

## SYNTHETIC CENSUS DATA GENERATION VIA MULTIDIMENSIONAL MULTISSET SUM\*

CYNTHIA DWORK, KRISTJAN GREENEWALD, AND MANISH RAGHAVAN

Department of Computer Science, Harvard University

MIT-IBM Watson AI Lab; IBM Research, Cambridge, MA, USA

MIT Sloan School of Management and Department of EECS  
*e-mail address:* mragh@mit.edu

**ABSTRACT.** The US Decennial Census provides valuable data for both research and policy purposes. Census data are subject to a variety of disclosure avoidance techniques prior to release in order to preserve respondent confidentiality. While many are interested in studying the impacts of disclosure avoidance methods on downstream analyses, particularly with the introduction of differential privacy in the 2020 Decennial Census, these efforts are limited by a critical lack of data: the underlying “microdata,” which serve as necessary input to disclosure avoidance methods, are kept confidential.

In this work, we aim to address this limitation by providing tools to generate synthetic microdata solely from published Census statistics, which can then be used as input to any number of disclosure avoidance algorithms for the sake of evaluation and carrying out comparisons. We define a principled distribution over microdata given published Census statistics and design algorithms to sample from this distribution. We formulate synthetic data generation in this context as a knapsack-style combinatorial optimization problem and develop novel algorithms for this setting. While the problem we study is provably hard, we show empirically that our methods work well in practice, and we offer theoretical arguments to explain our performance. Finally, we verify that the data we produce are “close” to the desired ground truth.

### 1. INTRODUCTION

Scholars, practitioners, and policy-makers rely on US Decennial Census data for a wide range of research and decision-making tasks. For privacy reasons, Census data are not released in full. All released data are instead subject to a variety of disclosure avoidance practices. For example, data may be perturbed before release or have their location information coarsened. The unperturbed “microdata” are kept secret for 72 years after their collection.

*Key words and phrases:* Decennial Census, multidimensional subset sum, synthetic data generation.

\* A version of this paper was published as a one-page extended abstract in the 2024 Symposium on Foundations of Responsible Computing (Dwork et al., 2024).

The Census Bureau updated its disclosure avoidance system to use differential privacy for the release of the 2020 Decennial Census instead of its prior swapping-based methodology (Abowd, 2018b; Abowd et al., 2022). This sparked renewed interest in the properties of various privacy-preserving methods and their impacts on downstream consumers of Census data products. For example, decisions involving budgeting, voting rights and redistricting, and planning rely on accurate and consistent Census data (Cohen et al., 2021; Steed et al., 2022). Social scientists also rely on Census data for a wide range of research including health and social mobility (Ruggles et al., 2019). These stakeholders have begun to ask a critical question: How reliable can these analyses be under privacy-preserving techniques?

A key obstacle to answering this question is the secrecy of the data themselves: while the TopDown algorithm used by the Census in 2020 is public, the underlying data on which it is run are not. Ideally, one could simply simulate multiple runs of the TopDown algorithm (or any other proposed alternative) to characterize its effects on downstream quantities of interest. But the algorithm(s) in question take as input the underlying *microdata*, which are only released 72 years after they are collected.

Prior research on disclosure avoidance and the Census has used a variety of workarounds, including a limited sample of public Census demonstration data (Kenny et al., 2024; Dick et al., 2023), older Census data (Bailie et al., 2023; Petti and Flaxman, 2019), or heuristic methods to generate microdata (Christ et al., 2022; Cohen et al., 2021). We discuss these heuristics in greater detail in Section 3.

In this work, we seek to enable comprehensive research into the US Census, including research on privacy, by providing a principled method to generate synthetic Census microdata from publicly available data sources. Our aim is to enable research into the impacts of privacy-preserving technology on Census data beyond the above-mentioned heuristic workarounds. Moreover, our tools could provide a starting point for Census data consumers to estimate and potentially correct for the biases induced by disclosure avoidance algorithms by estimating how they affect quantities of interest, which we discuss further in Section 7.

At a high level, we combine block-level aggregate statistics with a random sample of microdata containing only coarse location information (Beckman et al., 1996). Importantly, our goals and methods differ from those of many Census-specific “reconstruction attacks,” which seek to analyze privacy-preserving methods by testing whether and how many rows of the microdata can be reconstructed from publicly released information (Abowd, 2018a, 2021; Dick et al., 2023; Francis, 2022) (see also: Dinur and Nissim (2003)). Reconstruction attacks typically seek to find the *most likely* microdata given the available information, which will in general lead to a more homogeneous dataset at a population level. In contrast, our aim is to sample from a representative distribution over microdata. We do not use any auxiliary information (i.e., non-Census data products), and we seek to generate representative, state-wide synthetic microdata, instead of a fraction of the rows. We intend for researchers to perform downstream analyses over multiple samples of this microdata. If a finding (e.g., that a particular disclosure avoidance method biases a statistic of interest) holds across multiple samples from this distribution, we may be more concerned that it holds for the ground truth data as well. Note that we do not intend for our synthetic data to be interpreted as “ground truth”; our goal is to provide synthetic data that are both faithful to published information and statistically plausible.

At a technical level, we formulate synthetic data generation in this context as a knapsack-style combinatorial optimization problem. Given aggregate statistics for each Census block (e.g., number of households, number of individuals of each race, . . .), we seek to sample

households that, when put together, exactly match the aggregate statistics reported by the Census. We design a Markov Chain Monte Carlo algorithm to sample appropriate households. While the problem we seek to solve is NP-hard, we provide both theoretical and empirical evidence to show that our methods perform sufficiently well to be viable in this setting, allowing us to sample datasets across entire US states. We provide code and detailed instructions for others to generate their own synthetic data at <https://github.com/mraghavan/synthetic-census>. Our implementation is specific to the 2010 US Decennial Census, but our broader framework can be adapted to general population synthesis tasks.<sup>1</sup>

**Organization of the paper.** In Section 2, we formalize synthetic data generation in our setting as a combinatorial optimization problem. We discuss related work in Section 3. In Section 4, we present and analyze a pair of Markov Chain Monte Carlo (MCMC) algorithms, finding that our approach empirically performs much better than generalizations of prior work. We describe our overall hybrid integer linear programming-MCMC algorithm in Section 5. In Section 6, we evaluate the representativeness of data sampled via this algorithm. We discuss the implications and usage of our methods in Section 7.

## 2. PROBLEM FORMULATION

**Notation.** We will denote nonnegative integer vectors with bold lower-case letters (e.g.,  $\mathbf{x}$ ) and use calligraphic upper-case letters (e.g.,  $\mathcal{X}$  or  $\mathcal{D}$ ) to denote sets and probability distributions. We will use bold upper-case letters for nonnegative integer matrices (e.g.,  $\mathbf{V}$ ). We will write  $\mathbf{x}_i$  to denote the  $i$ th vector in a set and  $\mathbf{x}[i]$  to denote the  $i$ th entry of vector  $\mathbf{x}$ , indexing vectors beginning with 1. We use  $\preceq$  and  $\succeq$  to denote a vector being element-wise  $\leq$  (resp.  $\geq$ ) another vector. We will use  $\mathbf{x}_{i \leftarrow g}$  to denote  $\mathbf{x}$  with its  $i$ th entry replaced by the value  $g$ .  $[n]$  refers to the set  $\{1, \dots, n\}$ . For a random variable  $\mathbb{X}$  with distribution  $\sigma$  over a discrete set  $\mathcal{X}$ , we will write  $\sigma(\mathbf{x}) = \Pr[\mathbb{X} = \mathbf{x}]$  for  $\mathbf{x} \in \mathcal{X}$  and  $\sigma(\mathcal{S}) = \Pr[\mathbb{X} \in \mathcal{S}]$  for  $\mathcal{S} \subseteq \mathcal{X}$ . To refer to the conditional distribution of  $\sigma$  on  $\mathcal{S} \subseteq \mathcal{X}$ , we write  $\sigma \mid \mathbf{x} \in \mathcal{S}$ .

**Empirical setting.** Our goal is to generate synthetic microdata based on the 2010 US Decennial Census. Cleaned Census responses are collected in a dataset known as the Census Edited File (CEF) often referred to as “microdata.” To meet its statutory privacy obligations, the Census Bureau does not release this dataset. Instead, they apply a suite of disclosure avoidance techniques (including adding noise, censoring outliers, etc.) before releasing aggregate statistics. Often, statistics are released at the Census block level, where a Census block typically consists of at most a few hundred households. In particular, “Summary File 1” (SF1) provides granular demographic information for each Census block *after* disclosure avoidance techniques are used.

For each block, we will seek to sample a collection of households whose characteristics match statistics reported in SF1.<sup>2</sup> As a result, our data will exactly match SF1 along all the attributes we choose (detailed below). Importantly, our methods preserve structural zeros:

<sup>1</sup>Changes to the 2020 Census prevent our methods from being directly applicable. In particular, the 2020 Census includes the “Privacy-Protected Microdata File” (PPMF), which is a 100% enumeration of synthetic persons and households (U.S. Census Bureau, 2024). However, because the persons and household files are separate, future work could adapt our techniques to create synthetic microdata by combining these datasets.

<sup>2</sup>We obtain SF1 data via IPUMS (Ruggles et al., 2024).

if SF1 reports zero people with certain characteristics in a block, we ensure that our data have the same property.

The statistics included in SF1 are fairly detailed and include counts of individuals with various attributes (e.g., number of Hispanic persons) and detailed household types (e.g., number of households with three members headed by a householder of two or more races).<sup>3</sup> For a given block, we will denote these statistics by a nonnegative integer vector  $\mathbf{c} \in \mathbb{Z}_{\geq 0}^d$ . An important feature we will rely on is the fact that  $\mathbf{c}$  encodes the total number of households in a block, which we will refer to as  $m$ . In our case, each vector has dimension  $d = 135$ , where each dimension encodes a particular count. For more details, see Appendix A.

In addition to SF1, the Census Bureau also releases the Public Use Microdata Sample (PUMS), consisting of a 10% sample of households across each state.<sup>4</sup> Crucially, for privacy reasons, the PUMS does not contain granular geographic locations for each household; each household is annotated with the Public Use Microdata Area (PUMA) in which it resides. For our purposes, we ignore PUMA information and treat the PUMS as simply a statewide sample.<sup>5</sup> We will treat each distinct household in the PUMS dataset as a vector  $\mathbf{v}_i \in \mathbb{Z}_{\geq 0}^d$  (encoded in the same  $d = 135$ -dimensional space as each block). As before, each entry represents the count of a certain property, which can either be the number of individuals in a household satisfying a particular demographic property or the (binary) indicator for whether the household as a whole satisfies a property. The frequency of each distinct household type in this encoding yields a distribution  $\mathcal{D}$ . Let  $n$  be the number of distinct household types, each represented by a vector  $\mathbf{v}_i$ , in a given state (in our case,  $n$  is on the order of a few thousand), and define  $\mathbf{V} \in \mathbb{Z}_{\geq 0}^{d \times n}$  to be the matrix with columns  $\mathbf{v}_i$ . It will sometimes be convenient to refer to the set of household type vectors, which we will denote  $\mathcal{V} \triangleq \{\mathbf{v}_1, \dots, \mathbf{v}_n\}$ .

With this data, our goal is as follows. A “solution” to a block  $b$  is a multiset  $\mathbf{x} \in \mathbb{Z}_{\geq 0}^n$  (represented as a vector of cardinalities of each element) such that  $\mathbf{V}\mathbf{x} = \mathbf{c}_b$ . Note that this linear equality exactly captures the constraint that, when summed together, the characteristics of the multiset of households exactly match those reported in SF1.

**Handling multiplicity.** If each block had a unique  $\mathbf{x}$  satisfying  $\mathbf{V}\mathbf{x} = \mathbf{c}_b$ , then this would suffice—we could find  $\mathbf{x}$  for each block  $b$ , producing the entire microdata. This is not the case.<sup>6</sup> As a simplified example, consider a block with four individuals, two white and two Asian, split into two households of size two. Without further information, there are two possibilities: the block could contain either two racially homogeneous households or two multiracial households.

When faced with multiple possible solutions underlying a given block, what should we do? We could of course choose arbitrarily. The Census’ own demonstration reconstruction attack appears to take this approach (Abowd, 2018a). But in the above example, we might believe that statistically speaking, it is more likely that a block contains two racially homogeneous households than two multiracial households. This intuition is borne out in

<sup>3</sup>We choose of a subset of the SF1 statistics to match, which we describe in more detail in Appendix A.

<sup>4</sup><https://www.census.gov/data/datasets/2010/dec/stateside-pums.html>

<sup>5</sup>Each PUMA has at least 100,000 individuals in it. We aggregate to the state level because otherwise, the data are too sparse for our methods to be effective. As a result, we will fail to capture regional variation in household composition that are not explained by SF1. Intuitively, our formulation makes the assumption that, conditioned on the SF1 counts, the distribution of households is location-invariant within the state.

<sup>6</sup>For privacy reasons, this is to be expected.

the PUMS microdata sample: multiracial households are far less frequent than racially homogeneous households. How should this information inform our sample when multiple possible reconstructions exist?

