http://pubsonline.informs.org/journal/opre/

OPERATIONS RESEARCH
Vol. 67, No. 4, July–August 2019, pp. 950–964
ISSN 0030-364X (print), ISSN 1526-5463 (online)

CROSSCUTTING AREAS

Technical Note—Optimizing Foreclosed Housing Acquisitions in Societal Response to Foreclosures

Senay Solak,[a] Armagan Bayram,[a] Mehmet Gumus,[b] Yueran Zhuo[a]

[a] Isenberg School of Management, University of Massachusetts, Amherst, Massachusetts 01003; [b] Desautels Faculty of Management, McGill University, Montreal, Quebec H3A 1G5, Canada

Contact: solak@isenberg.umass.edu, https://orcid.org/0000-0002-6182-1852 (SS); abayram@som.umass.edu, https://orcid.org/0000-0003-3933-5388 (AB); mehmet.gumus@mcgill.ca, https://orcid.org/0000-0003-3814-896X (MG); yzhuo@som.umass.edu, https://orcid.org/0000-0002-0595-452x (YZ)

Received: August 13, 2013
Accepted: September 29, 2018
Published Online in Articles in Advance: June 24, 2019

Subject Classifications: government: programs; industries: real estate; planning: urban
Area of Review: Policy Modeling and Public Sector OR

https://doi.org/10.1287/opre.2018.1827

Abstract. A dramatic increase in U.S. mortgage foreclosures during and after the great economic recession of 2007–2009 had devastating impacts on the society and the economy. In response to such negative impacts, nonprofit community development corporations (CDCs) throughout the United States use various resources, such as grants and lines of credit, in acquiring and redeveloping foreclosed housing units to support neighborhood stabilization and revitalization. Given that the cost of all such acquisitions far exceeds the resources accessible by these nonprofit organizations, we identify socially optimal policies for CDCs for dynamically selecting foreclosed properties to target for potential acquisition as they become available over time. We evaluate our analytical results in a numerical study involving a CDC serving a major city in the United States and specify social return-based thresholds defining selection decisions at different funding levels. We also find that, for most foreclosed properties, CDCs should not offer more than the asking price and should typically consider overbidding only when the total available budget is low. Overall, comparisons of optimal policies with historical acquisition data suggest a potential improvement of around 20% in expected total impacts of the acquisitions on nearby property values. Considering a CDC with annual fund availability of $4 million for investment, this corresponds to an estimated additional value of around $280,000 for the society.

Funding: This work was supported by the National Science Foundation [Contract SES-1024909] and the University of Massachusetts Amherst Faculty Research Grant/Healey Endowment [Grant P1FRG0000000109].

Supplemental Material: The online appendix is available at https://doi.org/10.1287/opre.2018.1827.

Keywords: public sector • nonprofit operations • public policy • foreclosure • resource allocation

1. Introduction

The United States has experienced various economic recessions over the last several decades with adverse effects in all sectors of the economy, including residential housing. The most recent recession occurred between 2007 and 2009, and a root cause of this major recession was a dramatic increase in mortgage foreclosures originating in home price depreciation that amplified the effects of increases in mortgage rates and the number of risky mortgages (Bernanke 2008). The impacts of this crisis on the U.S. economy were broad and profound: there were substantial decreases in housing values, home equity, home sales, and total housing starts (Joint Center for Housing Studies 2009). As a result, home foreclosures resulted in massive losses of consumer wealth: on average, U.S. households lost more than $2.2 trillion per year in home value between 2008 and 2012, with losses totaling around $3 trillion in 2011 (Christie 2009, Lubin 2011). Although the markets have seen improvements since 2012, the recovery has been slow given the profound impacts of the recession.

Policies to mitigate such losses are implemented on a continuous basis independent of the state of the economy. These policies include efforts to reduce the number of foreclosed homes in neighborhoods, which, in turn, will result in the appreciation of home prices and stabilization of the housing market. To this end, many initiatives exist at federal, regional, or local levels. Key actors in these efforts are nonprofit community development corporations (CDCs), which acquire and rehabilitate foreclosed properties using available resources. CDCs exist in nearly every major urban area of the United States, with approximately 5,000 CDCs spread throughout all 50 states (Community Wealth 2012). Given the large number of CDCs operating in different parts of the United States, the decision problems that they face in their efforts to respond to the foreclosure problem have implications for the overall economy and public good.

In this paper, we focus on CDC-led initiatives in urban neighborhoods involving acquisition and redevelopment of foreclosed housing units for resale or rental to support neighborhood stabilization and revitalization. Our fundamental research question is the following: given the limitations in the amount of accessible funds and the uncertainty on the impacts of foreclosed property acquisitions, what are socially optimal policies that a CDC should implement while considering foreclosed properties for potential acquisition? Our approach to this problem involves a dynamic and stochastic resource allocation model where potentially accessible funds are allocated dynamically to maintain an optimal portfolio of acquired properties. We use data and other input obtained from a CDC operating in Boston, Massachusetts and identify analytical and numerical results that characterize optimal policies for a CDC.

1.1. Literature Review

1.1.1. Related Research on Nonprofit Foreclosed Housing Policy. Existing research on nonprofit foreclosed housing policy includes both descriptive and prescriptive studies. Although our research fits mostly into the latter category, several complementary descriptive analyses exist that typically investigate the impacts of foreclosures on neighborhoods and communities (Immergluck and Smith 2006; Harding et al. 2009; Leonard and Murdoch 2009a, b; Swanstrom et al. 2009; Campbell et al. 2011). The challenges faced by CDCs in their acquisition activities are specified by Bratt (2009) and NeighborWorks (2009), where the authors note that foreclosed property acquisition is different from traditional community development and highlight the need for CDCs to implement efficient and effective policies in foreclosed housing acquisition. Our study is aimed at providing such policy insights to CDCs or any other organization involved in similar activities.

Prescriptive models on foreclosed housing acquisition policies are very few and involve several restrictive assumptions. The first such study is by Johnson et al. (2010), in which the authors describe the formulation and solution of a multiobjective integer program to guide foreclosed housing acquisition decisions by CDCs. However, the presented analysis is both static and deterministic, because it assumes that the set of available properties and their characteristics are known a priori. A similar framework with stochastic property costs and returns, but with special cases involving probability distributions that have dominance relationships, is used by Bayram et al. (2011) in deriving basic analytical results for a two-period foreclosed housing acquisition problem. Moreover, similar to Johnson et al. (2010), the decision problem studied by Bayram et al. (2011) considers acquisition decisions over a known set of available properties. Other than these studies, there are no known analyses that address the management of foreclosed housing acquisition portfolios. Adding to this limited literature, in this paper, we develop policy insights under a comprehensive stochastic dynamic decision structure.

1.1.2. Related Research on Dynamic and Stochastic Resource Allocation Problems. One stream of research related to our problem involves dynamic resource allocation decisions on a set of given investment options, attributes of which evolve stochastically over time (Loch and Kavadias 2002, Bertsimas and Popescu 2003, Calafiore 2008, Mild and Salo 2009). However, considering that investment options in the foreclosed housing acquisition problem are not known a priori, the results of this stream of research on stochastic dynamic resource allocation models are not directly applicable to our framework. The foreclosed housing acquisition decision process involves knapsack-type offer/no offer decisions on properties that become available over time. Hence, a more relevant stream of research is the literature on dynamic stochastic knapsack problems. These problems have been formally defined by Kleywegt and Papastavrou (1998): the authors study optimal acceptance policies for equal-sized items that arrive randomly over time with stochastic return structures. This framework is then extended by Kleywegt and Papastavrou (2001) to include items with random sizes. These two papers form the basis for some other studies in the literature that involve models and solution approaches proposed for various applications with a dynamic stochastic knapsack structure. One such application is the study by Kilic et al. (2010), in which the authors study a model for raw material allocation in food production and propose a heuristic algorithm to derive an optimal policy. Similarly, Dizdar et al. (2011) derive simple policies under some limiting assumptions for revenue maximization-based general resource allocation problems. In another general analysis, Lu (2005) proposes a computational procedure to calculate the optimal policy for infinite horizon dynamic stochastic knapsack problems. In addition, Lin et al. (2008) apply stochastic knapsack-type models in revenue management and dynamic pricing, whereas Nikolaev and Jacobson (2010) study resource allocation to a random number of jobs and present optimal policies for the sequential stochastic assignment problem. All of these studies involve variations of the dynamic stochastic knapsack problem with different characteristics or algorithmic implementations in various applications. Our study adds to this stream of research by considering a distinct application in the nonprofit sector and extends this general framework by considering additional decisions, such as the selection of an overbid rate when making offers to selected foreclosed properties. The inclusion of the overbid rate decision in the problem framework results in an additional layer of optimization over the classical dynamic stochastic knapsack problems. Hence, both the acceptance policies and qualitative results provide unique aspects for similarly structured problems.

