxury neighborhood), which meant the simulation ran based solely on the user-inputted value hierarchies and program distribution targets. As such it could also act as a control scenario against which other approaches would be compared and gauged.

In the base model, inputs defining "product types," "construction types," "proximity parameters," and "zone products" were established. Working with the client, the URBANUS team conceived of seven program categories that encompassed twenty-nine types of end-state products. Each product shared a common set of parameters that allowed them to be evaluated against the corresponding categories of "site survey" (e.g., "maximum/minimum lot area" and "maximum/minimum building condition") and "street rules" (e.g., car accessibility). The acceptable values were then designated for each parameter for each product. As with all iterations, estimations, and assumptions of existing lot value and the value added by each product type were declared.[27]

Another category, "construction types," outlined six possible strategies, impacts and purposes of potential construction. For example, a lot that is historically protected could only undergo "full restoration" or "combined restoration and new construction."[28] Proximity parameters were then set to allow for lot mergers and to measure the influence of value appreciation effects. In this case, any two lots could potentially combine if the resultant area was less than 500m^2 and both lots' building condition and historic feature ratings were at one or below. Target program and product ratios were also set for the entire site thereby minimizing top-down influence on the resultant design.[29]

In Simulation A, there is no overlaid zoning, so the primary spatial rules concerned the influence of existing street-based drivers, including street condition (level of improvement), capacity (pedestrian alley to two-way vehicular road), infrastructure (public utility lines), and dominant program (commercial, residential, etc.).

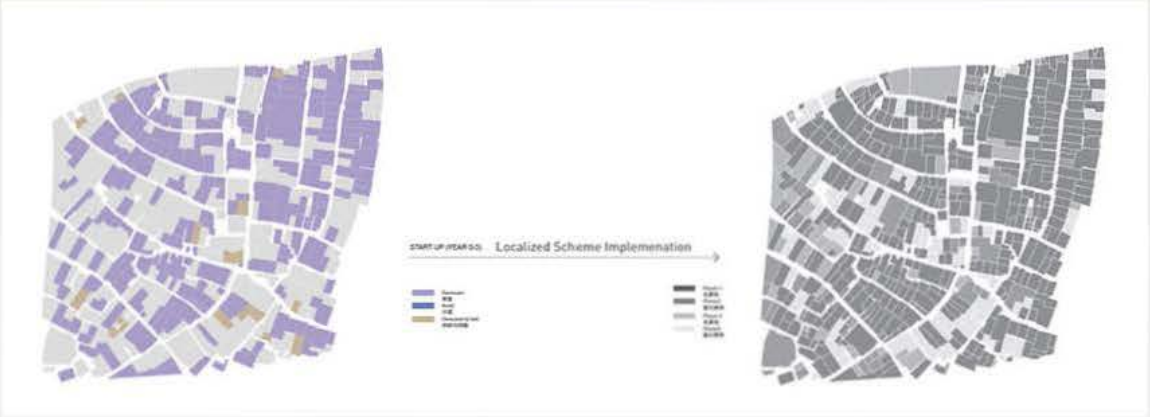

Parse-L Localized Scheme Implementation." Bunt/Staudt, 2014.

The Localized Scheme was a primary testing ground for us to develop and finesse the Parse-L model. After modifying several dozen iterations of the model's underlying parameters (rebuilding from the ground up once or twice) and experimenting with various input values, we were able to reliably simulate a distributed, individualized redevelopment of Qianmen. Interestingly enough, the more the functionality (and logical coherence) of the model was improved, the less the user's value hierarchies mattered to the end result. Simply put, when you allow the lot to decide what it wants to be, there are only so many

rational options available for it.