One approach would be to try to find the “most likely” reconstruction (for some definition of likelihood we will have to make formal). If our goal was a reconstruction attack, this might be the right choice. But if we do this for all blocks, our overall sample will be quite biased. For example, if racially homogeneous households are much more frequent than multiracial households, our resulting dataset will contain very few multiracial households relative to what we know about the overall population.

Instead, we take a different approach designed to produce a more representative sample. Let  $\mathcal{X}_b \triangleq \{\mathbf{x} : \mathbf{V}\mathbf{x} = \mathbf{c}_b\}$  be the set of all valid solutions for a given block. In other words,  $\mathcal{X}_b$  is the set of combinations of households that, when summed together, match the aggregate block-level counts. We will specify a distribution  $\pi$  over  $\mathcal{X}_b$  based on the PUMS distribution  $\mathcal{D}$  and seek to sample from  $\mathcal{X}_b$  according to  $\pi$ . Ideally, our choice of  $\pi$  would lead to a representative statewide sample, meaning that when sampling across the entire state, the expected frequency of each household in our sampled microdata matches its empirical PUMS frequency in  $\mathcal{D}$ . Unfortunately, this would require  $\pi$  to depend on  $\mathbf{c}_b$  for all  $b$  in a state, which would be prohibitively expensive. Instead, we specify a natural generative model to induce a distribution  $\pi$ . We evaluate the representativeness of samples produced by  $\pi$  in Section 6.

**A generative model.** Consider a generative model in which we are given the PUMS distribution  $\mathcal{D}$  over the set of household types  $\mathcal{V} \triangleq \{\mathbf{v}_1, \dots, \mathbf{v}_n\}$ . Assume that for a given Census block  $b$ , aggregate statistics  $\mathbf{c}_b$  are chosen exogenously. (We will often drop the subscript  $b$  for ease of notation.) Then, a multiset  $\mathbf{x}$  of households is sampled i.i.d. according to  $\mathcal{D}$ . We are interested in the distribution over multisets of households that this produces conditioned on the event that the aggregate characteristics of sampled households exactly matches the reported statistics: that is, conditioned on  $\mathbf{V}\mathbf{x} = \mathbf{c}$ .

Intuitively, we can think of this as the distribution induced via rejection sampling; indeed, a basic (and prohibitively inefficient) algorithm to sample from this distribution is to repeatedly sample households i.i.d. from  $\mathcal{D}$  until either  $\mathbf{V}\mathbf{x} = \mathbf{c}$  or  $(\mathbf{V}\mathbf{x})[i] > \mathbf{c}[i]$  for some  $i$ , and accept the first  $\mathbf{x}$  that satisfies  $\mathbf{V}\mathbf{x} = \mathbf{c}$ . (See Algorithm 1 in Appendix B.) In this generative model, for  $\mathbf{x}$  such that  $\mathbf{V}\mathbf{x} = \mathbf{c}$ , the posterior  $\pi$  induced by rejection sampling is given (up to a normalizing constant) by

$$\pi(\mathbf{x}) \propto f(\mathbf{x}) \triangleq \binom{\|\mathbf{x}\|_1}{\mathbf{x}} \prod_{i=1}^n \Pr_{\mathcal{D}}[\mathbf{v}_i]^{\mathbf{x}[i]}, \quad (2.1)$$

where the multinomial coefficient, defined  $\binom{\|\mathbf{x}\|_1}{\mathbf{x}} = \|\mathbf{x}\|_1! / (\mathbf{x}[1]! \dots \mathbf{x}[n]!)$  accounts for the multiplicity of elements in  $\mathbf{x}$ . (Recall that  $m = \|\mathbf{x}\|_1$ , the number of households in a block, is known for each block.) For  $\mathbf{x}$  such that  $\mathbf{V}\mathbf{x} \neq \mathbf{c}$ , we define  $\pi(\mathbf{x})$  to be 0. With this, we have specified a combinatorial optimization problem: sample from  $\mathcal{X} = \{\mathbf{x} : \mathbf{V}\mathbf{x} = \mathbf{c}\}$  according to the distribution  $\pi$  in (2.1). We formalize this below.

**Formal problem definition.** Sampling from  $\pi(\cdot)$  over  $\mathcal{X}$  is closely related to the classical subset sum problem. In subset sum, the goal is to choose a subset of integers that sum to a desired target. This can be generalized to the multidimensional subset sum problem, in which we choose a subset of integer vectors that sum to a target vector. The problem we seek to solve (sampling from  $\mathcal{X}$ ) can be thought of as a (nonnegative) “multidimensional multiset sum” problem, similar to multidimensional subset sum but with a slight modification: we may choose vectors with replacement. We define the decision version of the multidimensional subset sum (**MMS**) problem and its sampling analogue (**MMS-Sample**) below:

**Problem 1 (MMS).** Given a matrix  $\mathbf{V} \in \mathbb{Z}_{\geq 0}^{n \times d}$  and vector  $\mathbf{c} \in \mathbb{Z}_{\geq 0}^d$ , determine whether  $\mathcal{X} = \{\mathbf{x} \in \mathbb{Z}_{\geq 0}^n : \mathbf{V}\mathbf{x} = \mathbf{c}\}$  is non-empty.

**Problem 2 (MMS-Sample).** Given a matrix  $\mathbf{V} \in \mathbb{Z}_{\geq 0}^{n \times d}$  and a vector  $\mathbf{c} \in \mathbb{Z}_{\geq 0}^d$ , let  $\mathcal{X} = \{\mathbf{x} \in \mathbb{Z}_{\geq 0}^n : \mathbf{V}\mathbf{x} = \mathbf{c}\}$ . Given a distribution  $\sigma$  over  $\mathcal{X}$  specified up to a normalizing constant, sample from  $\sigma$ .

**MMS** is NP-hard (see Claim 1 in Appendix F for details), making it hard to approximate  $|\mathcal{X}|$  to within a constant factor. As a result, **MMS-Sample** is NP-hard (Jerrum et al., 1986). Moreover, relaxing the constraint to  $\mathbf{V}\mathbf{x} \approx \mathbf{c}$  is unlikely to make the problem significantly easier, since approximate multidimensional subset sum is NP-hard for  $d > 2$  (Emiris et al., 2017; Magazine and Chern, 1984; Kulik and Shachnai, 2010). Our key technical contribution is an algorithm for **MMS-Sample** that works sufficiently well in practice for the purposes of generating synthetic Census microdata. We focus on sampling from  $\pi$  given by (2.1), but our algorithms work for any distribution  $\sigma$  specified up to a normalizing constant. Before describing our solution, we discuss related work on population synthesis, Census disclosure avoidance, and similar combinatorial optimization problems.

### 3. RELATED WORK

**Population synthesis.** Researchers have sought to combine microdata with tabular data to generate synthetic populations for a variety of purposes including agent-based simulation and transportation analyses (Beckman et al., 1996). At a high level, this body of literature focuses on two steps: building models of population distributions and sampling from those distributions to meet given constraints (Müller and Axhausen, 2010). Much of the focus of this literature has been on building better models of populations (e.g. Creecy, 2009; Farooq et al., 2013; Casati et al., 2015; Sun and Erath, 2015; Sun et al., 2018), often by modeling complex relationships between demographic attributes or hierarchical household-individual relationships. This is particularly valuable when dealing with rich, high-dimensional data. More recent work has also used machine-learning approaches to model populations (Albiston et al., 2024; Gussenbauer et al., 2024). In contrast, we take a particularly simple approach to population modeling, treating the PUMS as a population model.

Given a population model, the literature contains a few high-level strategies to produce a sample satisfying aggregate constraints like the ones imposed by SF1. Early techniques like iterative proportional fitting (IPF) (Deming and Stephan, 1940) seek to fit weights for each member of a population distribution such that sampling from the population distribution according to those weights would yield the desired aggregate counts in expectation (Birkin and Clarke, 1988; Beckman et al., 1996). In our notation, these methods find real-valued

$\mathbf{x} \in \mathbb{R}_{\geq 0}^n$  such that  $\mathbf{V}\mathbf{x} = \mathbf{c}$ , and they seek to align  $\mathbf{x}$  with the population model  $\mathcal{D}$ . They do not guarantee that the constraints are met *ex post*, i.e., after an actual dataset is sampled. For large geographic regions, sampling this way will approximately satisfy the given constraints. But for small geographic regions (e.g., Census blocks), sampling schemes yield high variance, meaning generated datasets will differ significantly from the given constraints. A class of heuristics known as “deterministic reweighting” schemes seek to reduce the variance produced by sampling (Ballas et al., 2005; Creecy, 2009; Lovelace and Ballas, 2013; Casati et al., 2015), but these methods tend to either fail to provide exact guarantees or introduce bias into the sample (Müller and Axhausen, 2010).

Most related to ours, a line of work draws on combinatorial optimization heuristics like hill-climbing, simulated annealing, and genetic algorithms to construct a sample that comes as close as possible to matching the aggregate counts (Williamson et al., 1998; Voas and Williamson, 2000; Harland et al., 2012; Ma and Srinivasan, 2015; Wu et al., 2022; Whitworth, 2022; Gussenbauer et al., 2024). The quality of these approaches is often measured by the (sometimes squared) difference between the aggregate counts and the sample counts  $\|\mathbf{V}\mathbf{x} - \mathbf{c}\|_1$ . In general, these heuristics fail to guarantee  $\mathbf{V}\mathbf{x} = \mathbf{c}$ .

Our primary technical contribution lies in sampling, not in modeling the population distribution (since we adopt the PUMS as our population model). In contrast to prior work, our work guarantees exact matches to the aggregate counts. If this was our only goal, standard integer linear programming (ILP) would suffice, since our problem setting is simple enough that we do not need to rely on heuristic methods from combinatorial optimization. However, instead of simply finding a single solution, our methods are designed to sample from a known distribution over all possible solutions. We could in principle apply our techniques to use more complex population models than the simple PUMS distribution used here; we defer such investigations to future work.

**Analyzing Census disclosure avoidance systems.** Prior work that has generated synthetic microdata for the purposes of analyzing Census disclosure avoidance systems (Cohen et al., 2021; Christ et al., 2022) has relied on heuristics that do not produce reliable household-level data. Cohen et al. (2021) explicitly note that their synthetic data do not contain household information, preventing them from fully replicating the Census Bureau’s disclosure avoidance system. (They do perform experiments in which they arbitrarily group individuals into households of size 5 in their synthetic data.) Christ et al. (2022) use a combination of heuristics to produce a limited sample of data by randomly selecting blocks and pooling data. This enables them to generate individual-level data that bear some resemblance to the ground truth. In contrast, our methods generate state-wide microdata at the household level. Household-level data are strictly more general than individual-level microdata, since we can produce an individual-level dataset simply by enumerating individuals in each household.

Our work bears some resemblance to recent “reconstruction attacks” on Census data, which attempt to reconstruct rows of the dataset given Census statistics and potentially external information (Abowd, 2021; Dick et al., 2023; Francis, 2022). In particular, the Census Bureau conducted a reconstruction attack which used integer programming to combine information across multiple tables to produce microdata (Abowd, 2018a). Our work extends the scope of these efforts in a number of ways. First, in order to facilitate the analysis of the impacts of disclosure avoidance techniques like swapping, we produce household-level data instead of individual-level data. Second, we integrate information about the distribution

of households from the PUMS, allowing us to produce a more representative dataset. As discussed earlier, our goal of representativeness is quite different from reconstruction: maximizing reconstruction “accuracy” would suggest choosing the most likely sample, which would be fairly homogeneous and unrepresentative; in contrast, we seek to produce a dataset that matches state-wide household statistics, which introduces additional complexity. In principle, if we were interested in adapting our methods to produce an effective reconstruction attack at the household level, we could sample  $\mathbf{x}^* = \arg \max_{\mathbf{x} \in \mathcal{X}} \pi(\mathbf{x})$ . In Appendix D, we provide a linear approximation to  $\pi(\cdot)$  that could be used to solve this via integer linear programming.

**Multidimensional knapsack, subset sum, and related combinatorial optimization problems.** Prior work considers closely related problems of knapsack sampling/counting and systems of linear Diophantine equations (equations of the form  $\mathbf{V}\mathbf{x} = \mathbf{c}$  with integrality constraints). While there exist polynomial-time approximation schemes for knapsack counting (Morris and Sinclair, 2004; Dyer, 2003; Gopalan et al., 2011; Rizzi and Tomescu, 2014; Gawrychowski et al., 2018; Kayibi et al., 2018; Lawler, 1979), these techniques do not generalize to our setting because (1) we require *exact*, not just *feasible* solutions (i.e.,  $\mathbf{V}\mathbf{x} = \mathbf{c}$  instead of  $\mathbf{V}\mathbf{x} \leq \mathbf{c}$ ); and (2) these algorithms scale exponentially with dimension  $d$ . Because **MMS-Sample** is NP-hard to approximate, we should not expect a polynomial-time approximation scheme. Other related work includes generalizations of the knapsack problem (Hendrix and Jones, 2015) and dynamic programming approaches (Bossek et al., 2021), which are empirically too slow in our high-dimensional setting.