1.2. Contributions

In this paper, we design, implement, and evaluate decision models that yield foreclosed housing acquisition policies for community-based organizations. A comparative analysis of our proposed policies with historical acquisition data suggests a potential improvement of around 20% in expected total impacts of the acquisitions on nearby property values. Considering an annual fund availability of $4 million for investment by a CDC, this corresponds to an estimated additional value of around $280,000 for the society. Moreover, to the best of our knowledge, this paper is the first comprehensive study in the public or private sector that specifically addresses optimal decision making in housing acquisitions under uncertainty. Decision modeling for housing and community development is situated in a long-standing literature of public sector operations management (Larson and Odoni 1981, Kaplan 1984, Pollock et al. 1994, Johnson 2011), and our research contributes to this area by focusing on a new public policy-related problem with important social and economic implications. Through its distinct application area involving nonprofit management and extensions over the general stochastic dynamic knapsack framework, our analysis also complements the existing literature on stochastic dynamic resource allocation problems.

2. Dynamic Foreclosed Housing Acquisition Problem

We consider a CDC that faces decisions on potential acquisitions of foreclosed properties that become available over time in their service area. By being available, we refer to the case where a property is placed on market for potential sale by a bank or other mortgage holder and the property is potentially approvable for acquisition by a funding source. When a foreclosed property is put on sale, owner occupants and public entities, such as CDCs, get an early opportunity to bid on that property. Moreover, regulations for such sales require that the asking price for the property is at least 10% below the appraised property value and that the final sales price is not more than 99% of the appraised value. These lower rates reflect the cost savings of a quick sale for financial institutions, whereas for nonprofit agencies, they provide an opportunity to acquire properties without competing with private investors—because it is required that the seller accepts the highest offer made by an eligible organization during the first 15 days of property availability (Stevens 2011). Hence, this allows for a higher likelihood of a successful offer for CDCs due to the relatively fewer competitors in the process (Hersh 2012).

We assume that the availability of foreclosed properties for potential acquisition follows a Poisson process with rate λ as suggested by analysis of historical data.[1] The availability rate λ is typically related to the conditions of the economy and the housing market, and it is assumed to be stationary (i.e., we do not model a dynamic environment where the state of the economy fluctuates, because such fluctuations have long-term dynamics). However, we analyze later in the paper how optimal policies change for different availability rate levels. At the beginning of a planning period, a CDC has an estimate of the total funds that it can access through various sources during that period. This funding level, which we denote by B, represents the total amount of credit or other funds that the CDC can assume to be available for foreclosed property acquisition. Suppose that $T \in (0, \infty]$ denotes an expiration time for the available funds, after which any unused funds will have no value. This time limitation typically depends on the funding source. For example, certain government funds may have deadlines that they need to be used by, whereas other resources, such as donations, do not have any such stipulations.

When a foreclosed property is placed on the market by the lender, it has an associated *asking price* usually based on the price opinion of a broker with experience in the area. Without loss of generality, we assume that the asking price is a lower bound on a foreclosed property's market value as well as on the amount required for a successful offer. This is quite typical in the regular operation of the real estate market, because banks or other lenders would often have lower asking prices on foreclosed properties due to their desire to sell these properties quickly. Hence, such properties would typically sell at or above the asking price, and offers would typically involve overbids. We let $C \in [\underline{C}, \bar{C}]$ denote the asking price for a foreclosed property and assume that it is defined by a probability density $f(c)$.

Estimating the *social return* from the acquisition and redevelopment of a foreclosed property is clearly difficult. Johnson et al. (2013) highlight this challenge and develop a measure validated by some CDCs, which is based on the impact of the acquisition of a foreclosed property on the appreciation of the value of nearby properties. More formally, the property value impact (PVI) measure is defined as the expected impact on proximate property values due to the acquisition and redevelopment of a given foreclosure. This measure is directly related to the geographical location of a property, and it can easily be calculated through a procedure described by Johnson et al. (2013). However, the social return from a property is also a function of the number of previously acquired properties in the service area by the CDC. This is because of the contagion effect described by Harding et al. (2009), where the cumulative effects due to multiple foreclosures are not equal to the sum of the effects that each individual foreclosure would have on nearby properties. This suggests that any return from a foreclosed property acquisition would be higher if there are fewer remaining foreclosures nearby. Hence, we denote the social return from a property as $R(s) \in [\underline{R}_s, \bar{R}_s]$, where $s \in [0, S]$ is the number of previous acquisitions by the CDC in the service area. In this representation, S defines an upper bound on the total number of acquisitions, which may be set by the budget constraint or some other criteria. Johnson et al. (2013) model the contagion effect by using a discounting factor over the cumulative returns from multiple acquisitions, where the discounting factor varies based on the number of acquisitions within limited radii. Similarly, we use the average discounting values over different distances defined by Harding et al. (2009) and develop the following relationship: $R(s) = R \cdot \min\{1, -0.00055s^2 + 0.03s + 0.63\} = R \cdot M(s)$, where $R \in [\underline{R}, \bar{R}]$ corresponds to the individual return from the property in the absence of any contagion effects. The individual returns R are stochastic, and they are characterized by a probability density $f(r)$ or $f(r, c)$ if social returns and the asking prices are correlated.

Observing a foreclosed property entering the market at time $t \in [0, T]$, a CDC with available funds of $b \leq B$ decides whether it should consult with the funding source and make an offer on the property and if so, how much to offer. We model the offer amount decision for a property through an overbid rate parameter $\delta(b, s, c, r, t) \in [0, \bar{\delta}_c]$, which corresponds to the percentage difference between the offer amount and the asking price c. This generic notation implies that the overbid rate can potentially vary over the planning horizon for different funding and previous acquisition levels, and the limits on the overbid rate can differ based on the asking price of a given property. Our modeling considers and compares two practical cases based on different assumptions for the selection of δ. In the first case, for tractability and simplification purposes in the decision process, we assume partial-state dependence for δ, such that the overbid rate is price dependent only. The second case assumes a fully state-dependent structure, where the overbid levels are both price and PVI dependent. Moreover, we denote δ as a percentage (e.g., $\delta = 0.05$ implies that the offer amount is 5% over the asking price). The probability of success for an offer (i.e., the probability of winning a bid) is an increasing function of the overbid rate, and it is denoted $p(\delta)$. As part of our analysis, we assume that $p(\delta)$ can be any general increasing bounded function.

When an offer is made on a foreclosed property, an overhead cost corresponding to the time and other expenses required to make the offer is incurred by the CDC. This cost is typically a certain percentage G of the offer amount, and it can be defined as $g(c, \delta) = (c + c\delta)G$, where $c + c\delta$ is the offer amount. If the offer is not accepted by the seller, then this cost is sunk. However, if the offer is successful, the cost is still incurred and deducted from the PVI-based returns from the acquisition, because they represent funds that can be used by the CDC for social value creation. We note that these costs normally come out of the operating budget of the CDC. Hence, from a modeling perspective, a direct inclusion of them would make it necessary to have the operating budget as part of the model, which is outside the scope of this paper. Furthermore, deducting these costs from the acquisition budget b would not be practically as accurate and also, would result in tractability issues in the modeling framework. In addition, we show through the numerical analyses later in the paper that these indirect costs are not very influential in the decision-making process due to the disparity between the PVI levels and the typical overhead costs. Nonetheless, we include them in our model for completeness purposes, because such inconsequentiality would not be the case if the overhead costs were to be larger than their currently observed levels in practice.

Given this decision framework, the objective for the CDC is to determine a policy involving offer/no offer decisions and overbid rates that maximizes the expected total social return accumulated over a given planning horizon.

2.1. Model Formulation Under the Price-dependent Overbid Policy

In this section, we build a dynamic programming model to capture the timeline of events and decisions as described in the previous paragraphs under the simplifying assumption that the CDC's decision process will consist of two stages. In the first stage, the CDC observes the asking price $c \leq b$ for the property and commits to an overbid rate $\delta \leq \frac{b}{c} - 1$ without calculating the expected value to be generated from the property. In the second stage, an estimate of the return r is assumed to be available, and the CDC decides whether to proceed with the offer using the predetermined overbid rate δ. This two-stage problem repeats itself until the budget is fully spent or until it expires, whichever happens first. In what follows, we build our dynamic programming model by starting with the second decision stage.

Given the overbid rate δ in the *second stage*, the CDC's decision is whether to make an offer on a foreclosed property with an asking price of c and an expected social return of r. This decision depends on whether the CDC would generate a higher social value by making an offer or not making an offer on the property. Let $J_t(\delta|b,s,c,r)$ be the expected total social value if the CDC submits a bid with an overbid rate δ. This expected social value, or the PVI level, can be expressed as follows:

$$J_t(\delta|b,s,c,r)=\Big[r\cdot M(s)-g(c,\delta)+V_t(b-c-c\delta,s+1)\Big]p(\delta) + \Big[V_t(b,s)-g(c,\delta)\Big](1-p(\delta)), \quad (1)$$

where $V_t(b,s)$ is the optimal expected PVI to go to be generated through an available fund level b when s foreclosed properties have already been acquired in the service area. The first term in the brackets corresponds to the value when the offer is accepted, and the second term in the brackets corresponds to the value when it is not. Note that, if the CDC decides not to submit an offer, then its expected return will simply be equal to $V_t(b,s)$. Hence, the CDC's second-stage decision problem can be expressed as a discrete choice between $J_t(\delta|b,s,c,r)$ and $V_t(b,s)$. To summarize, the CDC would submit an offer if and only if

$$\overbrace{\begin{aligned}&[r\cdot M(s)-g(c,\delta)+V_t(b-c-c\delta,s+1)]p(\delta)\\ &+[V_t(b,s)-g(c,\delta)][1-p(\delta)]\end{aligned}}^{J_t(\delta|b,s,c,r)} \geq V_t(b,s). \quad (2)$$

Through some algebraic manipulation, condition (2) can be expressed as a threshold policy based on the individual PVI of the available property. More specifically, we expand the products on the left-hand side of the inequality and then leave the PVI term r alone on the left-hand side of the equation by moving all other terms to the right-hand side. This results in the conclusion that the CDC should make an offer on a foreclosed property available for acquisition if accessible funds are sufficient and the individual PVI value of the property is greater than a threshold: that is, if

$$r \geq x_t(\delta|b,s,c)=\frac{1}{M(s)}\left[V_t(b,s)-V_t(b-c-c\delta, s+1)+\frac{g(c,\delta)}{p(\delta)}\right]. \quad (3)$$

Although the threshold value $x_t(\delta|b,s,c)$ is dependent on δ, b, s, c, and t, for the sake of notational simplicity throughout the paper, we denote the PVI threshold by $x_t(\delta)$.