Nevertheless, though the end products did not vary significantly from simulation to simulation, the real opportunity for variation under this scheme was in phasing. By applying different strategies to the timing of infrastructure improvements and decisions to rent or sell, we could readily test which approaches were likely to offer the greatest return on investment. In its most recent iteration, Parse-L advocates renovating and renting early to boost property values, and saving larger-scale luxury construction for later years, when sales will demand the greatest premiums.[30]

Simulation B: Street-Based Scheme

The second simulation followed an urban design strategy for creating defined programmatic corridors throughout the area. A green belt would prioritize recreational amenities, parks, and cultural institutions, and a "Main Street" would prioritize retail and commercial usage. Lots that did not have frontage on either road would prioritize residential construction or renovation. Chronological phasing would favor earlier development on the two primary axes.

The key drivers in this scheme were the two spines, hence inputted street variables were the primary vehicle for the designers' intervention. By first defining specific values for what makes a green belt or Main Street, then nominating certain streets within the model to follow those parameters, the simulation could recalibrate projections based on the new priorities.

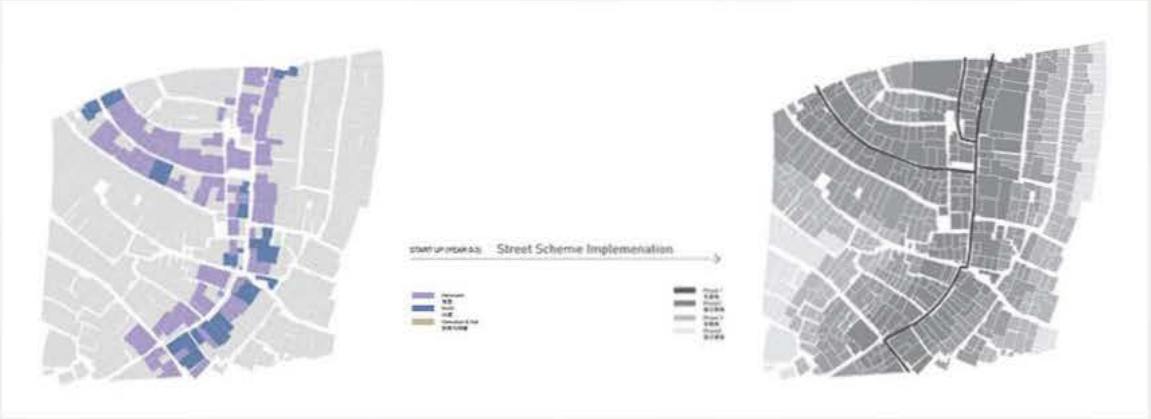

Parse-L Street Scheme Implementation." Bunt/Staudt, 2014.

As anticipated, the end-state projection of this scheme showed different behaviors in the lots adjacent to the spines compared to those farther away. One somewhat unexpected result was the marked increase in demolition carried out along the spines to make way for new cultural, recreational, and commercial developments. These new developments also gave rise to a lot value "heat map" that radiated from the two main spines. Unlike Simulation A, which through its wider distribution creates a proverbial "rising tide lifts all boats" situation, the focus of improvements in Simulation B creates corridors of high value in some areas while leaving significant portions of the neighborhood undervalued.

Simulation C: Pocket Scheme

The third, and most heavy-handed approach, simulated a conventional planning method in which the site was divided into different zones, or micro-neighborhoods, each with a different preference in terms of scale and land use. In each of these zones, priority would be given to certain programs and products. For example, in the "boutique living" zone, multiple small lots would be assigned a residential program (below minimum thresholds used in other zones), while "luxury house" and "smart living" zones would prioritize larger-format residential products (villas in the former, multi-family in the latter). Chronological phasing would be assigned and prioritized by zone, and could also be adjusted to favor certain "start-up zones" to test which micro-neighborhoods to invest in first.

The driving factors in Simulation C are the delineation of distinct zones and their

The user would be required to clearly define the boundary of each zone and create differing priorities for each; in this simulation, seven distinct neighborhoods were created.