Several specialized algorithms for knapsack and subset sum-style problems have appeared in the literature, some of which extend to multidimensional settings (e.g., Cabot, 1970; Ingargiola and Korsh, 1977; Puchinger et al., 2010; Salkin and De Kluyver, 1975; Martello and Toth, 1987; Pisinger, 1999). We do not experiment with them here since our own implementations of these methods are unlikely to compete with general-purpose but highly optimized integer linear programming packages, which we make heavy use of. Future work might be able to take advantage of these to further optimize the methods we develop here.

Our work draws most closely on the MCMC algorithm of Dyer et al. (1993). A direct adaptation of their algorithm (Section 4.1) is still too slow in our setting, but we develop a new MCMC approach that works better empirically (Section 4.2). We characterize the performance of our algorithms using a long line of theoretical results on the mixing time of Markov chains (Diaconis and Stroock, 1991; Lawler and Sokal, 1988; Jerrum and Sinclair, 1988; Sinclair and Jerrum, 1989; Jerrum and Sinclair, 1989).

**MMS-Sample** can also be described as sampling nonnegative solutions to a system of linear Diophantine equations. While these systems have been studied extensively (Blankinship, 1966; Bradley, 1971; Chou and Collins, 1982; Lazebnik, 1996; Brádler, 2016; Aardal et al., 2000; Sánchez-Roselly Navarro, 2016), algorithms in this setting generally produce integer solutions, not *nonnegative* integer solutions. (Recall that we require nonnegativity because a multiset of households cannot contain a negative number of copies of a household.) The literature contains some results on determining the existence of or bounding the number of nonnegative solutions (Brádler, 2016; Mahmoudvand et al., 2010), but these are not algorithmic in nature.

In addition to all of these techniques, we have one more tool at our disposal: Integer Linear Programming (ILP). For sufficiently small problems, we can use highly optimized

software packages<sup>7</sup> for ILP to simply enumerate  $\mathcal{X}$  and sample according to  $\pi$  as desired. In all of our instances (where each instance is a Census block), we find that ILP suffices to determine whether  $\mathcal{X}$  is non-empty. Enumeration, however, has a clear downside: its complexity scales with  $|\mathcal{X}|$ , which may be exponentially large in  $n$ , the number of possible households. For example, using ILP to enumerate up to 5000 elements of  $\mathcal{X}_b$  from each Census block in  $b$  Alabama and Nevada, we plot the distribution of  $\min(|\mathcal{X}_b|, 5000)$  in Figure 1. Our results suggest that  $|\mathcal{X}_b|$  has a heavy-tailed distribution,<sup>8</sup> meaning that enumerating  $\mathcal{X}_b$  completely for each block is likely to be computationally infeasible.<sup>9</sup> For instances where  $|\mathcal{X}_b|$  is too large, we need more efficient sampling algorithms. We turn to Markov Chain Monte Carlo methods for this.

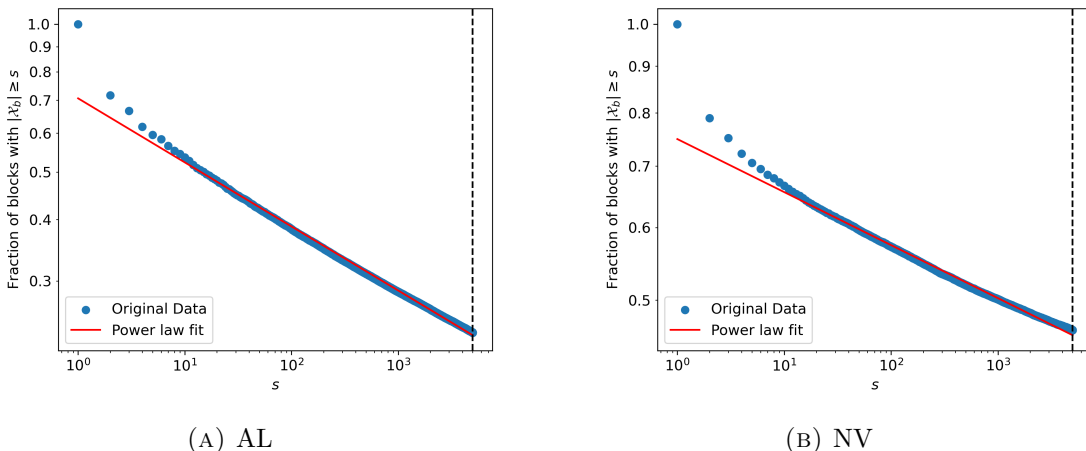


FIGURE 1.  $|\mathcal{X}_b|$ , the number of solutions in each block  $b$ , appears to be heavy-tailed.

#### 4. MARKOV CHAIN MONTE CARLO METHODS

Markov Chain Monte Carlo methods can allow us to efficiently sample from exponentially large state spaces and have been used in prior work on knapsack sampling (Dyer et al., 1993; Morris and Sinclair, 2004). At a high level, MCMC works by defining a Markov chain over the solution space, performing a random walk on this Markov chain, and yielding a solution after a fixed number of random walk steps. Here, we develop MCMC techniques to solve **MMS-Sample**. We present two approaches: the ‘simple’ chain, a modification of the algorithm of Dyer et al. (1993), and the ‘reduced’ chain, which can be interpreted as a truncation of the simple chain. These come with trade-offs that we evaluate both theoretically and empirically. At a high level, the simple chain is guaranteed to converge to the desired stationary distribution but may mix slowly. In contrast, the reduced chain may not converge

<sup>7</sup>We use Gurobi (<https://www.gurobi.com/>) in our experiments. Free alternative solvers can also be used in its place.

<sup>8</sup>We do not attempt to evaluate whether the power-law distribution fits these data well. Figures 1a and 1b are simply meant to be illustrative.

<sup>9</sup>For reference, enumerating up to 5000 solutions from  $\mathcal{X}_b$  for each block  $b$  takes thousands of CPU-hours for Alabama and Nevada, which have 135,838 and 35,916 non-empty Census blocks respectively.

to the desired distribution, but empirically performs better in our setting. Neither comes with strong theoretical guarantees; this is to be expected, since **MMS-Sample** is provably hard in the general case.

**Preliminaries.** As is standard for MCMC, we will evaluate our algorithms in terms of the computation required to produce an approximately random sample from  $\pi$  on  $\mathcal{X}$ . For distributions  $\sigma, \sigma'$  over  $\mathcal{X}$ , define  $d_{\text{TV}}$  to be the total variation distance:

$$d_{\text{TV}}(\sigma, \sigma') \triangleq \frac{1}{2} \|\sigma - \sigma'\|_1 = \max_{\mathcal{S} \subset \mathcal{X}} |\sigma(\mathcal{S}) - \sigma'(\mathcal{S})|. \quad (4.1)$$

When using MCMC to sample from a state space, we typically begin with a starting state chosen from some initial distribution  $\sigma_0$ , randomly transition for  $t$  steps, and return the final state. Let  $\mathbb{X}$  denote a random sample generated this way, and let  $\sigma_{\mathbb{X}}$  be its distribution. Our goal is to produce an  $\varepsilon$ -approximate sample from a target distribution  $\sigma$ :<sup>10</sup>

**Definition 1** ( $\varepsilon$ -approximate sample). A random variable  $\mathbb{X}$  with distribution  $\sigma_{\mathbb{X}}$  is an  $\varepsilon$ -approximate sample from a distribution  $\sigma$  if  $d_{\text{TV}}(\sigma, \sigma_{\mathbb{X}}) \leq \varepsilon$ .

This is closely related to the *mixing time* of a Markov chain with transition matrix  $P$  with stationary distribution  $\sigma$ :

$$\tau_{\text{mix}}(\varepsilon; P) \triangleq \min \left\{ t \in \mathbb{Z}_{\geq 0} : \max_{\sigma_0} d_{\text{TV}}(\sigma_0 P^t, \sigma) \leq \varepsilon \right\}.$$

However, because the algorithms we present each have a known initial state  $\sigma_0$ , we are not interested in the worst-case over all  $\sigma_0$ . Instead, we use a variant of the mixing time given a known starting distribution  $\sigma_0$  with probability 1 on initial state  $\mathbf{x}_0$  and 0 elsewhere. Define

$$\tau_{\text{mix}}(\varepsilon; P, \mathbf{x}_0) \triangleq \min \left\{ t \in \mathbb{Z}_{\geq 0} : d_{\text{TV}}(\sigma_0 P^t, \sigma) \leq \varepsilon \right\} \quad (4.2)$$

to be the number of iterations required to generate an  $\varepsilon$ -approximate sample from the stationary distribution  $\sigma$ . Standard spectral techniques yield bounds for  $\tau_{\text{mix}}(\varepsilon; P, \mathbf{x}_0)$ . Let  $P$  be the transition matrix of an irreducible Markov chain. Assume that the chain is ‘lazy,’ i.e.,  $P(\mathbf{x}, \mathbf{x}) \geq 1/2$  for all  $\mathbf{x}$ .<sup>11</sup> Let  $\lambda_2(P)$  be the second-largest eigenvalue of  $P$ . Then, the *relaxation time* of  $P$  is defined as  $\tau_{\text{rel}} \triangleq \frac{1}{1 - \lambda_2(P)}$ . Classical results tell us that  $\tau_{\text{rel}}$  provides a tight characterization of  $\tau_{\text{mix}}(\varepsilon; P, \mathbf{x}_0)$ :

**Theorem 1** (E.g., Guruswami (2016, Prop. 2.3); see also Levin et al. (2017, Thm. 12.5)). For an irreducible, aperiodic Markov chain,

$$(\tau_{\text{rel}} - 1) \log \left( \frac{1}{2\varepsilon} \right) \leq \tau_{\text{mix}}(\varepsilon; P, \mathbf{x}_0) \leq \tau_{\text{rel}} \log \left( \frac{1}{\varepsilon \sigma(\mathbf{x}_0)} \right). \quad (4.3)$$

With these definitions, we are ready to describe and analyze our two MCMC-based algorithms.

#### 4.1. The simple chain.

<sup>10</sup>Throughout this paper, we will choose the target distribution to be  $\pi$  as defined in (2.1).

<sup>11</sup>This guarantees that all eigenvalues are nonnegative.

4.1.1. *Defining the simple chain.* We begin by adapting the algorithm of Dyer et al. (1993) which samples (approximately) uniformly from the set of feasible knapsack solutions  $\mathcal{Y} \triangleq \{\mathbf{x} : \mathbf{V}\mathbf{x} \preceq \mathbf{c}\}$ . Our adaptation has two key differences: we are interested in exact solutions (i.e.,  $\mathcal{X} = \{\mathbf{x} : \mathbf{V}\mathbf{x} = \mathbf{c}\}$ ), and we want to sample from a particular distribution  $\pi$  as defined in (2.1). At a high level, we will design a Markov chain  $M_\gamma$  parameterized by  $\gamma \in \mathbb{R}$  with stationary distribution  $\tilde{\pi}_\gamma$ :

$$\tilde{\pi}_\gamma(\mathbf{x}) \propto f(\mathbf{x}) \exp(-\gamma \|\mathbf{V}\mathbf{x} - \mathbf{c}\|_1) \quad (4.4)$$

over  $\mathcal{Y}$ . Observe that the conditional distribution of  $\tilde{\pi}_\gamma$  over  $\mathcal{X}$  is  $\pi$ , i.e.,  $(\tilde{\pi}_\gamma \mid \mathbf{x} \in \mathcal{X}) = \pi$ . As a result, if we can sample efficiently from  $\tilde{\pi}_\gamma$ , we can use rejection sampling to generate samples from  $\pi$  over  $\mathcal{X}$ . In other words, our plan will be to repeatedly generate samples from  $\mathcal{Y}$  according to  $\tilde{\pi}_\gamma$  and accept the first sample that happens to lie in  $\mathcal{X} \subset \mathcal{Y}$ . Our choice of  $\tilde{\pi}_\gamma$  to penalize lower-quality solutions is a standard technique (see, e.g., Porod (2024)), and  $\gamma$  is often referred to as an ‘inverse temperature’ parameter. We next specify  $M_\gamma$  with the desired stationary distribution  $\tilde{\pi}_\gamma$ .