In the *first stage*, the CDC observes the asking price c and commits to how much to overbid based on the distribution of the return r. Let $W_t(\delta|b,s,c)$ denote the expected PVI, where the expectation is taken with respect to $r\in[\underline{R},\overline{R}]$:

$$\begin{aligned}W_t(\delta|b,s,c)&=\int_{\underline{R}}^{\overline{R}}\max\Big\{V_t(b,s),J_t(\delta|b,s,c,r)\Big\}f(r,c)\mathrm{d}r\\ &=\int_{\underline{R}}^{x_t(\delta|b,s,c)}V_t(b,s)f(r,c)dr\\ &\quad+\int_{x_t(\delta|b,s,c)}^{\overline{R}}J_t(\delta|b,s,c,r)f(r,c)\mathrm{d}r\\ &=V_t(b,s)P(r\leq x_t(\delta|b,s,c))\\ &\quad+\int_{x_t(\delta|b,s,c)}^{\overline{R}}J_t(\delta|b,s,c,r)f(r,c)\mathrm{d}r.\end{aligned} \quad (4)$$

Given $W_t(\delta|b,s,c)$, the CDC's first-stage problem can be cast into the following optimization:

$$W_t^*(b,s,c)=\max_{\delta:0\leq\delta\leq\frac{b}{c}-1}W_t(\delta|b,s,c). \quad (5)$$

Finally, the value function $V_t(b,s)$ can be expressed by conditioning on the possible events that can happen over a small interval of time Δt. Considering that foreclosed properties arrive at the rate of λ, the CDC receives a new arrival over the next time interval Δt with probability approximately $\lambda\Delta t$ and no arrival with probability approximately $1-\lambda\Delta t$. Given that a new property arrives at $t+\Delta t$ and an overbid rate δ is set, the CDC's offer succeeds with probability $p(\delta)$ and generates expected value $E_c\big[W_{t+\Delta t}^*(b,s,c)\big]$, where c is distributed between $\underline{C}$ and $\overline{C}$. Finally, assuming a discount rate of $\alpha\Delta t$ and using the principle of optimality, we can write $V_t(b,s)$ in terms of discounted expectation of value function over possible states in $t+\Delta t$:

$$V_t(b,s)=\frac{1}{1+\alpha\Delta t}\Big[\lambda\Delta t E_c\big[W_{t+\Delta t}^*(b,s,c)\big] + (1-\lambda\Delta t)V_{t+\Delta t}(b,s)\Big].$$

We can then rearrange the terms to obtain

$$V_t(b,s)-V_{t+\Delta t}(b,s)=\lambda\Delta t\Big[E_c\big[W_{t+\Delta t}^*(b,s,c)\big] - V_{t+\Delta t}(b,s)\Big]-\alpha\Delta t V_t(b,s).$$

Dividing both sides by Δt and taking the limit as $\Delta t\to 0$ lead to the following Hamilton–Jacobi–Bellman (HJB) equation for $V_t(b,s)$:

$$\frac{\partial V_t(b,s)}{\partial t}=\lambda\Big[E_c\big[W_t^*(b,s,c)\big]-V_t(b,s)\Big]-\alpha V_t(b,s). \quad (6)$$

To facilitate the analysis, we assume that the funds accessible by CDCs for foreclosed property acquisitions do not have usage deadlines. This implies that the value function is stationary over time: that is, $\frac{\partial V_t(b,s)}{\partial t} = 0$. Under the case without fund expiration, Equation (6) transforms into the following stationary HJB equation:

$$0 = \lambda\Big[E_c[W^*(b,s,c)] - V(b,s)\Big] - \alpha V(b,s). \qquad (7)$$

The above derivation of stationary HJB in Equation (7) is not rigorous, because we have not justified the interchanging of expectation and limit operations. However, these conditions can be justified formally using the results of Kleywegt and Papastavrou (2001), where general dynamic stochastic knapsack problems are studied. Thus, as we show later in Theorem 1, a solution to Equation (7) is indeed the optimal value function $V^*(b,s)$, and the overbid rate δ^* that achieves the maximization in Equation (5) is the optimal overbid rate for CDC.

Before we structurally analyze this two-stage model, we describe the comprehensive formulation where it is assumed that the offer submission decision and the determination of the overbid rate are made simultaneously under a price- and PVI-dependent structure.

2.2. Model Formulation Under the Price- and PVI-dependent Overbid Policy

In this extended formulation, the timeline of events is assumed to be similar to the case above, except that the CDC makes the overbid decision after the realization of asking price c and calculation of PVI r. Let $V_t^F(b,s)$ denote the value function under the price- and PVI-dependent overbid policy. We also start modeling this case from the last step. Noting that $J_t(\delta|b,s,c,r)$ is the expected social value if the CDC makes an offer using an overbid rate δ, the optimal overbid policy can be determined through the following optimization problem:

$$J_t^*(b,s,c,r) = \max_{\delta: 0\le\delta\le\frac{b}{c}-1} J_t(\delta|b,s,c,r). \qquad (8)$$

Hence, the CDC's decision on whether to submit a bid can be expressed as a discrete choice between $J_t^*(b,s,c,r)$ and $V_t^F(b,s)$. In other words, the CDC would submit an offer if and only if $J_t^*(b,s,c,r) \ge V_t^F(b,s)$. Finally, similar to Section 2.1, the value function $V_t^F(b,s)$ can be expressed by conditioning on the possible events that can happen over a small interval of time Δt. To avoid repetition, we provide the (nonstationary) HJB equation for the value function $V_t(b,s)$ as follows:

$$\frac{\partial V_t^F(b,s)}{\partial t} = \lambda\Big[E_{c,r}\left[\max\{V_t^F(b,s), J_t^*(b,s,c,r)\}\right] - V_t^F(b,s)\Big] - \alpha V_t^F(b,s), \qquad (9)$$

where $E_{c,r}$ denotes the expectation taken with respect to c and r. Similarly, under the case without fund expiration, the above equation transforms into a stationary HJB equation as follows:

$$0 = \lambda\Big[E_{c,r}\left[\max\{V^F(b,s), J^*(b,s,c,r)\}\right] - V^F(b,s)\Big] - \alpha V^F(b,s). \qquad (10)$$

3. Structural Results

In this section, considering the case without fund expiration, we first explore the structural properties of the price-dependent overbid policy and use the insights from this exploration to analyze the price- and PVI-dependent case. The implications of imposing expiring deadlines on fund availabilities are discussed at the conclusion of the section.

3.1. The Price-dependent Overbid Policy

As mentioned in Section 2.1, the value function for the case without fund expiration satisfies the stationary HJB equation provided in (7). Our first result shows that a solution to Equation (7) is indeed the optimal value function $V(b,s)$, and the overbid rate δ^* that achieves the maximization in Equation (5) forms the optimal overbid rate for CDC.

Theorem 1. *If accessible funds for foreclosed housing acquisition for a CDC do not expire, then the optimal expected total PVI for a given fund level b when there are s previous acquisitions in the service area, denoted as $V(b,s)$, is the solution of the following equation:*

$$V(b,s) = \frac{\lambda}{\alpha}\int_{\underline{C}}^{b} \max_{\delta\in[0,\bar{\delta}_c]}\left\{\int_{\underline{R}}^{\overline{R}}\Big([r\cdot M(s) + V(b-c-c\delta, s+1) - V(b,s)]p(\delta) - g(c,\delta)\Big)^+ f(r,c)\,\mathrm{d}r\right\}\mathrm{d}c, \qquad (11)$$

where $(x)^+ = \max(0,x)$.

Proof. All proofs are included in the online appendix. □

The optimality Equation (11) has some intuitive explanation that might help in understanding the rationale behind its form. Namely, the value function $V(b,s)$ can be thought of as being equal to the product of the arrival rate of new foreclosures and the discounted expectation of incremental value from a newly available foreclosed property—in an analogous way to Little's Law in queueing theory (Little 1961). Here, the incremental value can be expressed as the integrand of the inner integral in Equation (11). The integrations thus correspond to the expected incremental value of a newly available foreclosure when the amount of accessible funds is b and there are s previous acquisitions in the service area. Although the value function $V(b,s)$ in Theorem 1 is neither convex nor concave in b under a general setting, we are able to show its convexity under certain conditions involving the following.