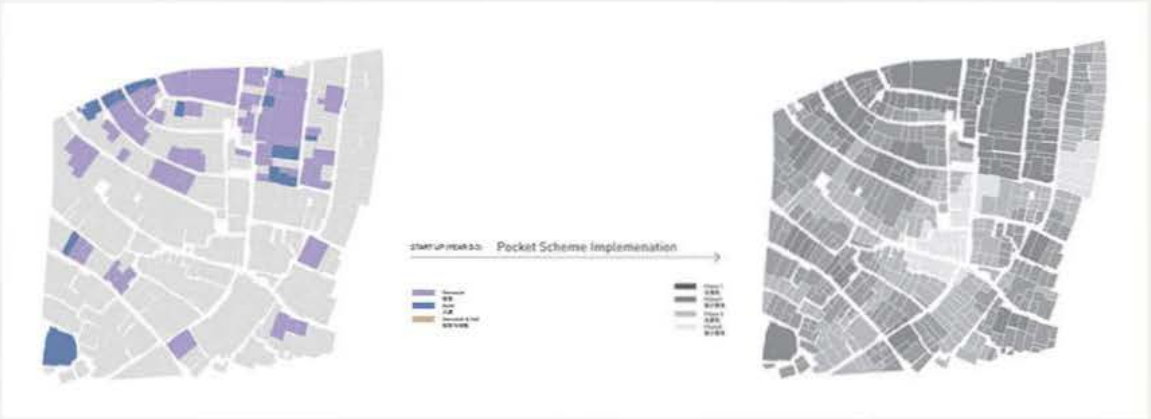

Parse-L Pocket Scheme Implementation." Bunt/Staudt, 2014.

The resultant maps matched the team's expectations by concentrating particular programs and products in ways that followed zone priorities. Financial maps also showed that higher values began to radiate from the spots where high value-added products, such as an elite luxury house, were assigned. However, as with Simulation B, value-generation was concentrated in the zones drawn by the designer.