We begin with a few definitions. For  $\mathbf{x}, \mathbf{x}' \in \mathcal{Y}$ , let  $h(\mathbf{x}, \mathbf{x}')$  denote the Hamming distance, or the number of entries in which  $\mathbf{x}$  and  $\mathbf{x}'$  disagree (i.e.,  $h(\mathbf{x}, \mathbf{x}') \triangleq \|\mathbf{x} - \mathbf{x}'\|_0$ ). For  $\mathbf{x}, \mathbf{x}' \in \mathcal{Y}$  such that  $h(\mathbf{x}, \mathbf{x}') = 1$ , let  $\delta(\mathbf{x}, \mathbf{x}')$  be the index  $i$  on which they disagree, so  $\mathbf{x}[i] \neq \mathbf{x}'[i]$ . Finally, let  $\Delta(\mathbf{x}, i) \triangleq \{\mathbf{x}_{i \leftarrow g} : \mathbf{V}\mathbf{x}_{i \leftarrow g} \preceq \mathbf{c}, g \in \mathbb{Z}_{\geq 0}\}$ . (Recall that  $\mathbf{x}_{i \leftarrow g}$  denotes replacing the  $i$ th entry of  $\mathbf{x}$  with the value  $g$ .) Intuitively,  $\Delta(\mathbf{x}, i)$  yields the set of feasible knapsack solutions obtained by changing the  $i$ th entry of  $\mathbf{x}$  to some nonnegative integer  $g$ . ( $g$  need not differ from the existing  $i$ th entry of  $\mathbf{x}$ .) Using a variant of Gibbs sampling, we define our Markov chain  $M_\gamma$  over the state space  $\mathcal{Y}$  to have transition probabilities

$$P_\gamma(\mathbf{x}, \mathbf{x}') \triangleq \begin{cases} \frac{\tilde{\pi}_\gamma(\mathbf{x}')}{2n \sum_{\mathbf{x}'' \in \Delta(\mathbf{x}, i)} \tilde{\pi}_\gamma(\mathbf{x}'')} & h(\mathbf{x}, \mathbf{x}') = 1 \wedge \delta(\mathbf{x}, \mathbf{x}') = i \\ 0 & h(\mathbf{x}, \mathbf{x}') > 1 \\ 1 - \sum_{\mathbf{x}'' \neq \mathbf{x}} P_\gamma(\mathbf{x}, \mathbf{x}'') & \mathbf{x} = \mathbf{x}' \end{cases}.$$

In what follows, we describe useful properties of  $P_\gamma$  and provide additional intuition behind it.

4.1.2. *Properties of the simple chain.* In order to show that running MCMC on  $M_\gamma$  yields an  $\varepsilon$ -approximate sample from  $\tilde{\pi}_\gamma$ , we will show that:

- $M_\gamma$  is lazy and irreducible.
- The stationary distribution of  $M_\gamma$  is  $\tilde{\pi}_\gamma$ .
- We can implement the transitions  $P_\gamma$  with an efficient algorithm.

First, observe that by construction,  $M_\gamma$  is “lazy,” meaning  $P_\gamma(\mathbf{x}, \mathbf{x}) \geq 1/2$ . This is because

$$\begin{aligned} P_\gamma(\mathbf{x}, \mathbf{x}) &= 1 - \sum_{\mathbf{x}' \neq \mathbf{x}} P_\gamma(\mathbf{x}, \mathbf{x}') \\ &= 1 - \sum_{i=1}^n \sum_{\mathbf{x}' \in \Delta(\mathbf{x}, i) \setminus \{\mathbf{x}\}} \frac{\tilde{\pi}_\gamma(\mathbf{x}')}{2n \sum_{\mathbf{x}'' \in \Delta(\mathbf{x}, i)} \tilde{\pi}_\gamma(\mathbf{x}'')} \\ &\geq 1 - \sum_{i=1}^n \sum_{\mathbf{x}' \in \Delta(\mathbf{x}, i)} \frac{\tilde{\pi}_\gamma(\mathbf{x}')}{2n \sum_{\mathbf{x}'' \in \Delta(\mathbf{x}, i)} \tilde{\pi}_\gamma(\mathbf{x}'')} \\ &= \frac{1}{2}. \end{aligned}$$

Moreover, any state is reachable from any other state since all states are connected to the empty solution  $\mathbf{x} = \mathbf{0}$ . The simple chain is thus both lazy and irreducible, which guarantees that  $M_\gamma$  is aperiodic and  $P_\gamma$  has nonnegative eigenvalues.

Next, we will show that  $\tilde{\pi}_\gamma$  is the stationary distribution of  $M_\gamma$ . For an aperiodic, irreducible Markov chain, a distribution  $\sigma$  that satisfies the so-called ‘detailed balance equations’ given by (4.5) is its unique stationary distribution (see, e.g., Levin et al. (2017, Cor. 1.17 and Prop. 1.20)): For all  $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$ ,

$$\sigma(\mathbf{x})P(\mathbf{x}, \mathbf{x}') = \sigma(\mathbf{x}')P(\mathbf{x}', \mathbf{x}). \quad (4.5)$$

In our case, this is clearly true for  $\mathbf{x} = \mathbf{x}'$  and when  $h(\mathbf{x}, \mathbf{x}') > 1$  (since  $P_\gamma(\mathbf{x}, \mathbf{x}') = P_\gamma(\mathbf{x}', \mathbf{x}) = 0$  in this case). For  $h(\mathbf{x}, \mathbf{x}') = 1$ , let  $i = \delta(\mathbf{x}, \mathbf{x}')$ . Note that  $\Delta(\mathbf{x}, i) = \Delta(\mathbf{x}', i)$  by construction. Therefore,

$$\tilde{\pi}_\gamma(\mathbf{x})P_\gamma(\mathbf{x}, \mathbf{x}') = \frac{\tilde{\pi}_\gamma(\mathbf{x})\tilde{\pi}_\gamma(\mathbf{x}')}{2n \sum_{\mathbf{x}'' \in \Delta(\mathbf{x}, i)} \tilde{\pi}_\gamma(\mathbf{x}'')} = \frac{\tilde{\pi}_\gamma(\mathbf{x})\tilde{\pi}_\gamma(\mathbf{x}')}{2n \sum_{\mathbf{x}'' \in \Delta(\mathbf{x}', i)} \tilde{\pi}_\gamma(\mathbf{x}'')} = \tilde{\pi}_\gamma(\mathbf{x}')P_\gamma(\mathbf{x}', \mathbf{x}).$$

Thus,  $\tilde{\pi}_\gamma$  is the stationary distribution of  $M_\gamma$ .

Finally, we must show that given  $\mathbf{x}$ , we can efficiently sample from  $P_\gamma(\mathbf{x}, \cdot)$ . The following algorithm does so: with probability  $1/2$ , remain at  $\mathbf{x}$ . With the remaining probability  $1/2$ , choose  $i \in [n]$  uniformly at random and enumerate the set  $\Delta(\mathbf{x}, i)$ . (In our instances, entries of  $\mathbf{c}$  are on the order of hundreds at most, meaning  $|\Delta(\mathbf{x}, i)|$  is relatively small.) Then, sample  $\mathbf{x}'$  proportional to  $\tilde{\pi}_\gamma(\cdot)$  from  $\Delta(\mathbf{x}, i)$  and transition to  $\mathbf{x}'$ .<sup>12</sup> We write this formally as Algorithm 2 in Appendix C.1.

**4.1.3. The number of samples needed for rejection sampling.** To generate a sample from  $\mathcal{X}$  (not just from  $\mathcal{Y} \supset \mathcal{X}$ ), given some hyperparameter  $t$ , we run  $M_\gamma$  for  $t$  iterations beginning with the empty solution as the start state (i.e.,  $\mathbf{x}_0 = \mathbf{0}$ ).<sup>13</sup> This yields a sample  $\mathbf{x} \sim \pi_{\gamma, t}$ , where we define  $\pi_{\gamma, t} \triangleq \pi_0 P_\gamma^t$  to be the distribution of the random walk after  $t$  steps. If  $\mathbf{x} \in \mathcal{X}$ , then we return  $\mathbf{x}$ ; if not, we reject the sample and repeat the process until we find some  $\mathbf{x} \in \mathcal{X}$ . See Algorithm 2 for details.

Define  $p_\gamma(t) \triangleq \pi_{\gamma, t}(\mathcal{X}) = \Pr_{\mathbf{x} \sim \pi_{\gamma, t}}[\mathbf{x} \in \mathcal{X}]$  to be the probability that a sample from  $\pi_{\gamma, t}$  is in  $\mathcal{X}$ . Because the number of samples needed for rejection sampling is geometrically distributed with parameter  $p_\gamma(t)$ , the total expected number of Markov chain iterations to

<sup>12</sup>For general  $\sigma$ , this entire process may require time proportional to  $|\Delta(\mathbf{x}, i)|$ . However, our implementation takes advantage of the structure of  $\tilde{\pi}_\gamma(\cdot)$  to do this in time proportional to  $\log(|\Delta(\mathbf{x}, i)|)$ .

<sup>13</sup>When we evaluate the simple chain, the starting state will not matter for the lower bounds we show.

produce a sample from  $\mathcal{X}$  is  $t/p_\gamma(t)$ . To fully specify our algorithm, we must choose  $t$  to be sufficiently large to produce an  $\varepsilon$ -approximate sample from  $\mathcal{X}$ .

We might think that it suffices to choose  $t \geq \tau_{\text{mix}}(\varepsilon; P_\gamma, \mathbf{x}_0)$ , since this implies that  $d_{\text{TV}}(\pi_{\gamma,t}, \tilde{\pi}_\gamma) \leq \varepsilon$ . However, this is insufficient: We want an  $\varepsilon$ -approximate sample from  $\mathcal{X}$ , not from  $\mathcal{Y}$ . Roughly speaking, a sample with approximation error  $\varepsilon$  for  $\tilde{\pi}_\gamma$  may have approximation error on the order  $\varepsilon/\tilde{\pi}_\gamma(\mathcal{X})$  on  $\mathcal{X}$ . We formalize this in Lemma 1, which we prove in Appendix C.2.

**Lemma 1.** Let  $\sigma$  and  $\sigma'$  be distributions defined on  $\mathcal{Y}$ . Let  $\sigma_{\mathcal{X}}$  and  $\sigma'_{\mathcal{X}}$  be their respective conditional distributions on  $\mathcal{X} \subset \mathcal{Y}$ , i.e., for  $\mathbf{x} \in \mathcal{X}$ ,  $\sigma_{\mathcal{X}}(\mathbf{x}) = \sigma(\mathbf{x})/\sigma(\mathcal{X})$ . If  $d_{\text{TV}}(\sigma, \sigma') \leq \varepsilon$ , then  $d_{\text{TV}}(\sigma_{\mathcal{X}}, \sigma'_{\mathcal{X}}) \leq 3\varepsilon/(2\sigma(\mathcal{X}))$ . For any  $\varepsilon$ , there exist instances for which this is tight to within a constant factor.

Let  $p_\gamma^* \triangleq \lim_{t \rightarrow \infty} p_\gamma(t) = \tilde{\pi}_\gamma(\mathcal{X})$ . By Lemma 1, to generate an  $\varepsilon$ -approximate sample from  $\mathcal{X}$ , it suffices to choose  $t \geq \tau_\gamma^*(\varepsilon) \triangleq \tau_{\text{mix}}(2p_\gamma^*\varepsilon/3; P_\gamma, \mathbf{x}_0)$ , since

$$d_{\text{TV}}(\pi_{\gamma,t} \mid \mathbf{x} \in \mathcal{X}, \tilde{\pi}_\gamma \mid \mathbf{x} \in \mathcal{X}) \leq \frac{3d_{\text{TV}}(\pi_{\gamma,t}, \tilde{\pi}_\gamma)}{2\tilde{\pi}_\gamma(\mathcal{X})} \leq \frac{3(2p_\gamma^*\varepsilon/3)}{2p_\gamma^*} = \varepsilon.$$

In other words, a  $2p_\gamma^*\varepsilon/3$ -approximate sample from  $\tilde{\pi}_\gamma$  over  $\mathcal{Y}$  yields an  $\varepsilon$ -approximate sample on from  $\pi$  over  $\mathcal{X}$ . This is tight to within a constant factor. Of course, we know neither  $p_\gamma^*$  nor  $\tau_{\text{mix}}(\cdot; P_\gamma, \mathbf{x}_0)$  *a priori*; we experimentally determine them for a subset of our instances in Section 4.1.4. Choosing  $t$  to be  $\tau_\gamma^*(\varepsilon)$ , the expected number of MCMC iterations needed to produce an  $\varepsilon$ -approximate sample from  $\mathcal{X}$  is

$$N_\gamma(\varepsilon) \triangleq \frac{\tau_\gamma^*(\varepsilon)}{p_\gamma(\tau_\gamma^*(\varepsilon))}, \quad (4.6)$$

where the numerator is the number of MCMC iterations needed per sample from  $\mathcal{Y}$  and the denominator is the probability that we accept a sample (i.e., the probability that it lies in  $\mathcal{X}$ ).