Assumption 1. *The probability of a successful bid $p(\delta)$ and the cost $g(c,\delta)$ for making an offer on a property are linear in δ.*

We use the above assumption along with the fact that convexity is closed under summation with nonnegative weights, pointwise maximization, and expectation. Although this assumption may sound technical and restrictive, indeed the real data collected from the CDC largely validate the linearity of the probability and cost functions. With these assumptions, we now show that the value function is convex and increasing in b.

Theorem 2. *Suppose that Assumption* 1 *holds. Then the value function $V(b,s)$ is convex and increasing in b.*

By evaluating the integral in Equation (11), the set of corresponding recursive relationships can be numerically solved to determine the optimal expected total PVI for each $b \in [0,B]$ and $s \in [0,S]$ and thus, the optimal overbid rates δ^* as well as the PVI thresholds $x(\delta^*)$. For the optimal overbid rate, it is possible to numerically evaluate the recursion in (11) by considering a discrete set of overbid options and then, selecting the rate that results in the maximum expected incremental value. However, it is also possible to characterize the optimal overbid rate analytically under certain conditions as follows.

Theorem 3. (1) *Under general distribution for the returns, if the condition $p'(\bar{\delta}_c)\bar{K} - L \le 0$ holds, then the CDC's offer should be at the asking price. However, if the condition $\bar{L} \le 0$ holds, then the CDC's offer should be at an overbid level of $\bar{\delta}_c$, where $\bar{K}$, L, and $\bar{L}$ are as defined in the online appendix.*

(2) *If the returns in the foreclosed housing acquisition problem are uniformly distributed and $p'(\bar{\delta}_c)(\bar{R} - \underline{R}) - 2\underline{L} \le 0$ holds, then the CDC's offer should be at the asking price.*

Through additional analysis, we also note some qualitative characteristics related to the optimal acquisition policy under the no fund expiration case. As part of this analysis, we introduce two measures of practical relevance, which we refer to as the "critical fund level" and "critical time." The critical fund level is defined as the specific funding level such that the optimal policy for funds larger than that level is to make offers to all available properties. Similarly, the critical time is the time period such that the optimal policy after that time period is to make offers to all available properties. For the no fund expiration case, the optimal thresholds are constant over time, and therefore, the critical time is either zero or ∞. Hence, this measure becomes more relevant when the funds expire at a certain time, which we discuss in the online appendix.

Given these additional definitions, we summarize some important characteristics for the optimal foreclosed housing acquisition policy as follows.

Theorem 4. *The following conditions hold for the foreclosed housing acquisition problem with no fund expiration.*

1. *The larger the amount of accessible funds, the higher the total expected property value impact to be realized from foreclosed property acquisitions.*
2. *The higher the availability rate of foreclosed properties, the lower the overbid rate should be.*
3. *Let $p(\delta) = p_0 + p_1\delta$. Then, if $p_0 \ge p_1$ holds, then the higher the PVI threshold used by the CDC, the higher the overbid rate should be.*
4. *Let $p(\delta) = p_0 + p_1\delta$. Then, if $p_0 \ge p_1$ holds, the higher the overbid rate used, the higher the critical fund level.*

The first item is an intuitive conclusion that the higher the amount of accessible funds, the more value they have in terms of the total PVI that can be achieved from acquired properties. Similarly, item 2 in Theorem 4 is also somewhat intuitive, because it indicates that a CDC should lower the overbid rates if acquisition options arrive at a higher rate. The property described in the third item in Theorem 4 implies that the relationship between PVI thresholds and overbid rates can vary over the range of overbid rates, but a statement can be made on the relationship when the condition in the item is satisfied. In that case, a CDC should use higher overbid rates when an offer is made on a property with a larger PVI value. Finally, we note similarly that, under the stated condition, if the CDC uses high overbid rates and high thresholds, then the critical fund level will get higher as the overbid rate increases. In other words, usage of higher overbid rates would result in the CDC stopping its selectivity earlier.

3.2. The Price- and PVI-dependent Overbid Policy

We now consider the case in which the CDC decides the overbid rate based on both asking price and PVI. Similar to Theorem 1, we can also show that a solution to Equation (10) is indeed the optimal value function $V^F(b,s)$ under the price- and PVI-dependent overbid policy.

Theorem 5. *Let $V^F(b,s)$ denote the optimal expected total PVI in the service area if the CDC decides on the optimal overbid rate after observing both the asking price c and return r. Then $V^F(b,s)$ solves the following recursion equation:*

$$V^F(b,s) = \frac{\lambda}{\alpha}\int_{\underline{C}}^{b}\int_{\underline{R}}^{\overline{R}} \max_{\delta\in[0,\bar{\delta}_c]}\Big(\big[r\cdot M(s) + V^F(b-c-c\delta, s+1) - V^F(b,s)\big]p(\delta) - g(c,\delta)\Big)^+ f(r,c)\,\mathrm{d}r\,\mathrm{d}c, \quad (12)$$

where $(x)^+ = \max(0,x)$.

As in the previous case, our numerical studies show that the value function is neither convex nor concave under a general setting. Hence, we are motivated to use a similar approach and show that the value function is convex in b under linear probability and cost functions. To avoid repetition, we delegate the details of the convexity analysis of the value function $V^F(b, s)$ to the online appendix. Finally, the convexity of the value function can be used to determine comparative statistics of the optimal overbid rate with respect to b and r as follows.

Theorem 6. *Suppose that Assumption* 1 *holds. Then the optimal overbid rate is increasing in r and decreasing in b.*

3.3. Policies Under Fund Expiration

Another case in the foreclosed housing acquisition process of a CDC is when the CDC faces deadlines for using non-loan-type acquisition funds. This is typically the case when the providers of the funds stipulate that they are to be used within a given timeframe. For example, the funds that are made available to CDCs by the federal government as part of the Neighborhood Stabilization Program require that these funds are used by a certain deadline (Stable Communities 2012). Utilizing a similar structure to the infinite horizon case, we can characterize the optimal foreclosed housing acquisition policies for CDCs when they have such time-based limitations for using the accessible funds. We broadly discuss here the implications of fund expiration on the problem. Imposing fund expiration introduces nonstationarity into both of our models. This, in turn, implies that one needs to solve a partial differential equation to characterize the optimal value function at each time period. In addition to the computational challenges, characterizing bounds on the nonstationary value function becomes analytically quite challenging, because our numerical investigations yield that the resulting value function becomes neither convex nor concave under all cases. We present specific observations from numerical investigations under fund expiration in the online appendix.

4. Real-Life Implementation and Implications for Housing Acquisition

4.1. Description of Data

Data-based implementation and analysis of the models were performed in close coordination with a CDC located in the city of Boston, Massachusetts, and the CDC was also involved in the model-building phase of the study. Although there is no comprehensive quantitative information that compares the operating framework for this CDC with that of other CDCs, there is anecdotal evidence that the CDC studied can be representative of typical CDC operations in other major cities (NeighborWorks 2009, Gass 2011). This is also supported by the fact that PVI distributions calculated for a different CDC in an entirely different area of Boston, Massachusetts, reflect similar characteristics as those obtained for the service area studied in this paper (Johnson et al. 2013).

To characterize the actual decision problem parameters as well as the uncertainty in property costs and returns, historical foreclosed property availabilities and acquisitions over an eight-month period were used. Based on these data, an average of approximately 10 properties was observed to be entering the market each month following a Poisson distribution, which corresponds to an average availability rate of 2.5 properties per week. A histogram showing the distribution of the time between property availabilities over an eight-month period is included in Figure 1(a). In addition, the results of the statistical goodness of fit tests for all estimated distributions are included in the online appendix. The data were also checked for seasonality effects through a time series and autocorrelation analysis, which suggested no seasonality-based pattern in the market entry rates of foreclosed properties. Similarly, the asking price data were also studied in detail to identify the best-fitting distribution. In Figure 1(b), we show a histogram of asking prices (i.e., C for the properties considered in the data set), which was fitted with a triangular distribution with parameters of \$120,000, \$250,000, and \$380,000.

The distribution of the PVI values (i.e., R for the same data set) as calculated through the methodology described in Johnson et al. (2013) is shown in Figure 2(a). The best-fitting distribution for the PVI values was found to be a uniform distribution between \$60,000 and \$120,000. Moreover, the PVI values of the properties are found to be independent of their asking prices based on statistical testing, which are also reflected in the scatter plot and regression line in Figure 2(b). Specifically, a Pearson correlation test was performed where the correlation coefficient was calculated as 0.002 with $p = 0.992$, indicating strong statistical consistency with the hypothesis that PVI and asking prices are uncorrelated. Hence, we assume independent distributions for asking prices and the PVI measures.