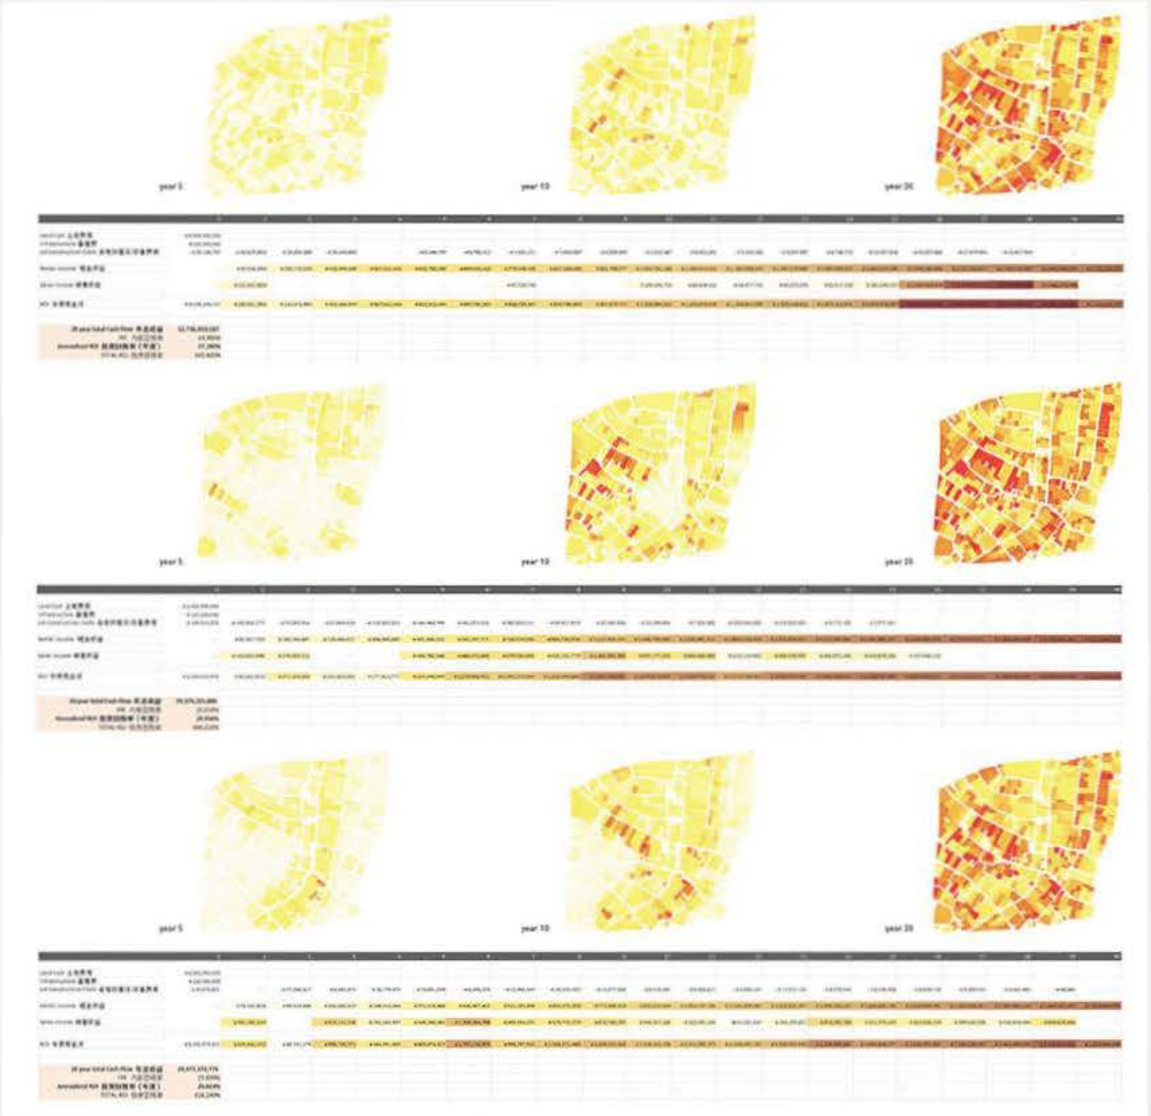

Parse-L 3 Scheme Pro Forma Comparison." Bunt/Staudt, 2014.

COMPARISONS AND CONCLUSIONS

By overriding a more straightforward application of "highest and best use" exemplified in the first simulation, the second and third simulations delivered a less varied mix of program and a less profitable result. This is to be expected, given that, as explained above, design inputs are necessarily associated with the *removal* of potential options. Besides the obvious formal distinctions between the three simulated approaches, what is critical about the two master-planned strategies is the manner in which they link phasing to their respective intervention. In other words, for the Street-Based Scheme to function as intended by the designer, the nominated green belt and main street areas would demand early-stage investment. Likewise, the Pocket Scheme would require an ordered succession of targeted development zones. Design intent in these two schemes precludes the use of financial incentives to determine the timing of infrastructure investments. Invariably, both consistently underperform the Localized Scheme in their return on investment.

In short, Parse-L delivers as much variability into the neighborhood plan as the user allows. Even the comparatively heavy-handed schemes B and C are truly hybrids, improving considerably on any tabula rasa alternative both in terms of design sensibility and financial return. Increasing localized variability, while certainly requiring more planning and management effort on the part of a development company, creates a better product and a more profitable project over the long run.

Operating in China or elsewhere in the developing world—where both the ethos of preservation and legal protections for historical sites can be weak—traditional ethnographic appeals are often insufficient to prevent the wanton destruction of existing building stock executed in the pursuit of Gross Domestic Product growth. In this arena, a financial and operational case must be made. The creation of a tool like Parse-L is a step toward disconnecting preservation strategies from social arguments alone, while facilitating the same ultimate goals. When an approach of preservation, restoration, and adaptation can stand on its own as an economic case rather than a philanthropic one, it is far likelier to see success.