4.1.4. *Empirical results for the simple chain.* To evaluate the performance of Algorithm 2, our goal will be to determine  $N_\gamma(\varepsilon)$  for our instances. To do so, we will bound both  $\tau_\gamma^*(\varepsilon)$  and  $p_\gamma(\tau_\gamma^*(\varepsilon))$ . Observe that because  $d_{\text{TV}}(\pi_{\gamma, \tau_\gamma^*(\varepsilon)}, \tilde{\pi}_\gamma) \leq 2p_\gamma^*\varepsilon/3$ , by definition of total variation distance (4.1),

$$\left(1 - \frac{2\varepsilon}{3}\right) p_\gamma^* \leq p_\gamma(\tau_\gamma^*(\varepsilon)) \leq \left(1 + \frac{2\varepsilon}{3}\right) p_\gamma^*. \quad (4.7)$$

Let  $\tau_{\text{rel}, \gamma} \triangleq \frac{1}{1 - \lambda_2(P_\gamma)}$  be the relaxation time of  $M_\gamma$ . Combining (4.3) and (4.7) yields

$$\underline{N}_\gamma(\varepsilon) \triangleq \frac{(\tau_{\text{rel}, \gamma} - 1) \log\left(\frac{3}{4\varepsilon p_\gamma^*}\right)}{\left(1 + \frac{2\varepsilon}{3}\right) p_\gamma^*} \leq N_\gamma(\varepsilon) \leq \frac{\tau_{\text{rel}, \gamma} \log\left(\frac{3}{2\varepsilon p_\gamma^* \tilde{\pi}_\gamma(\mathbf{x}_0)}\right)}{\left(1 - \frac{2\varepsilon}{3}\right) p_\gamma^*} \triangleq \overline{N}_\gamma(\varepsilon). \quad (4.8)$$

We will use these tight lower and upper bounds for  $N_\gamma(\varepsilon)$  to characterize the performance of the simple chain. We choose  $\varepsilon = 1/(2e)$  by convention and write  $N_\gamma = N_\gamma(1/(2e))$ .

To compute  $\underline{N}_\gamma$  for a given Census block, we will need to compute  $\lambda_1(P_\gamma)$  and  $p_\gamma^*$ . Unfortunately, this requires  $\Omega(|\mathcal{Y}|)$  time and space just to write  $P_\gamma$ . This can be prohibitive for our instances since  $\mathcal{Y}$  can be exponentially (in  $n$ ) large. Our approach is as follows: we choose a subset of ‘small’ instances and compute the full transition matrix  $P_\gamma$  for a random

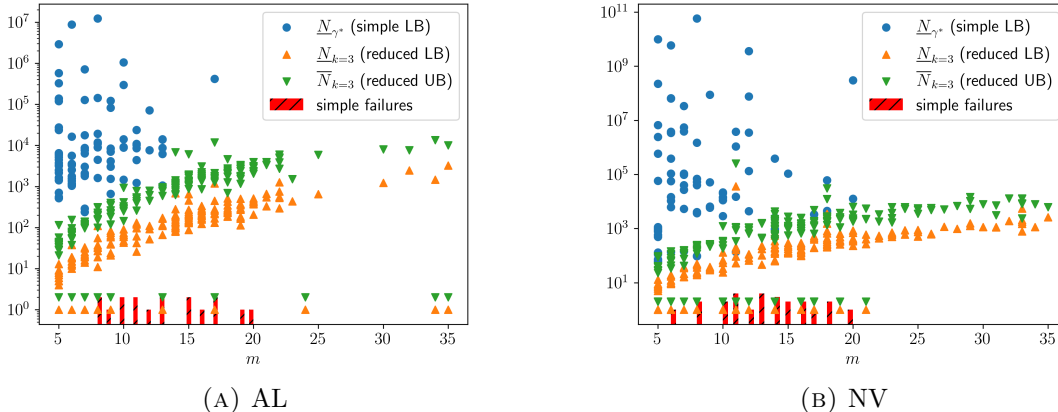


FIGURE 2. There are blocks for which the number of iterations required to generate an  $\varepsilon$ -approximate sample from  $\mathcal{X}$  using the simple chain is large. The reduced chain (described in Section 4.2) requires orders of magnitude fewer iterations in the worst case.

sample of those instances. We will show that in this sample, there exist instances for which  $\underline{N}_\gamma$  is prohibitively large, making Algorithm 2 impractical in our setting.

Recall that  $m$  is the number of households in a Census block. We sample 100 blocks where  $5 \leq m \leq 20$  and  $|\mathcal{X}| < 5000$  in two states: Alabama and Nevada (chosen arbitrarily). For each block, we seek to compute  $P_\gamma$  for  $\gamma \in \Gamma \triangleq \{0.0, 0.2, 0.4, 0.6, 0.8, 1.0, 1.2\}$ . For a significant number of blocks (indicated on Figure 2 in red bars; 25 in NV and 17 in AL), we fail to compute  $P_\gamma$  within a time limit of 8 hours. In Appendix E, we show that the globally optimal choice for  $\gamma$  on our sample is  $\approx 0.8$ .<sup>14</sup>

In Figure 2, we plot  $\underline{N}_{\gamma^*}$  for the blocks in our sample, where for each block,  $\underline{N}_{\gamma^*} \triangleq \min_{\gamma \in \Gamma} \underline{N}_\gamma$  using blue  $\bullet$  markers. (This is simply taking the “best” choice of  $\gamma$  for each block.) Observe that  $\underline{N}_{\gamma^*}$  can be quite large, even for small instances. There are blocks with fewer than 10 households for which  $\underline{N}_{\gamma^*}$  is on the order of  $10^{11}$ . In contrast, the ‘reduced’ approach we will develop in Section 4.2 appears to perform much better; the analogous lower and upper bounds for the reduced chain are small for all of our instances (Figure 2, orange  $\blacktriangle$  and green  $\blacktriangledown$  markers respectively).<sup>15</sup>

4.1.5. *Theoretical results for the simple chain.* Before we describe this improved approach, we briefly provide a theoretical characterization of how  $\underline{N}_\gamma$  varies with  $\gamma$  to gain some intuition as to why the simple chain performs poorly and how  $N_\gamma$  depends on  $\gamma$ . Intuitively, we will see that varying  $\gamma$  creates a trade-off between  $p_\gamma^*$  and  $\tau_{\text{rel},\gamma}$ . (Recall that  $\tau_{\text{rel},\gamma}$  is the relaxation time of  $M_\gamma$ .) Since  $N_\gamma \approx \tau_{\text{rel},\gamma}/p_\gamma^*$ , these opposing forces empirically make the simple chain impractical in our setting.

<sup>14</sup>Due to numerical instability in computing eigenvalues of  $P_\gamma$ , we sometimes underestimate  $\underline{N}_\gamma$  for larger values of  $\gamma$ .

<sup>15</sup>While we compare the number of MCMC iterations as opposed to computation time here, our experiments show that computation time per iteration is similar for the two Markov chains we consider. See Appendix E for details.

We begin by characterizing how  $p_\gamma^*$  varies with  $\gamma$ . All proofs are deferred to Appendix C.2.

**Lemma 2.**  $p_\gamma^*$  is monotonically increasing in  $\gamma$ , and

$$\lim_{\gamma \rightarrow -\infty} p_\gamma^* = 0 \quad \text{and} \quad \lim_{\gamma \rightarrow \infty} p_\gamma^* = 1.$$

As we might expect, large values of  $\gamma$  make it more likely that we sample an exact solution  $\mathbf{x} \in \mathcal{X}$ , since larger values of  $\gamma$  penalize inexact solutions more. Unfortunately, this comes at a cost, which we show in Lemma 3: in the limit, as  $\gamma$  increases,  $\tau_{\text{rel},\gamma}$  increases exponentially.<sup>16</sup>

**Lemma 3.** For the stationary distribution  $\tilde{\pi}_\gamma$  as defined in (4.4),  $\tau_{\text{rel},\gamma} = \Omega(\exp(\gamma \min_i \|\mathbf{v}_i\|_1))$ .

We prove Lemma 3 using a conductance argument and Cheeger’s inequality (Lawler and Sokal, 1988; Jerrum and Sinclair, 1988). Taken together, Lemmas 2 and 3 describe the trade-off in our choice of  $\gamma$ : for small values of  $\gamma$ , samples from  $\tilde{\pi}_\gamma$  rarely fall in  $\mathcal{X}$ , making rejection sampling inefficient. For large values of  $\gamma$ , the mixing time of the simple chain increases exponentially, requiring many iterations to generate each sample. These results tell us that the optimal choice of  $\gamma$  is finite.

**Lemma 4.** For every instance, there is some finite  $\gamma$  that minimizes  $N_\gamma$ .

Figure 3 visualizes this trade-off. Each line represents a different Census block. Red triangle  $\blacktriangle$  markers show low-precision estimates for the spectral gap  $1 - \lambda_2$ , which will typically mean we underestimate  $N_\gamma$ . Observe that each line appears to be quasiconvex, with a minimum in our range of choices for  $\gamma$ .

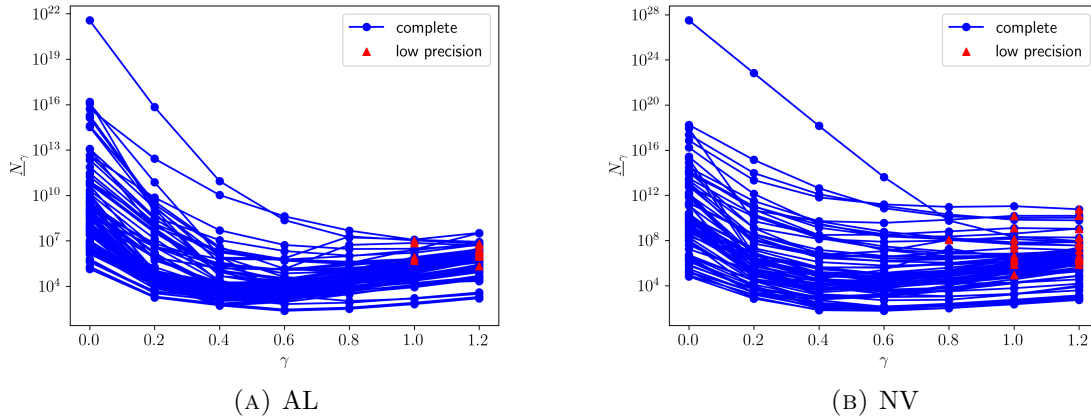


FIGURE 3.  $N_\gamma$  as a function of  $\gamma$ .

<sup>16</sup>A stronger version of this result would claim that, analogously to Lemma 2, mixing time increases monotonically with  $\gamma$ . While this may be true, proving monotonicity over the temperature parameter is notoriously difficult (see, e.g., Nacu (2003); Kargin (2011) and Levin et al. (2017, Ch. 26 Open Question 2)).

**4.2. The reduced chain.** Part of the reason why the simple chain requires many iterations in expectation to produce a valid sample from  $\mathcal{X}$  is its potentially large state space: because  $\mathcal{Y}$  can be much larger than  $\mathcal{X}$ , a random walk according to  $P_\gamma$  may yield samples from  $\mathcal{Y} \setminus \mathcal{X}$  with very high probability. This motivates our approach: the ‘reduced’ chain. We design a Markov chain  $M_k$  with state space  $\mathcal{X}$  instead of  $\mathcal{Y}$ .

**4.2.1. Defining the reduced chain.** The reduced chain is parameterized by an integer  $k \geq 2$ . Intuitively, given a solution  $\mathbf{x} \in \mathcal{X}$ , we randomly remove  $k$  elements from  $\mathbf{x}$  and replace them with another multiset of  $k$  elements (found via ILP) such that the resulting sum still exactly equals the constraint  $\mathbf{c}$ . For small  $k$ , this can be done fairly quickly. In our experiments, we use  $k \in \{2, 3, 4\}$ .

We again use a variant of Gibbs sampling to induce the desired stationary distribution. We define the following transition matrix  $P_k$  for our Markov chain  $M_k$ . Let  $\mathcal{A}(\mathbf{x}, \mathbf{x}')$  be the set of distinct pairs of multisets  $(\mathbf{z}, \mathbf{z}')$ , each of size  $k$ , such that we can transform  $\mathbf{x}$  into  $\mathbf{x}'$  by removing elements from  $\mathbf{z}$  and replacing them with elements from  $\mathbf{z}'$ . Formally, this is  $\mathcal{A}(\mathbf{x}, \mathbf{x}') \triangleq \{(\mathbf{z}, \mathbf{z}') \in \mathbb{Z}_{\geq 0}^n \times \mathbb{Z}_{\geq 0}^n : (\mathbf{V}\mathbf{z} = \mathbf{V}\mathbf{z}') \wedge (\mathbf{x} - \mathbf{z} = \mathbf{x}' - \mathbf{z}' \succeq \mathbf{0}) \wedge (\|\mathbf{z}\|_1 = \|\mathbf{z}'\|_1 = k)\}$ . Let  $\mathcal{Z}(\mathbf{z}) \triangleq \{\mathbf{z}' \in \mathbb{Z}_{\geq 0}^n : (\mathbf{V}\mathbf{z}' = \mathbf{V}\mathbf{z}) \wedge (\mathbf{z}' \succeq \mathbf{0}) \wedge \|\mathbf{z}'\|_1 = k\}$ .<sup>17</sup> Then, for  $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$ ,<sup>18</sup>

$$P_k(\mathbf{x}, \mathbf{x}') \triangleq \begin{cases} \frac{1}{2 \binom{m}{k}} \sum_{(\mathbf{z}, \mathbf{z}') \in \mathcal{A}(\mathbf{x}, \mathbf{x}')} \frac{\pi(\mathbf{x}')}{\sum_{\mathbf{z}'' \in \mathcal{Z}(\mathbf{z})} \pi(\mathbf{x} - \mathbf{z} + \mathbf{z}'')} & 2 \leq h(\mathbf{x}, \mathbf{x}') \leq 2k \\ 0 & h(\mathbf{x}, \mathbf{x}') > 2k \\ 1 - \sum_{\mathbf{x}'' \neq \mathbf{x}} P_k(\mathbf{x}, \mathbf{x}'') & \mathbf{x} = \mathbf{x}' \end{cases}.$$

Because  $\mathcal{Z}(\mathbf{z}) = \mathcal{Z}(\mathbf{z}')$  for  $(\mathbf{z}, \mathbf{z}') \in \mathcal{A}(\mathbf{x}, \mathbf{x}')$ ,  $\pi$  satisfies the detailed balance conditions  $\pi(\mathbf{x})P_k(\mathbf{x}, \mathbf{x}') = \pi(\mathbf{x}')P_k(\mathbf{x}', \mathbf{x})$ , making it a stationary distribution of  $M_k$ . Unfortunately,  $\pi$  is not necessarily the *unique* stationary distribution of  $M_k$ , which we will discuss below. Note that  $P_k$  is lazy by construction. In Algorithm 3 in Appendix D.1, we show that we can sample efficiently from  $P_k(\mathbf{x}, \cdot)$  as long as we can efficiently compute  $\mathcal{Z}(\mathbf{z})$  (via ILP), which empirically we can in our setting. The complexity of this computation depends on the problem dimension  $d$  and the parameter choice  $k$ , *not* on the total counts  $\mathbf{c}$ .