The overhead costs for the CDC for each offer that it makes on a property are calculated to be around 1.5% of the amount offered on the property. As discussed in the model description, the probability of acquisition after an offer is made on a property depends on the overbid rate used for a given asking price. Using data from 24 previous acquisitions and also based on consultations with the CDC staff, we define this probability as $p(\delta) = 0.22\delta + 0.37$. We note that the 24 data points were the only data available for estimating $p(\delta)$, and they may not be large enough for a perfect estimation of this function. However, the structure of

Figure 1. Distributions of Asking Price and Market Entry Times of Foreclosed Properties in the CDC's Service Area

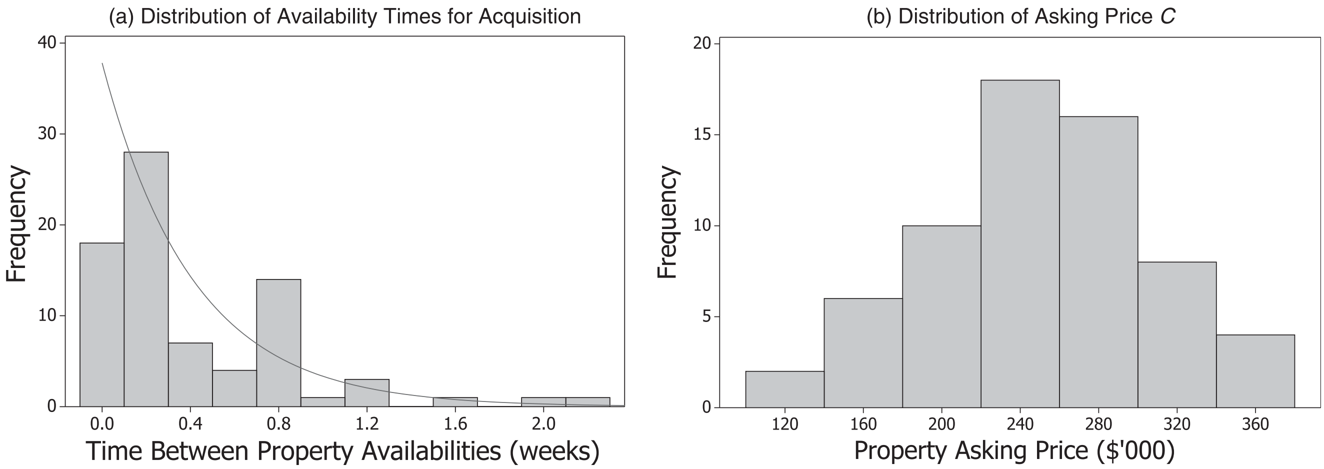

this probability of success function is consistent with the discussions and function descriptions by Holt and Sherman (2000) and Caskey and Aobdia (2012) for similar settings.

4.2. Implementation Results and Policy Implications

4.2.1. Expected Value Under Price-dependent vs. Price- and PVI-dependent Overbid Policies. We use the data described above to compare the expected PVIs under the two policies and assess the value of implementing the more complex price- and PVI-dependent overbid policy. In Figure 3(a), we show this value as a function of the funds accessible by the CDC for the case with no fund expiration deadlines and no previous acquisitions (i.e., for $s = 0$). It can be observed that the difference between the expected total PVIs under the two policies is minimal, with the percentage difference being less than 0.7% for all cases. The differences are calculated in percentages as $\frac{V^F(b,s)-V(b,s)}{V^F(b,s)} \times 100\%$ for $s = 0$, where $V(b,s)$ and $V^F(b,s)$ correspond to the value functions under price-dependent and price- and PVI-dependent overbidding policies, respectively. Moreover, for larger budgets, this difference is even smaller (i.e., around 0.1%). The sigmoidal behavior at lower fund levels is likely due to the integer knapsack-type effects in the problem framework, where the left-over funds, which are not enough to purchase even the lowest-cost property, form a larger portion of the available funds. The relative impact of these on the overall value can be larger when the budget level is low. For the typical funding levels shown in Figure 3(a), the PVI differences correspond to values around a few thousand dollars in absolute terms. Hence, it can be concluded that, for typical practical settings, the price-dependent overbid policy performs very well and that a CDC does not necessarily need a PVI calculation to determine the overbid rate when an offer decision is made. This will allow for a simpler

Figure 2. Distribution of PVI Values and Their Dependency on Asking Prices for Foreclosed Properties in the CDC's Service Area

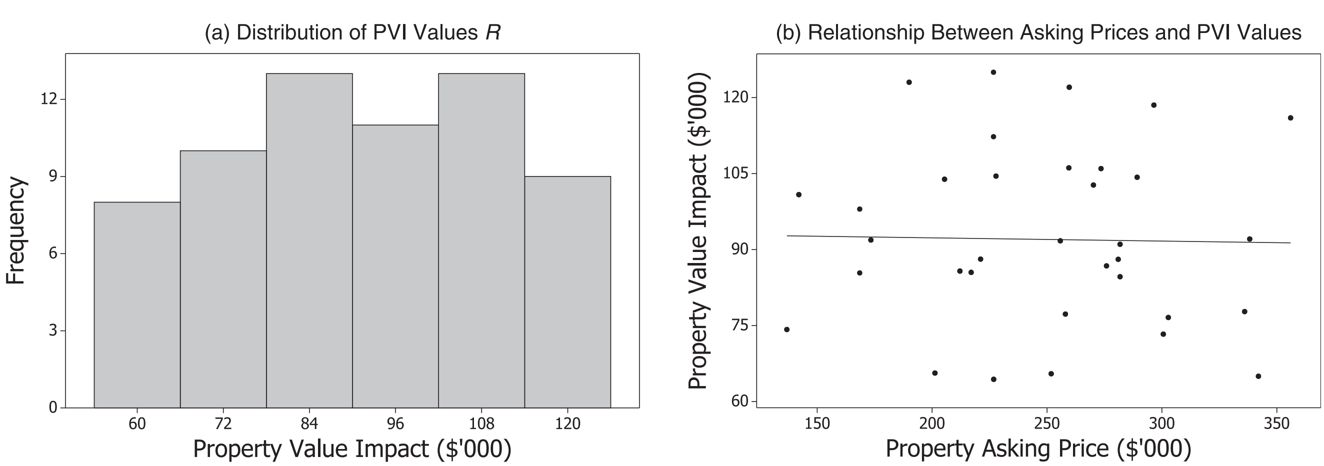

Figure 3. Expected Total PVI for Different Funding and Overbid Levels Under the Two Policies with No Fund Expiration and No Previous Acquisitions

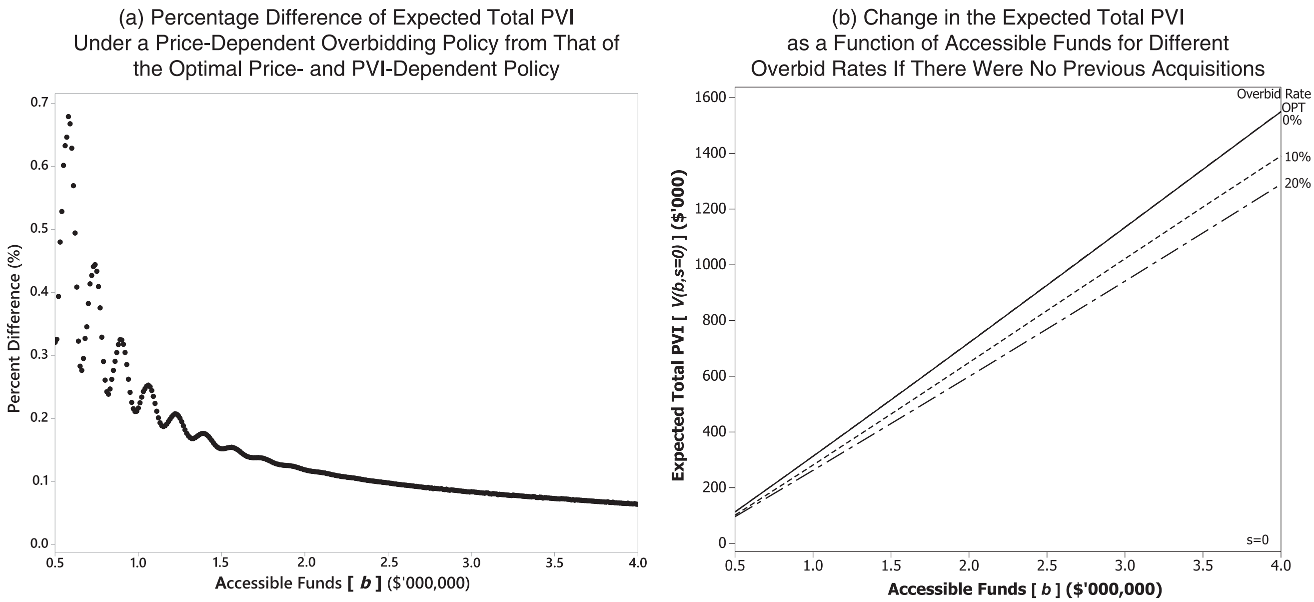

Notes. The vertical axis for panel (a) corresponds to the percentage difference between the value functions computed under the price-dependent and the price- and PVI-dependent bidding policies: that is, $\frac{V^F(b,s)-V(b,s)}{V(b,s)} \times 100\%$, where $s = 0$. The vertical axis for panel (b) corresponds to the value function $V(b,s)$ computed under different overbid rates for the price-dependent policy. $V(b,s)$ and $V^F(b,s)$ are calculated as defined by Equations (11) and (12), respectively, and correspond to the case with no fund expiration.

overbidding strategy to be used in practical decision-making situations. The reason for such similarity of results in our problem framework is likely due to the impact of PVI on the bidding threshold. More specifically, under the price- and PVI-dependent policy, the PVI value affects the bidding threshold $x(\delta|b,s,c,r)$ both directly (through r) and indirectly (through overbid rate δ), whereas the price-dependent policy captures only the direct effect (through r). Because the direct effect dominates the indirect effect, we observe that the decision to bid or not to bid is well captured by the price-dependent policy. Thus, for most practical settings, both policies are likely to produce similar results and value. The following analyses are based on the numerically more amenable price-dependent overbid policy.