---

1. Established in 2011, the URB is a research-oriented department of URBANUS Architecture and Design in Shenzhen, China. ^
2. Originally Zhengyangmen, 正阳门. ^
3. Due to the political sensitivity of hutong redevelopment, the name of the client organization as well as the specific terms of the consultant contract remain confidential, per government policy. The appropriately generic project title was 前门东区城市研究项目 (Urban Research for the Eastern Qianmen District), and URBANUS was one of multiple consultants tasked with different portions of design and research for the larger area. ^
4. URL is the sister organization to URB. It is located in Hong Kong and focuses on "extra-architectural" experimentation in addition to traditional urban research. ^
5. "Organic" in that actions run "bottom-up" from a series of individual actors making individualized decisions, rather than "top-down" as in a traditionally imagined master plan. Actions within the model occur in response to the specific conditions of each lot, building, street, and courtyard. ^
6. "Highest and Best Use" is a standard term in real estate appraisal that refers to the maximum income-producing potential of a given property within the law. ^
7. In the case of Qianmen, a lot with no buildings or in ruin beyond repair is given a *historic feature* rating of 0. Any existing building has at least a *historic feature* rating of 1. ^
8. It should be noted that China's landmark/preservation statutes differ in major respects from Western approaches. In many cases, new ground-up construction that replicates a demolished historic structure is an allowed "preservation" strategy (and in practice, often encouraged). ^
9. While such an explicit and immediate ranking system could be potentially problematic, this method gave us a tool to help determine which lots had the most potential for strategic redevelopment versus preservation. Our position was that while these designations should be disputed and scrutinized, if this ranking is done thoroughly involving all stakeholders, such a ranking system can be productive in determining the best fate of each lot. ^
10. There is a not-insignificant amount of data cleanup that must be done manually as part of this process. An obstacle to achieving a fully automated workflow is the nature of most urban CAD plans (particularly in China) where lot boundaries are typically assembled with irregular polylines of various directionalities and segment lengths. Hence some manual interventions had to be made in order to attain consistent orientations of the geometries. Efforts are under way to explore better methods of geometry generation and more sophisticated form finding for Parse-L 2.0. ^
11. Where information needed to be expressed at a higher resolution than that of the lot, a subset of spatial data was recorded for individual buildings and units. For example, "historic protection" measurements are logged at the scale of the lot while "property ownership" is recorded at the scale of the unit. For the sake of our model, it was understood that any lot containing at least one officially protected structure would need to be treated with special sensitivity to historical context. Therefore, we labeled an entire lot as historically protected if at least one structure carried this designation. But when dealing with building occupancy and property rights, we felt a higher degree of resolution was necessary. For the purpose of modeling real estate exchanges, we found representing ownership and occupancy as a percentage of the entire lot was sufficient, but in the case of optimizing lot combinations, partial ownership and/or occupancy had to be described as geometric areas within each lot. ^
12. The classes are structured in parallel (i.e., non-hierarchical) such that design interventions can be made on multiple scales. The model is built such that the designer can input driving parameters at multiple scales (i.e., *lot, building, street, zone,* or *global*). In order for these drivers to have influence globally, model functions were designed to flow between object classes, allowing design drivers from one scale to influence another. From tenant units to circulation networks, the division of parcels does not necessarily limit the morphology of the urban fabric or limit design goals. In addition, the definition of lots is dynamic and can be modified during the simulation process. One functionality of the model is to reorganize lot boundaries in order to optimize lot areas. This process is discussed further in other sections. ^

13. For example, in the case of the *lot-lot network*, the mesh nodes represent center points of each parcel and a connection exists between two nodes if the two lots are adjacent. ^

14. Each lot was represented by a closed curve in Rhino which was assigned a distinct "Lot ID" (from ArcGIS) as its Rhino object name. This Lot ID was then utilized to pull out information from the corresponding row in the Site Survey. The curves and their corresponding survey data were later fed into Grasshopper in which an object of the Lot class was instantiated for each Lot, such that each Lot object in the script contained its corresponding survey data. ^