**Choosing an initial state  $\mathbf{x}_0$ .** To fully specify the algorithm, we must choose an initial state  $\mathbf{x}_0$ . In general, we could find an arbitrary  $\mathbf{x}_0 \in \mathcal{X}$  via ILP; however, in order to produce tighter bounds on performance, we aim to begin with a ‘better-than-average’  $\mathbf{x}_0$  (i.e., such that  $\pi(\mathbf{x}_0) \geq 1/|\mathcal{X}|$ ). To see why, observe that the upper bounds in (4.3) depend on  $\log(1/\pi(\mathbf{x}_0))$ , so larger values of  $\pi(\mathbf{x}_0)$  will reduce the number of MCMC iterations we need to produce a sample. We defer these details to Appendix D.1. Note that finding any  $\mathbf{x}_0 \in \mathcal{X}$  requires that we can solve the decision problem **MMS** via ILP despite its NP-hardness; empirically, we can for all of our instances.

<sup>17</sup>Equivalently, we could write  $\mathcal{Z}(\mathbf{z}) = \{\mathbf{z}' \in \mathbb{Z}_{\geq 0}^n : \exists \mathbf{x}' \in \mathcal{X} \text{ such that } (\mathbf{z}, \mathbf{z}') \in \mathcal{A}(\mathbf{x}, \mathbf{x}')\}$ .

<sup>18</sup>Because the elements in  $\mathcal{V}$  are unique, it is impossible to have  $h(\mathbf{x}, \mathbf{x}') = 1$  for  $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$ .

**The reduced chain is not necessarily irreducible.** In an irreducible Markov chain, any state is reachable from any other state, i.e., for any  $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$ , there exists  $t$  such that  $P_k^t(\mathbf{x}, \mathbf{x}') > 0$ . This is not necessarily the case for  $P_k$ ; it is possible to have disconnected components in the state space graph (see Figure 8 for an example). When this happens,  $\pi$  is not the only stationary distribution of  $M_k$ , and our algorithm will not sample from  $\pi$  as desired. We evaluate whether this is the case for our instances in what follows.

4.2.2. *Empirical results for the reduced chain.* We are interested in  $N_k(\varepsilon) \triangleq \tau_{\text{mix}}(\varepsilon; P_k, \mathbf{x}_0)$ , defined to be the number of Markov chain iterations required to generate an  $\varepsilon$ -approximate sample from  $\mathcal{X}$ . If  $M_k$  is not irreducible,  $N_k(\varepsilon)$  is undefined. As in Section 4.1.4, we derive bounds for  $N_k(\varepsilon)$  and evaluate these bounds for a random sample of instances. For irreducible  $M_k$ , let  $\tau_{\text{rel},k} \triangleq \frac{1}{1-\lambda_2(P_k)}$ . Using (4.3),

$$\underline{N}_k(\varepsilon) \triangleq (\tau_{\text{rel},k} - 1) \log \left( \frac{1}{2\varepsilon} \right) \leq N_k(\varepsilon) \leq \tau_{\text{rel},k} \log \left( \frac{1}{\varepsilon\pi(\mathbf{x}_0)} \right) \triangleq \overline{N}_k(\varepsilon). \quad (4.9)$$

We again choose  $\varepsilon = 1/(2e)$ . Because our choice of  $\mathbf{x}_0$  satisfies  $\pi(\mathbf{x}_0) \geq 1/|\mathcal{X}|$ , the lower and upper bounds are within an  $O(\log |\mathcal{X}|)$  factor of one another. In Lemma 5 in Section 4.2.3, we show that  $\log |\mathcal{X}| = O(m \log n)$ , meaning  $N_k = O(\tau_{\text{rel},k} \cdot m \log n)$  (see Corollary 1 below).

We re-use the same sample of 100 blocks per state that we used to evaluate the simple chain in Section 4.1.4. In addition, because  $|\mathcal{X}|$  can be much smaller than  $|\mathcal{Y}|$ , we are able to compute  $\lambda_2(P_k)$  for larger instances. We therefore sample another 100 blocks for each of AL and NV with  $14 \leq m \leq 35$  and  $|\mathcal{X}| < 5000$ .<sup>19</sup> For each block, we analyze  $k \in \{2, 3, 4\}$ .<sup>20</sup>




FIGURE 4. There is no clear relationship between  $|\mathcal{X}|$  and the central tendency of  $\underline{N}_k$  for  $k = 3$ .

Figure 2 shows  $\underline{N}_k$  and  $\overline{N}_k$  as a function of  $m$  for  $k = 3$ . The reduced chain requires orders of magnitude fewer iterations than the simple chain to generate an  $\varepsilon$ -approximate

<sup>19</sup>In Alabama and Nevada, 72% and 45% of blocks respectively have  $m \leq 35$ ,  $|\mathcal{X}| < 5000$ , and sufficiently complete data for our use (see Appendix A for details about excluded blocks).

<sup>20</sup>For some blocks, our computation times out (beyond a limit of 8 hours) for  $k = 4$ . Because our results suggest that  $k = 3$  suffices, these failures have little effect on our conclusions.

sample from  $\mathcal{X}$ . Interestingly,  $N_k$  appears to be much more consistent as a function of  $m$  than  $\underline{N}_\gamma$  is. Our sample includes two cutoffs (on  $m \leq 35$  and  $|\mathcal{X}| < 5000$ ), and we might worry that  $N_k$  will grow dramatically beyond these cutoffs. While we cannot extrapolate beyond them, Figure 2 provides evidence that  $N_k$  is not growing too quickly near the  $m = 35$  cutoff. Figure 4 shows that there is no clear relationship between the central tendency of  $N_k$  and  $|\mathcal{X}|$ , and some of the largest values of  $N_k$  come from blocks with small  $|\mathcal{X}|$ . Again, while we cannot extrapolate beyond the cutoffs in our sample, Figures 2 and 4 do not suggest that  $N_k$  will be prohibitively large outside of these cutoffs.




FIGURE 5.  $k = 3$  appears to be optimal because (1) there are blocks for which  $M_{k=2}$  is disconnected, and (2) there are blocks for which  $M_{k=2}$  is connected but has high mixing time.

Finally, we find that for  $k = 2$ ,  $M_k$  is not irreducible (i.e., the state space graph is disconnected) for 2 out of 200 blocks in AL and 3 out of 200 blocks in NV. On the other hand, all blocks in our sample yield irreducible Markov chains for  $k \geq 3$ . (By construction, if the reduced chain is irreducible for some  $k$ , it is irreducible for any  $k' > k$ .) Based on these results,  $k = 3$  appears to be a good choice in our setting. Figure 5 provides a more detailed comparison across  $k \in \{2, 3, 4\}$ . Purple plus  $\oplus$  markers denote blocks where the reduced chain is not irreducible for  $k = 2$ . Gold diamond  $\diamond$  markers denote blocks for which computation times out (beyond an 8-hour limit) for  $k = 4$ .

Finally, we comment on the computational cost of running the reduced chain in practice. Our results suggest that roughly  $10^5$  MCMC iterations suffices to generate an  $\varepsilon$ -approximate sample. In Appendix E, we provide experimental results showing that running this many iterations on an M1 MacBook Pro takes roughly 10 seconds per Census block. Because we process each Census block independently, this can be run efficiently in parallel on a standard computing cluster, requiring tens of days of computing for the states in question (see Appendix E for more details).

4.2.3. *Theoretical results for the reduced chain.* Here, we provide theoretical results for the reduced chain. In Appendix G, we construct families of pathological instances where (1) the reduced chain is disconnected for any  $k < |\mathcal{X}|$  (Example 2), and (2) the reduced chain is connected for  $k = 2$  but has exponentially high mixing time (Example 3). Despite this,

$M_k$  appears to perform quite well for our instances, even though we cannot obtain general subexponential bounds on  $N_k$  unless  $P = NP$ . Here, we attempt to provide intuition for why  $M_k$  performs well. First, we bound  $\log |\mathcal{X}|$ . All proofs are deferred to Appendix D.2.

**Lemma 5.**  $\log |\mathcal{X}| < \log |\mathcal{Y}| \leq m(1 + \log(n + 1))$ .

Recall that  $m$  is the number of households in a Census block. Given that  $m$  is reasonably small in our instances (typically on the order of tens, and at most a few thousand), Lemma 5 and (4.9) imply that  $N_k$  is primarily determined by the relaxation time  $\tau_{\text{rel},k} = 1/(1 - \lambda_2(P_k))$ .

**Corollary 1.** If  $\pi(\mathbf{x}_0) \geq \frac{1}{|\mathcal{X}|}$ , then  $N_k = \tau_{\text{rel},k} \cdot O(m \log n)$ .

We cannot in general provide bounds on  $\lambda_2(P_k)$ , but our empirical results suggest that it is sufficiently small (bounded away from 1) in our setting. We also cannot guarantee that  $M_k$  is irreducible; however, we next provide a general condition under which  $M_k$  is irreducible. This condition is empirically not met in our case, but as we will argue, it is *nearly* met, and might provide intuition as to why  $M_k$  appears to be irreducible in practice.

For attribute  $j \in [d]$ , let  $a_j^{\min}$  and  $a_j^{\max}$  denote the minimal and maximal value that any household vector  $\mathbf{v}_i$  takes in that position. Formally,  $a_j^{\min} \triangleq \min_{i \in [n]} \mathbf{v}_i[j]$ , and  $a_j^{\max}$  is defined analogously. Let  $\mathcal{A}_i \subset \mathbb{Z}_{\geq 0}$  be the set  $\{a_j^{\min}, a_j^{\min} + 1, \dots, a_j^{\max}\}$ . And finally, let  $\mathcal{A} \subset \mathbb{Z}_{\geq 0}^d$  be the set of integer vectors  $\mathcal{A}_1 \times \mathcal{A}_2 \times \dots \times \mathcal{A}_d$ . Informally,  $\mathcal{A}$  is the set of integer vectors contained in the hyperrectangle formed by the ranges of each attribute. By this definition  $\mathcal{V} \subseteq \mathcal{A}$ , since each household type vector must lie within  $\mathcal{A}$ . In what follows, we provide a condition under which  $M_k$  is irreducible.

**Lemma 6.** If  $\mathcal{V} = \mathcal{A}$ , then  $M_k$  is irreducible for  $k \geq 2$ .

Lemma 6 provides some insight as to why we might expect connectivity in  $M_k$ . Most of the attributes in our  $d$ -dimensional encoding represent counts of individuals within each household who belong to certain demographic groups (e.g., number of adults, number of Hispanic individuals, etc.). Given two solutions  $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$ , when should we be able to transition between them by only changing a small number of households (i.e.,  $k = 2$ ) at a time? Intuitively, if  $\mathbf{x}$  contains two households  $\mathbf{v}_i$  and  $\mathbf{v}_j$ , suppose we were to shuffle the individuals within those households, producing two new household types  $\mathbf{v}_{i'}$  and  $\mathbf{v}_{j'}$ . If  $\mathcal{V} = \mathcal{A}$ , both of these new household types will also be in  $\mathcal{V}$ . And because we're using the same set of individuals, we have not changed the overall demographic composition of the block. As a result, replacing  $\mathbf{v}_i, \mathbf{v}_j$  with  $\mathbf{v}_{i'}, \mathbf{v}_{j'}$  will still yield a valid solution  $\mathbf{x}'' \in \mathcal{X}$ . Doing this repeatedly allows us to find a path from  $\mathbf{x}$  to  $\mathbf{x}'$ , meaning  $M_k$  is irreducible.