4.2.2. Policy Implications for Return on Investment. In this section, we consider the expected social value (i.e., the total property value impact of funds available to a CDC) and how this value measures against the invested amounts.

When no previous acquisition-related effects are present (i.e., for $s = 0$), we find that, for a fund level of $4 million without an expiration deadline, the expected PVI to be realized through the foreclosed property acquisitions by the CDC is around $1.6 million. This value decreases at lower fund levels as shown in Figure 3(b) by the top curve, which corresponds to the optimal overbid rate usage under a price- and PVI-dependent policy. Although the optimal expected value is convex increasing at smaller budget levels, it becomes a mostly linear function for fund levels above $0.5 million. These results are also in line with Theorem 2. A relevant question involves the marginal value of the funds that the CDC can potentially access; for example, the marginal impact of a reduction in the acquisition funds can play a role when a CDC faces a decision on whether to use part of their available line of credit for purposes other than foreclosed property acquisition. We find that the marginal return increases as the fund level is increased up to around $0.5 million and then remains constant due to the linear structure. More specifically, we note the following conclusions. Each dollar of accessible funds is expected to result in about $0.35 of PVI returns at fund levels higher than $0.5 million and with no previous acquisitions. It is also observed that the value from each dollar invested increases synergistically as more properties are acquired. For a given fund level, the marginal PVI values increase up to around $0.45 per dollar if more acquisitions have been made in a service area. Hence, when multiple CDCs operate in a given service area with everything else being equal, a donor or funding agency should prefer to invest more on CDCs that have acquired more properties in the area. This finding can be used by successful CDCs to argue for additional donations, because they are likely to create more value to the society than a CDC that may not have been as active in property acquisitions in the corresponding service area. We note, however, that this conclusion does not take into account equity and fairness issues that could play a role when a competitive environment exists or when a new CDC is being established in a new area.

4.2.3. Policy Implications for PVI Thresholds in Offer/No Offer Decisions. A key finding of our models is that the CDC should work with the funding source and make an offer on all properties that satisfy a certain return threshold. These thresholds vary based on the asking price of the property, the fund level that the CDC has access to, and the number of previous acquisitions in the service area. In Figure 4, we show how the return thresholds, calculated as defined by Equation (3), behave as a function of each of these parameters under no fund expiration. We observe in Figure 4(a) that the thresholds decrease for each different asking price level as the amount of funds increases, and this decrease in thresholds is almost linear. For demonstration purposes, we select four different asking prices, namely c = \$120,000, \$200,000, \$280,000, \$360,000. We see that, at the lowest funding levels, the optimal policy is to bid on a \$120,000 property only if the return is greater than \$90,000. This criterion is gradually relaxed, and at the higher funding levels of around \$4 million, the optimal policy involves bidding on almost all properties with asking prices of \$200,000 or less, independent of their return values. As also mentioned in Section 3, a critical fund level can be identified such that, for any remaining fund level that is greater than the critical fund level, the CDC would make offers to all available properties within the given price range. However, a more selective policy can be used for fund levels that are less than the critical fund level based on the optimal PVI thresholds. For example, we observe through Figure 4(a) that the critical fund level for this numerical setup for an asking price of \$120,000 is around \$1.5 million. Overall, a CDC can use a spreadsheet-based tool or a series of charts similar to Figure 4(a) as a decision aid when choosing the properties that it will place offers for.

Notice that, because of contagion effects, the optimal PVI thresholds vary based on the previous foreclosed property acquisitions that a CDC might have made in a service area. This is reflected in the threshold plots shown in Figure 4(b), where it is assumed that the number of foreclosures in a given service area is bounded, and the CDC can be gradually decreasing the number of such properties through acquisitions. In such a case, as the number of acquisitions increases up to a certain level, the CDC should get more aggressive and become less selective in making offer decisions. This is because of the additional value that will be created through such synergistic acquisitions. However, selectivity should increase after a large number of previous acquisitions is reached, because the synergistic effects diminish in such cases. As an example, although the optimal policy is to never bid on the most expensive property when there are no previous acquisitions, this changes at higher acquisition states, and offers can be made on such costly properties if their expected property value impacts are reasonably large. As an additional issue, the modeling structure can also indirectly handle the case where the CDC simultaneously makes offer decisions on multiple properties. The discussion on this issue is included in the online appendix.

Figure 4. (Color online) Optimal PVI Thresholds for Different Funding Levels, Acquisition States, and Asking Prices

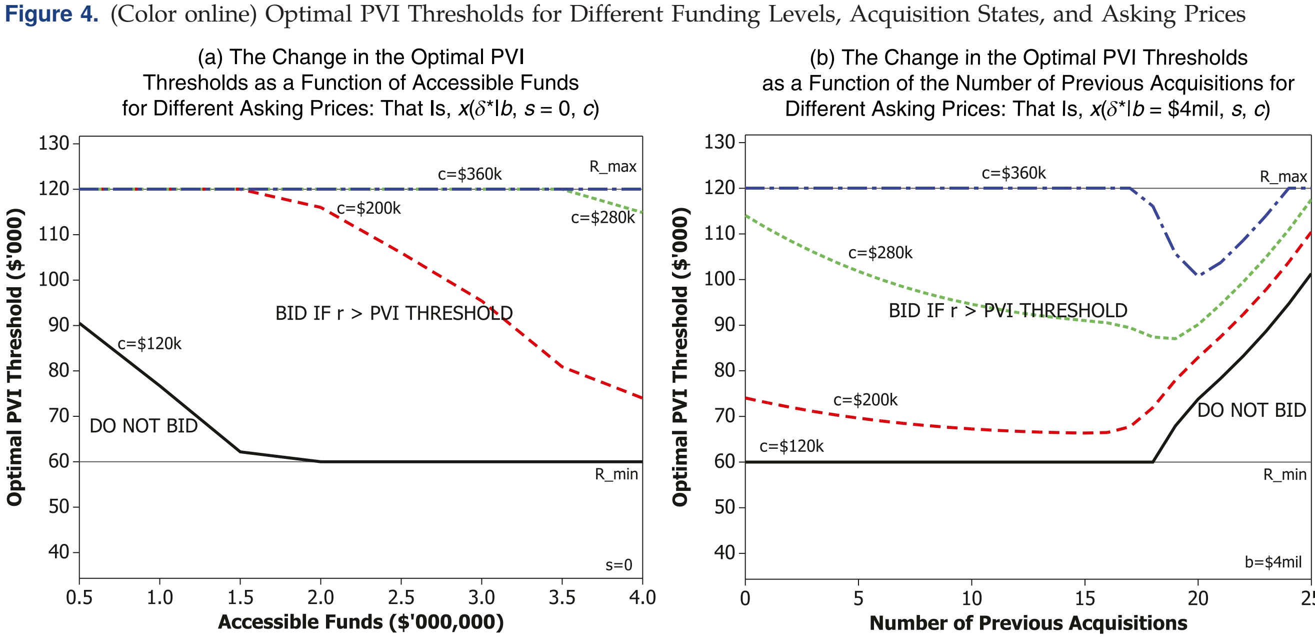

Note. For both figures, optimal PVI thresholds are calculated using price-dependent bidding under no fund expiration.

4.2.4. Policy Implications for Overbid Rates. The optimal overbid rate for the types of CDCs studied is mostly at the minimum possible level for all properties (i.e., the CDC should not offer more than the asking price for most foreclosed properties and should typically consider overbidding only when the available budget is low). This result implies a different policy than what is used in practice, and it is also in line with our findings in Theorem 6. We note in Theorem 3 some cases where the optimal overbid rate would be nonzero. In Figure 3(b), we illustrate the changes in the expected PVI as a function of the accessible funds for the optimal and some fixed overbid rates. The optimal rate in the figure corresponds to the price- and PVI-dependent policy results. These representations visually show the optimal overbid rate to be effectively zero, with the numerical results suggesting nonzero overbid rates at lower fund levels only. This observation is consistent with both our analytical findings and the expert views from the CDC staff. One potential reason for this conclusion is that the overhead costs are not that significant compared with other costs, such as the losses due to discounting, whereas at the same time, there exists a relatively large number of foreclosures in the CDC's service area. The result holds even when the overhead costs are increased to higher levels than 1.5% of the asking price. Therefore, even if the CDC's offer is not successful, it is likely that there will be other properties with comparable PVI values that will become available. Overall, the implication is that a CDC should make more frequent bids but do less overbidding to hedge against the risk associated with the uncertain bid outcome.