15. In this particular scenario, the total FAR for the entire neighborhood was fixed, meaning GFA could shift around but the total leasable area across all lots could not be increased. The height limits of hutong areas near Tiananmen Square are also strictly enforced, making this primarily a two-dimensional exercise. Given a scenario in which an increase in FAR and/or height was allowed, this portion of the modeling would naturally be calibrated differently. ^

16. In this case a simple linear fall-off was used. How this urban phenomenon could be modeled mathematically in a more precise way is beyond the scope of this research. Throughout the development of the tool, resources were focused on theorizing a broad stroke workflow that can help designers leverage urban data in their conceptual thinking. With a broad framework in place, calculation techniques can be refined in later iterations, such that an entirely new gentrification algorithm could be plugged in. ^

17. In addition to calculating walking distances between parcels, the parcel-street network is also utilized to establish what street each lot is most associated with. This gives the designer the ability to input design drivers at the level of the street. For example, a street can be designated for a particular program or target building type. The model can then use this information to narrow down what target types are suitable for each lot. ^

18. Dijkstra's algorithm, or "The Shortest Path Problem," as used in Parse-L is directly based upon open source code incorporated from Runestone Interactive at: http://interactivepython.org/courselib/static/pythonds/Graphs/graphshortpath.html (accessed June 1, 2014). ^

19. Where most urban design works as an inductive process, building up a solution from the extrapolation of existing context, Parse-L's process runs deductively, beginning with all theoretically possible end-states and eliminating unsuitable or undesirable outcomes according to the applied rule-set. ^

20. A region can be defined either by drawing individual outlines and/or by street. If the user classifies zones by street, lots will are considered within that street's "zone" if it is adjacent to the referred street. (Utilizing the lot-street network.) ^

21. In this case, the target range was {400,600}m^2. ^

22. Lots with historically significant structures, structures in relatively good condition and partially occupied, or partially owned by a third party were excluded from lot reorganization (*historic feature* > {1}, *building condition* > {2}, *percent owned* = 1, *num occupied units* = 0). ^

23. The original Parse-L model was built around a twenty-year projection of financial and construction developments. However, the client organization deemed this standard property development time horizon to be incompatible with the particular political circumstances in Beijing, where government officials typically have a five-year tenure in a given office. The organization must therefore present a convincing five-year case in order to get approval; any projects over five years run the risk of political jeopardy. A quick workaround was done by redefining the twenty-year timeline as twenty quarters (thus five years), and adjusting the affected categories such as appreciation multiplier, lot status actions, etc., to fit with the shortened time frame. Parse-L 2.0, currently under development, will build in intelligence and flexibility so that the user may specify the length of the projection period (in years), and the resolution of output (in years, quarters, or months). ^

24. All numeric data is formatted in terms of scalable ratios rather than explicit values. By doing this, multipliers and real value assessments can be applied independently. ^

25. With the exception of a user-inputted Assumptions sheet (a necessary component to any pro forma), which sets financial baselines, all other data flows from the prior simulation processes. ^

26. To "heat-map" a set of data, Excel considers the full range of values and assigns colors along a gradient, with each cell rendered relative to its rank within the spectrum. ^

27. Existing lots were classified into eleven categories according to their "Historic Feature" and "Building Condition" rankings. A Value Rating for each category was also calculated with the formula (HF*0.3+BC*0.7), which would in turn be multiplied by a baseline value (set at ¥80,000 for this simulation) to arrive at the lot value/m2. Each product was assumed to improve the building condition of each lot by a fixed amount. Using the same formula, a value rating of each product was calculated. The result was a ranking of product types and existing lot categories in terms of value rating and lot value/m2. This ranking was then fed into the model through Grasshopper for lot value calculations and output to pro forma. ^