Thus, if  $\mathcal{V}$  is very ‘dense,’ meaning all possible household types within some range have nonzero probability,  $M_k$  is irreducible. In practice,  $\mathcal{V}$  is not this dense; there exist some household types we never observe in the PUMS, meaning  $\mathcal{V} \subset \mathcal{A}$ . On the other hand, Lemma 6 provides a particularly strong notion of connectivity. In the proof, we show that when  $\mathcal{V} = \mathcal{A}$ , there are multiple disjoint paths between pairs  $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$ , not just the single path required for irreducibility. Future work may be able to relax this assumption to show irreducibility under weaker conditions, for example, when household types are missing with some probability  $p$  (see, e.g., Grimmer, 1999).

Finally, even when  $M_k$  is not irreducible, the reduced chain improves solution quality: a standard coupling argument shows that for any initial distribution  $\pi_0$ ,  $\|\pi_0 P_k - \pi\|_1 \leq \|\pi_0 - \pi\|_1$ , even when  $M_k$  is not irreducible. Moreover, we show that  $M_k$  converges to a distribution that is in a sense optimal given initial distribution  $\pi_0$ :

**Lemma 7.** Let  $\{\mathcal{C}_i\}_{i=1}^c$  be the connected components of the undirected graph induced by  $P_k$ . Given initial state  $\mathbb{X}_0 \sim \pi_0$ , Algorithm 3 converges to  $\pi^+(\pi_0) \triangleq \lim_{t \rightarrow \infty} \pi_0 P_k^t$  satisfying

$$\pi^+(\pi_0) \in \arg \min_{\pi': \pi'(C_i) = \pi_0(C_i) \forall i \in [c]} d_{\text{TV}}(\pi, \pi').$$

In other words,  $\pi^+(\pi_0)$  is as close as possible to  $\pi$  subject to the constraint that it puts the same probability mass in each connected component that  $\pi_0$  does.

4.2.4. *The reduced chain as a truncation of the simple chain.* We conclude this section with some intuition on the relationship between the simple chain and the reduced chain. We can think of the simple chain as operating on a lattice  $\mathcal{L}$  over  $\mathbb{Z}_{\geq 0}^n$ , where each  $\mathbf{x} \in \mathcal{Y}$  represents a lattice node. The constraint  $\mathbf{V}\mathbf{x} = \mathbf{c}$  is an intersection of hyperplanes  $\mathcal{I}$  over  $\mathbb{R}^n$ , and the set of exact solutions  $\mathcal{X}$  is the intersection of  $\mathcal{L}$  and  $\mathcal{I}$ . The simple chain corresponds to a random walk over the lattice that never goes ‘above’  $\mathcal{I}$ . The parameter  $\gamma$  effectively biases this random walk away from the origin by penalizing nodes exponentially in their distance from  $\mathcal{I}$ . In contrast, the reduced chain only operates on  $\mathcal{X} = \mathcal{L} \cap \mathcal{I}$ . But because transitions are defined by replacing  $k$  elements at a time, two states  $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$  have a transition under  $P_k$  if and only if there is a path between them in  $\mathcal{L}$  of length at most  $2k$ .

This intuition leads to a slightly different way to modify the simple chain: instead of penalizing inexact solutions with  $\exp(-\gamma \|\mathbf{V}\mathbf{x} - \mathbf{c}\|_1)$  as in (4.4), we could truncate the lattice with a threshold penalty like  $\mathbb{1}[\|\mathbf{V}\mathbf{x} - \mathbf{c}\|_1 \leq \omega]$  for some parameter  $\omega \in \mathbb{Z}_{\geq 0}$ . This would prevent the Markov chain from ever transitioning to states that are distance  $> \omega$  from  $\mathcal{I}$ , effectively restricting the state space of the simple chain to be

$$\mathcal{Y}_\omega \triangleq \{\mathbf{x} : \mathbf{V}\mathbf{x} \preceq \mathbf{c} \wedge \|\mathbf{V}\mathbf{x} - \mathbf{c}\|_1 \leq \omega\}.$$

We can write the reduced chain by simply setting  $\omega = 0$  in this definition (i.e.,  $\mathcal{Y}_{\omega=0} = \mathcal{X}$ ). If  $\mathcal{X}$  occupies sufficiently large mass within  $\mathcal{Y}_\omega$  for some  $\omega > 0$ , rejection sampling on this chain might lead to a practical algorithm in our setting. Of course, as is the case with the reduced chain, we would no longer be able to guarantee that this modification to the simple chain remains irreducible. We defer such investigations, as well as more complex methods to modify the state space, to future work.

## 5. A PRACTICAL ALGORITHM BASED ON THE REDUCED CHAIN

We conclude with a description of how we implement our algorithm in practice. Because we cannot guarantee that the reduced chain is irreducible, we draw inspiration from Lemma 7 to choose an initial distribution  $\pi_0$  over initial states to reduce total variation distance to the desired distribution even when the reduced chain is not irreducible. Our algorithm has three main hyperparameters:  $k$ ,  $t$ , and  $N$ . As before,  $k$  gives the number of elements we remove and replace at each iteration of the reduced chain. Larger values of  $k$  increase connectivity at the expense of increased computational cost. The number of MCMC iterations we run is  $t$ , where in our experiments (Section 4.2.2) we found that  $t$  on the order of  $10^5$  seems appropriate.

$N$  governs our initial state: instead of choosing a single initial state, we use ILP to enumerate the best  $N$  solutions according to a linear approximation to  $\pi$  (see Appendix D.1 for details). If a particular block has  $|\mathcal{X}_b| < N$ , we will enumerate all of them, meaning we can simply sample according to  $\pi$  with no need for MCMC. If, however,  $|\mathcal{X}_b| \geq N$ , we

will enumerate some  $\mathcal{S} \subseteq \mathcal{X}$  such that  $|\mathcal{S}| = N$ . We choose our initial distribution to be  $\pi_0 = (\pi \mid \mathbf{x} \in \mathcal{S})$ , meaning we sample a random  $\mathbf{x}_0$  from  $\mathcal{S}$  according to  $\pi$  and use this as our starting state for MCMC. For a formal description of this algorithm, see Algorithm 4 in Appendix H. This choice of  $\pi_0$  has two benefits:

- (1) If  $M_k$  is not irreducible, the stationary distribution we will converge to is  $\pi^+(\pi_0)$  as defined in Lemma 7.
- (2) If  $M_k$  is irreducible, this choice of  $\pi_0$  provides slightly better mixing time guarantees than starting with a single deterministic  $\mathbf{x}_0$ :

**Lemma 8.** If  $M_k$  is irreducible, then Algorithm 4 produces an  $\varepsilon$ -approximate sample from  $\pi$  after at most

$$\tau_{\text{mix}}(\varepsilon; P_k, \mathbb{X}) \leq \tau_{\text{rel},k} \log \left( \frac{1}{2\varepsilon \sqrt{\pi(\mathcal{S})}} \right)$$

Markov chain iterations.

(Recall that if  $|\mathcal{X}_b| < N$ , then Algorithm 4 does not use MCMC at all.) Thus, the number of MCMC iterations needed reduces as  $\mathcal{S}$  captures more probability mass. We prove Lemma 8 in Appendix H.

## 6. EVALUATING AND ENFORCING REPRESENTATIVENESS

We now return to our original goal: generating a dataset such that the marginal distribution over households approximates the PUMS distribution. The PUMS data are very granular, consisting of detailed information on each individual in each household. In contrast, the synthetic microdata we produce are much coarser. Consistent with the SF1 data release, we report the following pieces of information for each household in our microdata:<sup>21</sup>

- Number of persons of each of the 7 race categories,
- Number of Hispanic persons,
- Number of adults.

Let  $\mathcal{L} \subset \mathbb{Z}_{\geq 0}^9$  be the set of all low-dimensional household “types” defined this way. For each  $i \in [n]$ , let  $\ell_i$  be the low-dimensional projection of  $\mathbf{v}_i$  onto these features. Note that this means for  $i \neq j$ , it is possible to have  $\ell_i = \ell_j$  if  $\mathbf{v}_i$  and  $\mathbf{v}_j$  agree on these 9 attributes. Our goal will be to compare the distribution over types induced by our microdata to the distribution of types found in the PUMS.

To do so, we introduce additional notation. Let  $\mathcal{B}$  be the set of blocks in a given state. For a block  $b \in \mathcal{B}$ , let  $\pi_b$  be the distribution  $\pi$  over  $\mathcal{X}_b$ , the set of solutions for block  $b$ , as defined in (2.1). Let  $m_b$  be the number of households in block  $b$ . Then, let  $q_\ell$  be the expected frequency of households of type  $\ell$  in our microdata. For  $\ell \in \mathcal{L}$ ,

$$q_\ell \triangleq \frac{1}{|\mathcal{B}|} \sum_{b \in \mathcal{B}} \frac{1}{m_b} \sum_{\substack{i \in [m] \\ \ell_i = \ell}} \mathbb{E}_{\mathbb{X} \sim \pi_b} [\mathbb{X}[i]]. \tag{6.1}$$

---

<sup>21</sup>We could instead provide more granular information like household status and householder race. This would lead to an increase in total variation distance. We choose to omit them here because downstream analyses we seek to support (including running privacy-preserving algorithms) only use these attributes.

We can compare this directly with  $p_\ell$ , defined to be the frequency of households of type  $\ell$  in the PUMS:

$$p_\ell \triangleq \sum_{\substack{i \in [n] \\ \ell_i = \ell}} \Pr_{\mathcal{D}}[\mathbf{v}_i].$$

Ideally, we would have  $q_\ell = p_\ell$  for all  $\ell \in \mathcal{L}$ , meaning our methods are unbiased with respect to overall household distribution. However, this will in general not be the case for a variety of reasons. First, as detailed further in Appendix A, data may be incomplete, meaning we cannot always sample from  $\mathcal{X}_b$  for every block  $b$ . (In such cases, we coarsen the constraints we enforce to ensure we still get reasonable results.) More subtly, even with perfect data, our induced distribution is not unbiased. Consider the following example with three blocks:

**Example 1** (Empirical frequencies may be biased). Block A contains one household with two Black persons (type 1); block B contains two households, each of which contains one white person and one Black person (type 2); and block C contains one household with two white persons (type 3).

Overall, the frequency of the household types 1, 2, and 3 are  $1/4$ ,  $1/2$ , and  $1/4$  respectively. Given these frequencies, our approach will correctly reconstruct blocks A and C, since there is a unique solution to each of these blocks. Block A can only be reconstructed as a single household of type 1, and block C requires a single household of type 3. However, block B admits multiple possible solutions: it can be reconstructed either as one household of type 1 and one household of type 3, or as two households of type 2. Under the distribution  $\pi$  as defined in (2.1), these will each be sampled with probabilities  $1/3$  and  $2/3$  respectively. As a result, despite the fact that  $(p_1, p_2, p_3) = (1/4, 1/2, 1/4)$ , our induced frequencies will be  $(q_1, q_2, q_3) = (1/3, 1/3, 1/3)$ . Thus, in general,  $q_\ell \neq p_\ell$ .

**Empirical differences in the household distribution.** Given that our induced distribution will not be unbiased in general, we empirically evaluate just how close it is. To do so, we use Algorithm 4 to sample a synthetic dataset given by  $\{\mathbf{x}_b\}_{b \in \mathcal{B}}$  across an entire state. Note that there is an additional source of error here: we do not in general approximately sample from  $\pi_b$  because when we are unable to completely enumerate  $\mathcal{X}_b$  (which is precisely when we need the MCMC methods developed earlier) and the reduced chain is not connected, we instead sample from  $\pi^+(\pi \mid \mathbf{x} \in \mathcal{S})$  as defined in Lemma 7.

We will compare the PUMS distribution over types  $p_\ell$  to the empirical frequencies  $\hat{q}_\ell$  over the household types, defined analogously to (6.1):

$$\hat{q}_\ell \triangleq \frac{1}{|\mathcal{B}|} \sum_{b \in \mathcal{B}} \frac{1}{m_b} \sum_{\substack{i \in [n] \\ \ell_i = \ell}} \mathbf{x}_b[i].$$

In what follows, we only report statistics involving  $\hat{q}$  based on a single run of Algorithm 4. Due to the computational cost involved, we are unable to provide confidence intervals. Let  $\mathcal{P}$  be the distribution over  $\mathcal{L}$  with probabilities given by  $p_\ell$ , and let  $\hat{\mathcal{Q}}$  be the distribution with probabilities  $\hat{q}_\ell$ . We can now evaluate the distance between our empirical distribution  $\mathcal{P}$  and  $\hat{\mathcal{Q}}$  using the total variation distance  $d_{\text{TV}}(\mathcal{P}, \hat{\mathcal{Q}}) = \frac{1}{2} \sum_{\ell \in \mathcal{L}} |p_\ell - \hat{q}_\ell|$ .

For Alabama, we find that  $d_{\text{TV}}(\mathcal{P}, \hat{\mathcal{Q}}) = 0.126$ . However, there is a source of error that is beyond our control and, when accounted for, reduces this gap significantly: the PUMS

distribution is inconsistent with the distribution reported by SF1. For example, according to SF1, the proportion of single-person households in Alabama is 0.27, whereas according to the PUMS, it is 0.32. Because we are constrained to exactly match the number of single-person households reported by SF1, we are guaranteed to incur at least this error relative to the PUMS.