4.3. Value of Optimal Policies over Current Practice

In this section, we describe a comparative study aimed at showing the value of the optimal policies with respect to current practice and other simplistic heuristic procedures. The analysis is based on historical bidding and acquisition data obtained from the CDC with which we collaborated. These data consist of information on a set of properties that were considered for potential acquisition, with an offer/no offer decision on each property. In Table 1, we show part of this information, which includes the asking price and the estimated PVI value for each property, whether an offer was made on the property, and the overbid rate used. Current practices used by the CDC when considering an offer decision do not involve a systematic structure, and they are mostly based on qualitative and subjective opinions of the staff. An important observation is that the CDC is selective in making offers and almost always overbids on the properties for which an offer decision is reached. Contrary to this implementation, the optimal policies under both price-dependent and price- and PVI-dependent overbid rates suggest that an offer decision should be made on all candidate properties and with almost no overbidding. For our comparative study displayed in Table 2, we assume an accessible funding level of $1 million with no previous acquisitions and no fund expiration deadlines. We calculate the expected total discounted PVI that would have been achieved if the price-dependent and price- and PVI-dependent policies were implemented on the properties in the data set shown in Table 1. The calculations are based on Equations (11) and (12), respectively. These values

Table 1. Summary Information on a Set of Properties Considered for Potential Acquisition by the CDC

Property identification	Asking price (c)	Estimated PVI (r)	Offer made?	Overbid rate used (δ), %	Bid outcome
1	$164,500	$85,401	Yes	4.6	Lost
2	$170,360	$91,896	Yes	5.6	Acquired
3	$201,250	$75,628	Yes	4.5	Lost
4	$205,400	$103,884	Yes	9.6	Lost
5	$221,100	$88,122	No	—	—
6	$226,700	$112,285	No	—	—
7	$227,800	$104,507	No	—	—
8	$212,100	$85,767	No	—	—
9	$214,000	$85,513	Yes	5.2	Acquired
10	$190,000	$105,989	Yes	0	Acquired
11	$142,000	$100,850	Yes	7.2	Lost
12	$259,700	$122,034	No	—	—
13	$259,500	$106,138	No	—	—
14	$296,500	$108,519	Yes	3.2	Acquired
15	$338,200	$92,093	Yes	3.5	Lost
16	$226,800	$64,390	No	—	—
17	$125,000	$75,620	No	—	—
18	$150,000	$88,765	Yes	6.7	Lost
19	$184,900	$93,487	Yes	5.4	Lost

Table 2. Comparison of Different Policies Based on Historical Property Availability Data

Policy type (π)	Expected total PVI [$V^{\pi}(b = \$1\text{mil}, s = 0)$]	% Improvement over implemented [$\frac{V^{\pi}-V^{\text{imp}}}{V^{\text{imp}}} \times 100\%$]	Fund usage time (months)	% Difference from best [$\frac{V^{F}-V^{\pi}}{V^{F}} \times 100\%$]
Implemented (π = imp)	\$309,830	—	2.5	16.6
Price and PVI dependent ($\pi = F$)	\$371,515	19.9	1	—
Price dependent	\$370,403	19.6	1	0.3
Heuristic 1	\$339,606	9.6	5	8.6
Heuristic 2	\$327,162	5.6	4	11.9

Notes. Best policy corresponds to the price- and PVI-dependent overbid policy. The comparative analysis assumes a funding level of \$1 million with no previous acquisitions and no fund expiration deadlines.

are then compared with the total discounted PVI achieved through the actual offer and acquisition decisions, which we refer to as $V^{\text{imp}}(b,s)$. Similarly, we test the efficiency of two heuristic selection criteria by calculating the corresponding expected total PVI values if they were to be implemented for the data set. The two heuristics have a somewhat similar setup. In the first heuristic, an offer decision is reached if the standardized marginal return of a given property is greater than 0.5 and the offer price is equal to the asking price (i.e., no overbidding). In the second heuristic, offer/no offer decisions are made similarly, but the overbid rate used is assumed to be the average overbid rate in current practice (i.e., 5%). In Table 2, we show the differences in expected total discounted PVI values based on multiple simulated bid outcomes for the five cases as well as the average time that it takes to use the accessible funds of \$1 million for each case.

Based on the results of this comparative analysis, it can be observed that the optimal policies improve the total PVI of a funding level of \$1 million by about 20% compared with the PVI levels realized from the historically implemented offer/no offer decisions. Considering expected total PVI values and an annual fund availability of \$4 million for investment by a CDC, this corresponds to an estimated additional value of around \$280,000 for the society. The superiority of the optimal policies under both price-dependent and price- and PVI-dependent overbid rates is also visible over the heuristic policies that were based on marginal returns of properties. Similarly, it can be noted through a comparison of the two heuristics that overbidding still does not add value when it is used as part of a less aggressive strategy, such as the marginal return-based heuristics evaluated.

4.4. Robustness of Optimal Policies

An important issue deals with the robustness of the proposed results when the optimal policies are implemented based on the assumed distributions but the underlying real distributions deviate from these assumed distributions. To study this aspect of our policy findings, we have performed multiple tests aimed at measuring the effectiveness of the optimal policies with respect to the heuristic policies under different actual distributions of property asking prices and PVI values. Specifically, we considered five test cases and performed simulations for each case assuming a \$1 million accessible fund level. These cases and the corresponding test results are summarized in Table 3, where the baseline optimal policy in the table refers to the price-dependent overbid rate implementation with the originally assumed distributions. The first two cases correspond to the situation where the distributions that we assume are correct, but the means of these distributions are actually one standard deviation above or below the assumed means. The last three cases study the performance of the policies under different actual distributions of the asking prices and PVI values. The third case considers the same mean as the assumed means for different distributions, whereas the fourth and fifth cases address the different distribution situations with the means being above or below the assumed means, respectively.

Our findings in this robustness analysis are quite supportive of a conclusion of the robust nature of our optimal policies with respect to the simple heuristics with which they were compared. The value of the optimal policies remains high, especially under the same distribution, higher actual mean case, with a difference of between 6% and 9% over the heuristics. This value reduces down to around an average of around 2% with respect to the first heuristic policy for different distribution cases, which corresponds to an additional \$30,000 PVI under a fund level of \$4 million. The only case where a heuristic policy performs better than our proposed policies (which are based on specific distribution assumptions) is when the actual distributions are different and the actual means are also lower than the assumed means. Although not so significant, the first heuristic policy performs slightly better under this case. Overall, we conclude that our policies are quite robust under different actual

Table 3. Comparison of Different Policies Under Deviations from Assumed Distributions

Distributions for R and C	Policy type (π)	Expected total PVI [$V^{\pi}(b = \$1\text{mil}, s = 0)$]	% Difference from optimal [$\frac{V^{opt} - V^{\pi}}{V^{opt}} \times 100\%$]
$R \sim$ Uniform and $C \sim$ Triangular with means for R and C increased by one standard deviation: that is, $\mu_R = \bar{\mu}_R + \bar{\sigma}_R$, and $\mu_C = \bar{\mu}_C + \bar{\sigma}_C$	Baseline optimal	$349,095	—
	Heuristic 1	$328,336	6.0
	Heuristic 2	$315,947	9.5
$R \sim$ Uniform and $C \sim$ Triangular with means for R and C decreased by one standard deviation: that is, $\mu_R = \bar{\mu}_R - \bar{\sigma}_R$, and $\mu_C = \bar{\mu}_C - \bar{\sigma}_C$	Baseline optimal	$359,203	—
	Heuristic 1	$353,097	1.7
	Heuristic 2	$355,822	0.9
$R \sim$ Triangular and $C \sim$ Normal with means for R and C kept at the baseline level: that is, $\mu_R = \bar{\mu}_R$, and $\mu_C = \bar{\mu}_C$	Baseline optimal	$331,466	—
	Heuristic 1	$322,812	2.6
	Heuristic 2	$304,308	8.2
$R \sim$ Triangular and $C \sim$ Normal with means for R and C increased by one standard deviation: that is, $\mu_R = \bar{\mu}_R + \bar{\sigma}_R$, and $\mu_C = \bar{\mu}_C + \bar{\sigma}_C$	Baseline optimal	$308,371	—
	Heuristic 1	$306,471	0.6
	Heuristic 2	$272,960	11.5
$R \sim$ Triangular and $C \sim$ Normal with means for R and C decreased by one standard deviation: that is, $\mu_R = \bar{\mu}_R - \bar{\sigma}_R$, and $\mu_C = \bar{\mu}_C - \bar{\sigma}_C$	Baseline optimal	$376,282	—
	Heuristic 1	$377,348	−0.3
	Heuristic 2	$359,224	4.5

Notes. $\bar{\mu}_R$, $\bar{\sigma}_R$, $\bar{\mu}_C$, and $\bar{\sigma}_C$ correspond to baseline means and standard deviations for R and C as described in Section 4.1. Baseline optimal policy corresponds to the price-dependent policy optimized according to Equation (11), and V^{opt} refers to the expected total PVI under this optimal policy. For all computations presented, we consider $b = \$1\text{mil}$ and $s = 0$.

distributions of property asking prices and PVI values, and average values of 2% and 7% can be achieved compared with the first and second heuristics, respectively, at least regarding the scenarios considered in Table 3.