28. Interestingly, it was this particular function that garnered the most interest from the client organization. Immediately after the design team at URBANUS Shenzhen presented the three original iterations, a fourth approach was requested, this time by prioritizing Construction Types. As an effort to preserve the appearance of the neighborhood, the design team and the client organization agreed that "Hutong Façade Condition" should be added to the analysis and given a priority as high as those of Historic Feature and Building Condition. From the developer's point of view, Construction Types were most easily quantifiable in terms of initial investment. Therefore an alternative workflow was created in which Construction Types were first determined by Hutong Façade Condition, Historic Feature, Building Condition, and other design criteria. Depending on its parameters, a lot would be assigned an Interior Court Construction Type and a Hutong Façade Construction Type. Product types were in turn assigned based on these two Construction Types and other factors such as program distribution. The modified workflow was more linear than the original Parse-L, but it nevertheless provides insights into how real developers might adjust their applications of the model based on priorities in different sites. Plans for Parse-L 2.0 would add functionality that allows the user to choose from a flat hierarchy, a Construction Type prioritized workflow, or a Product Type prioritized workflow. ^

29. Although no top-down zoning was envisioned for the first scheme, it was still necessary to declare the whole site as a generic zone so that rules for program and product distributions could work. A zone was also represented in the Rhino space by a closed curve. Rhino layer names were used to differentiate between zones. In this simulation, all zone curves were assigned to the same layer (0. Anchor Core) and were in effect one generic zone. This assignment allowed the script to extract information from the corresponding "Zone Products" spreadsheet in the Program and Products Rules document. ^

30. While this may sound quite obvious to Western developers, it is generally the reverse of the schedule favored by Chinese development companies —including the Qianmen client organization, as detailed in the introduction to this article. ^

TRAVIS JARED MARMARELLIS BUNT is the principal of WORKSHOPOLIS, a consultancy specializing in the urban development of emerging economies. In 2014, he was recognized by *Perspective Magazine* with a 40-under-40 award as a "rising star" in the architecture and design industries, in part for his work establishing and leading the Hong Kong office of URBANUS and his role in curating the Hong Kong edition of the 2013–14 Bi-City Biennale of Architecture and Urbanism. He is an adjunct assistant professor for the real estate development program at GSAPP and the author of *A Life Away: Letters from Afghanistan*, which reflects on the U.S. military's role as nation-builder to the world.

MATHEW STAUDT is a designer specializing in computational techniques in design and fabrication. He works with artists and architects to realize custom digital workflows, design interfaces, and assembly systems. Mathew has been a member of several progressive studios including THEVERYMANY, MOS, Ogrydziak Prillinger Architects, and

URBANUS Labs (URL). As a technical consultant, Mathew's recent clients include Snohetta, AM Design, CW-ZW Design, Pliskin Architecture, and Studio SHAN SHAN SHENG.

TAT LAM has been the director of URBANUS Research Bureau (URB) since 2010, at URBANUS Architecture and Design Inc. based in Shenzhen, Beijing, and Hong Kong. In 2007, Dr. Lam was appointed research scholar at China Lab, a research group out of Columbia University GSAPP, where he has been collaborating with scholars on the research and publication of mega-block developments in China. He is currently an assistant professor teaching at the School of Architecture at the Chinese University of Hong Kong (CUHK). Dr. Lam additionally leads independent research projects on rural urbanization issues as part of the grass-roots platform Shanzai City.

TIM TSANG is a master of architecture candidate at Columbia University's Graduate School of Architecture, Planning and Preservation. He joined URBANUS Labs as a research assistant, where he worked on the development of the Parse-L model and the application of computational tools to real-world architectural problems. At GSAPP, his research focuses on the integration of computation with theory and practice. Prior to Columbia, Tim practiced in Hong Kong with a specialization in institutional projects. He received his bachelor of social science in architectural studies from the Chinese University of Hong Kong.