Last but not least, we test the robustness of our results for the case where the indirect costs are negligible. Although assessing the impact of indirect costs necessitates that we conduct a thorough analysis under a modified value function as we discuss in the online appendix, our representative numerical studies yield that, when the indirect cost of preparing a bid is not taken into account, this can potentially increase the frequency of bidding but only at minimal levels. Because the asking prices for the properties in our data set already exceed the baseline threshold function $x(\delta|b, s, c)$, in our context, ignoring indirect costs does not really generate any noticeable impact on the resulting bidding decisions.

5. Conclusions

In this paper, we develop, implement, and evaluate a dynamic and stochastic decision model that aims at assisting community-based organizations when choosing foreclosed properties to acquire in the service of community stabilization and revitalization. We derive analytical results for calculating optimal return thresholds that a CDC can use to determine which properties it should make an offer on and how much to overbid. The policies are implemented in a numerical setting based on operational data from a CDC, which is considered to be reflective of the operating conditions of many other CDCs in major cities. General policy guidelines have been suggested based on this numerical study, where it is estimated that a potential increase of around 20% can be achieved in expected total PVI through optimal policy implementations compared with historical acquisition data. The proposed decision rules can be used in the form of spreadsheets or visual charts to aid in the decision making for a property with a given asking price and estimated property value impact. Through multiple simulation analyses, we also conclude that the proposed policies are quite robust under potential deviations from the assumed distributions in the problem framework. Overall, the presented models and guidelines can potentially aid CDCs and other similar organizations when making investment decisions with social returns, such as in foreclosed housing acquisition and redevelopment. Given the absence of any such optimization-based analysis tool for this type of nonprofit decision making, our results can help improve the efficiency and effectiveness of the decisions by these nonprofit organizations, thus contributing to social value creation.

Acknowledgments

The authors thank Rachel Drew from University of Massachusetts Boston for help in providing the data used in the numerical analysis. The authors also thank Michael Johnson and Jeffrey Keisler from University of Massachusetts Boston as well as David Turcotte and Emily Vidrine from University of Massachusetts Lowell, for comments on the models presented in the paper.

Endnote

[1] The notation used throughout the paper is summarized in the online appendix for reference purposes.

References

Hersh MB (2012) The hard part. Accessed April 25, 2012, https://shelterforce.org/2012/04/25/the_hard_part/.

Bayram A, Solak S, Johnson MP, Turcotte D (2011) Managing foreclosed housing portfolios for improved social outcomes. *Proc. 22nd Annual Production Oper. Management Soc. Conf., Reno, NV*.

Bernanke BS (2008) Housing, mortgage markets, and foreclosures. Accessed March 26, 2019, https://www.federalreserve.gov/newsevents/speech/bernanke20081204a.htm.

Bertsimas D, Popescu I (2003) Revenue management in a dynamic network environment. *Transportation Sci.* 37(3):257–277.

Bratt R (2009) Challenges for nonprofit housing organizations created by the private housing market. *J. Urban Affairs* 31(1):67–96.

Calafiore GC (2008) Multi-period portfolio optimization with linear control policies. *Automatica* 44(10):2463–2473.

Campbell JY, Giglio S, Pathak P (2011) Forced sales and house prices. *Amer. Econom. Rev*. 101(5):2108–2131.

Caskey J, Aobdia D (2012) Investor cost basis and takeover bids. Working paper, University of California, Berkeley.

Christie L (2009) Home values plummet $500 billion. *CNN Money* (December 9), http://money.cnn.com/2009/12/09/real_estate/home_value_loss/index.htm.

Community Wealth (2012) Overview: Community development corporations. Accessed December 14, 2015, http://www.community-wealth.org/strategies/panel/cdcs/index.html.

Dizdar D, Gershkov A, Moldovanu B (2011) Revenue maximization in the dynamic knapsack problem. *Theoret. Econom.* 6(3):157–184.

Gass A (2011) Implementing the neighborhood stabilization program. Technical report, NeighborWorks America, Washington, DC.

Harding JP, Rosenblatt E, Yao VW (2009) The contagion effect of foreclosed properties. *J. Urban Econom.* 66(3):164–178.

Holt CA, Sherman R (2000) Risk aversion and the winner's curse. Working paper, University of Virginia, Charlottesville.

Immergluck D, Smith G (2006) The external costs of foreclosure: The impact of single-family mortgage foreclosures on property values. *Housing Policy Debate* 17(1):57–79.

Johnson M, Turcotte D, Sullivan F (2010) What foreclosed homes should a municipality purchase to stabilize vulnerable neighborhoods? *Networks Spatial Econom.* 10(3):363–388.

Johnson M, Solak S, Drew R, Keisler J (2013) Property value impacts of foreclosed housing acquisitions under uncertainty. *Socio-Economic Planning Sci.* 47(4):292–308.

Johnson MP (2011) *Community-Based Operations Research: Decision Modeling for Local Impact and Diverse Populations* (Springer, New York).

Joint Center for Housing Studies (2009) The state of the nation's housing. Accessed May 12, 2015, http://www.jchs.harvard.edu/publications/markets/son2009/son2009.pdf.

Kaplan E (1984) Managing the demand for public housing. PhD thesis, Massachusetts Institute of Technology, Cambridge.

Kilic OA, Donk DP, Wijngaard J, Tarim SA (2010) Order acceptance in food processing systems with random raw material requirements. *OR Spectrum* 32(3):905–925.

Kleywegt AJ, Papastavrou JD (1998) The dynamic and stochastic knapsack problem. *Oper. Res.* 46(1):17–35.

Kleywegt AJ, Papastavrou JD (2001) The dynamic and stochastic knapsack problem with random sized items. *Oper. Res.* 4(1):26–41.

Larson RC, Odoni AR (1981) *Urban Operations Research* (Prentice-Hall, Englewood Cliffs, NJ).

Leonard T, Murdoch J (2009a) The impact of foreclosures on neighboring housing sales. *J. Real Estate Res.* 31(4):455–479.

Leonard T, Murdoch J (2009b) The neighborhood effects of foreclosure. *J. Geographical Systems* 11(4):317–332.

Lin GY, Lu Y, Yao DD (2008) The stochastic knapsack revisited: Switch-over policies and dynamic pricing. *Oper. Res.* 56(4):945–957.

Little JDC (1961) A proof for the queuing formula: L=λ w. *Oper. Res.* 9(3):383–387.

Loch CH, Kavadias S (2002) Dynamic portfolio selection of NPD programs using marginal returns. *Management Sci.* 48(10):1227–1241.

Lu Y (2005) Solving a dynamic resource allocation problem through continuous optimization. *IEEE Trans. Automatic Control* 50(6):890–894.

Lubin G (2011) U.S. home values lost $798 billion last quarter, nearly $10 trillion destroyed since peak. *Business Insider* (February 9), https://www.businessinsider.com/zillow-fourth-quarter-798-billion-2011-2.

Mild P, Salo A (2009) Combining a multiattribute value function with an optimization model: An application to dynamic resource allocation for infrastructure maintenance. *Decision Anal.* 6(3):139–152.

NeighborWorks (2009) Responsible approaches to neighborhood stabilization: Case studies in action. Technical report, NeighborWorks America, Washington, DC.

Nikolaev AG, Jacobson SH (2010) Stochastic sequential decision-making with a random number of jobs. *Oper. Res.* 58(4):1023–1027.

Pollock SM, Rothkopf MH, Barnett A (1994) *Operations Research in the Public Sector* (North-Holland, Amsterdam).

Stable Communities (2012) Neighborhood stabilization program (NSP) strategies. Accessed May 20, 2015, http://www.stablecommunities.org/nsp-strategies.

Stevens DH (2011) Federal Housing Administration (FHA) first look sales method for grantees, nonprofit organizations, and subrecipients under the neighborhood stabilization programs (NSP). Technical report, U.S. Department of Housing and Urban Development, Washington, DC.

Swanstrom T, Chapple K, Immergluck D (2009) Regional resilience in the face of foreclosures: Evidence from six metropolitan areas. Working paper, University of California, Berkeley, Berkeley.

Senay Solak is an associate professor in the Department of Operations and Information Management at Isenberg School of Management, University of Massachusetts Amherst. His research interests involve portfolio management for technology and capital investment projects, with specific applications in the nonprofit sector, and air transportation planning.

Armagan Bayram is an assistant professor in the Department of Industrial and Manufacturing Systems Engineering at the University of Michigan Dearborn. Her research interests include the development of stochastic models and solution methods for capacity and resource allocation problems that involve nonprofit, healthcare, and mobility applications.

Mehmet Gumus is an associate professor of operations management and the academic director for the Masters of Management in Analytics Program at the Desautels Faculty of Management at McGill University. In his research, he explores the impact of customer behavior and information asymmetry on supply chain management, dynamic pricing, and risk management.

Yueran Zhuo is a postdoctoral researcher at Harvard Medical School and a clinical research coordinator at the Institute for Technology Assessment in Massachusetts General Hospital. Her research interests involve information security investments, medical decision making, and healthcare economics.