To account for this, in our analysis, we compare against a modified distribution  $\tilde{\mathcal{P}}$  derived by reweighting  $\mathcal{P}$  to match the SF1 data on characteristics that are completely determined by SF1. By doing so, we attempt to control for discrepancies between the PUMS distribution and SF1 that are outside our control. Specifically, we reweight  $\mathcal{P}$  to match SF1 on both overall household size (up to size 7) and ethnicity among single-person households (where the proportion of such households is exactly determined by SF1 table P28H), producing  $\tilde{\mathcal{P}}$ .<sup>22</sup> For example, for a household  $\ell \in \mathcal{L}$  containing two individuals, we define

$$\Pr_{\ell \sim \tilde{\mathcal{P}}}[\ell] \triangleq \Pr_{\ell \sim \mathcal{P}}[\ell] \cdot \frac{\Pr_{\ell' \sim \hat{\mathcal{Q}}}[\ell' \text{ contains two individuals}]}{\Pr_{\ell' \sim \mathcal{P}}[\ell' \text{ contains two individuals}]}.$$

By construction,  $\tilde{\mathcal{P}}$  will match the SF1 distribution on our chosen set of characteristics  $\mathcal{C}$ . In our case  $\mathcal{C} \triangleq \{1H, 1N, 2, 3, 4, 5, 6, 7+\}$ , where  $1H$  and  $1N$  represent single-person households containing a Hispanic individual and a non-Hispanic individual respectively. These characteristics form a partition of the set of all households, and SF1 completely determines  $\Pr_{\ell \sim \hat{\mathcal{Q}}}[\ell \text{ is of type } c]$  for all  $c \in \mathcal{C}$ .<sup>23</sup> Thus, for any  $c \in \mathcal{C}$ ,

$$\begin{aligned} \Pr_{\ell \sim \tilde{\mathcal{P}}}[\ell \text{ is of type } c] &= \sum_{\ell: \ell \text{ is of type } c} \Pr_{\ell \sim \tilde{\mathcal{P}}}[\ell] \\ &= \sum_{\ell: \ell \text{ is of type } c} \Pr_{\ell \sim \mathcal{P}}[\ell] \cdot \frac{\Pr_{\ell' \sim \hat{\mathcal{Q}}}[\ell' \text{ is of type } c]}{\Pr_{\ell' \sim \mathcal{P}}[\ell' \text{ is of type } c]} \\ &= \frac{\Pr_{\ell' \sim \hat{\mathcal{Q}}}[\ell' \text{ is of type } c]}{\Pr_{\ell' \sim \mathcal{P}}[\ell' \text{ is of type } c]} \sum_{\ell: \ell \text{ is of type } c} \Pr_{\ell \sim \mathcal{P}}[\ell] \\ &= \frac{\Pr_{\ell' \sim \hat{\mathcal{Q}}}[\ell' \text{ is of type } c]}{\Pr_{\ell' \sim \mathcal{P}}[\ell' \text{ is of type } c]} \cdot \Pr_{\ell \sim \mathcal{P}}[\ell \text{ is of type } c] \\ &= \Pr_{\ell' \sim \hat{\mathcal{Q}}}[\ell' \text{ is of type } c]. \end{aligned}$$

Again, because SF1 completely determines  $\Pr_{\ell \sim \hat{\mathcal{Q}}}[\ell \text{ is of type } c]$ , adjusting  $\mathcal{P}$  in this way effectively accounts for the fact that our two data sources do not exactly agree—that is,  $\Pr_{\ell \sim \mathcal{P}}[\ell \text{ is of type } c]$  is inconsistent with SF1.

Under this adjustment, again in Alabama,  $d_{\text{TV}}(\tilde{\mathcal{P}}, \hat{\mathcal{Q}}) = 0.097$ . Finally, for data incompleteness reasons detailed in Appendix A,  $\mathcal{X}_b$  is empty for some blocks, and we can only solve a coarsened version of our original problem. If we only consider the distribution  $\tilde{\mathcal{Q}}$  that excludes these blocks, we find  $d_{\text{TV}}(\tilde{\mathcal{P}}, \tilde{\mathcal{Q}}) = 0.083$ .

It's not immediately obvious how to interpret these TVD figures. On the one hand, we're only measuring TVD at the state level; TVD at the block level would necessarily be larger,

<sup>22</sup>Note that this does not affect our original data generation process since the constraints exactly specify the number of households of each size.

<sup>23</sup>For a small fraction of blocks with incomplete information, this is not strictly true. See Appendix A for details.

though we cannot measure it without access to complete microdata. On the other hand, we know that the PUMS sample is biased, meaning there is some irreducible TVD between it and any microdata consistent with SF1. While we have partially accounted for this by adjusting the PUMS on household size (going from  $\hat{Q}$  to  $\tilde{Q}$ ), there may be irreducible TVD along other dimensions. (For example, if the expected number of persons of a particular race given by the PUMS is different from the number of persons of that race reported by SF1 in the state, this would be another source of irreducible TVD.) While we cannot conclusively resolve the sources of error at play, we will find some evidence below that at least some of this error is due to our methods, not irreducible TVD.

A similar trend holds for NV, summarized in Table 1 (labeled “no reweighting”), though the total variation distance is substantially larger. This is consistent with a broader trend: Nevada requires more MCMC iterations per block, requires MCMC more often (i.e.,  $|\mathcal{X}_b|$  is more often greater than  $N$ ; see Figure 1), and requires more computation time overall relative to Alabama. We speculate that this is because Nevada tends to have substantially more people per block (an average of 75.8, as opposed to 35.3 for Alabama) as well as a larger support for the household distribution  $\mathcal{D}$  (i.e., more distinct household types). Replicating our work across more states will likely lead to more insight here.

State	$\lambda$	$d_{\text{TV}}(\mathcal{P}, \hat{Q})$	$d_{\text{TV}}(\tilde{\mathcal{P}}, \hat{Q})$	$d_{\text{TV}}(\tilde{\mathcal{P}}, \tilde{Q})$
AL	no reweighting	0.126	0.097	0.083
	$10^{-3}$	0.091	0.077	0.069
	$10^{-5}$	0.109	0.080	0.067
NV	no reweighting	0.229	0.219	0.205
	$10^{-3}$	0.217	0.212	0.198
	$10^{-5}$	0.222	0.217	0.201

TABLE 1. Total variation distance between empirical distribution and statewide PUMS.

**Improving representativeness via reweighting.** When the PUMS distribution differs substantially from our empirical frequencies, we can adjust the inputs to our algorithm to try to close the gap. Algorithm 4 depends on the input distribution  $\mathcal{D}$ , and we can *reweight*  $\mathcal{D}$  (drawing some inspiration from Wang et al. (2024)) to produce a new distribution  $\mathcal{D}_\lambda$  such that, when we run Algorithm 4, the resulting empirical frequencies  $\hat{Q}$  are closer to the original PUMS distribution  $\mathcal{P}$ . We provide a heuristic in what follows, showing that our reweighting scheme reduces total variation distance to  $\tilde{\mathcal{P}}$ .

Suppose we run Algorithm 4 once, resulting in empirical household type frequencies  $\hat{q}_\ell$ . Consider the distribution  $\mathcal{D}_\lambda$  defined as follows for some smoothing parameter  $\lambda$ :

$$\Pr_{\mathcal{D}_\lambda}[\mathbf{v}_i] \propto \Pr_{\mathcal{D}}[\mathbf{v}_i] \cdot \left( \frac{p_{\ell_i} + \lambda}{\hat{q}_{\ell_i} + \lambda} \right).$$

We might expect that if we run Algorithm 4 on  $\mathcal{D}_\lambda$  instead of  $\mathcal{D}$ , we will end up with a lower total variation distance in our household type distribution. Intuitively, this is because  $\mathcal{D}_\lambda$  attempts to ‘correct’ for the bias induced by Algorithm 4, boosting the frequency of households that appear infrequently in  $\hat{Q}$  relative to  $\mathcal{P}$  and reducing the frequencies of

households that are overrepresented. Table 1 shows that reweighting  $\mathcal{D}$  decreases total variation distance, though it does not bring it down to 0. It is possible that additional adjustments could reduce TVD further.

Given Example 1, we might wonder whether multiracial households are likely to be underrepresented in our data. This is indeed the case in both states we analyze, where the frequency of multiracial households drops from 3.36% (AL PUMS) to 2.52% (AL synthetic data) and from 11.59% (NV PUMS) to 6.27% (NV synthetic data). At least facially, this can be remedied by artificially boosting the frequencies of multiracial households in  $\mathcal{D}_\lambda$ . As a crude example, if we let  $\lambda = 10^{-5}$  and double the frequency of each multiracial household in  $\mathcal{D}_\lambda$  (before normalizing it), we raise the frequency of multiracial households in  $\hat{\mathcal{Q}}$  to 10.59% in Nevada, much closer to the figure of 11.59% found in the NV PUMS. Again, further work is needed to determine whether this has broader effects on the realism of our synthetic microdata. Practitioners interested in a particular aspect of the distribution (e.g., multiracial households or single-parent households) should explicitly verify that the distribution of those characteristics in the synthetic microdata matches the PUMS distribution.

## 7. DISCUSSION AND LIMITATIONS

We have presented algorithms to generate synthetic data and demonstrated that these algorithms are efficient in the Census context. While these data “look like” microdata, it is important to recognize that no sample from this distribution should be treated as ground truth; analyses should look for consistent findings over multiple samples from this distribution. The distribution chosen here is principled, but future work could test whether alternative distributions produce more realistic data, and whether this paradigm of synthetic data generation via combinatorial optimization is appropriate in contexts beyond the Census.

Finally, we comment on potential uses for our synthetic microdata. One clear use case is researchers who seek to **study the properties of disclosure avoidance systems**. For example, a researcher could generate synthetic datasets, run a disclosure avoidance method on it (e.g., TopDown (Abowd, 2018b)), and perform fine-grained analyses of its privacy and accuracy properties. A growing line of work seeks to do just this (e.g., Cohen et al., 2021; Christ et al., 2022; Kenny et al., 2021; Bailie et al., 2023), and the methods presented here will enable a more systematic approach. Moreover, scientists who depend on Census data may be interested in **biases induced by disclosure avoidance algorithms** (e.g. Mueller and Santos-Lozada, 2022; Ruggles et al., 2019; Hauer and Santos-Lozada, 2021; Santos-Lozada et al., 2020; Winkler et al., 2021); a potential use for our methods is to estimate and perhaps correct for these biases. If a researcher is interested in a particular statistic measured using Census data, they could estimate the effect that a disclosure avoidance tool will have on that statistic over the distribution of synthetic data we generate. As noted in Section 1, we suggest that all analyses be performed across multiple samples of synthetic data. Used appropriately, the methods presented here can be a valuable tool for researchers to understand the properties of privacy-preserving algorithms.

Our methods could easily be extended to other population synthesis settings, potentially with more complex population models (see, e.g., Chapuis et al., 2022). This might introduce additional complexity: Unlike our PUMS-derived distribution, population models in general can have exponentially large support (recall that  $n$  is the size of the support of  $\mathcal{D}$ ). If, for example, every possible household exists with nonzero probability, writing down the matrix  $\mathbf{V}$  (let alone solving an integer program with it) can be computationally infeasible. To

get around this, one could take  $K$  samples from the population distribution  $\mathcal{D}$  to get an empirical distribution  $\hat{\mathcal{D}}_N$  with bounded support. We could even make this more efficient by ruling out samples that are incompatible with the SF1 counts  $\mathbf{c}$ . (For example, if a block has no 5-person households, we would not need to include 5-person households when sampling  $\hat{\mathcal{D}}_N$ .) Doing so independently for each block would still yield an exact match to the aggregate counts, though we may need to make  $N$  sufficiently large to ensure that a valid solution exists.

Our work has important limitations, both technical and conceptual. Technically, our methods are not guaranteed to converge to the desired stationary distribution, a limitation we are unlikely to overcome due to the NP-hardness of the underlying sampling problem. As noted in Example 1 our choice of distribution  $\pi$  does not lead to an unbiased set of households. The heuristics we present in Section 6 reduce but do not eliminate this bias, and researchers should take care that synthetic data match the PUMS data on characteristics important to their analyses. We are also limited by the fact that the PUMS is only a sample of households, meaning low-frequency household types may never appear.

Future work could extend our methods in a number of different ways. Our particular problem formulation sought to match particular statistics on race, ethnicity, gender, and household size. SF1 data contain far more information than these, however, and in theory one could seek to match other characteristics as well. For instance, researchers interested in studying older populations may wish to match statistics from table P25 (“Households by Presence of People 65 Years and Over, Household Size, and Household Type”). Of course, the more constraints we add, the larger the dimension  $d$  of the problem. This could increase the computational cost of finding solutions or overconstrain the problem, making it impossible to solve exactly. A better understanding of this trade-off is needed to take advantage of all of the data provided by SF1. Beyond the Decennial Census, research and public policy often rely on other Census products like the American Community Survey. Extending our methods to these other data releases may be a fruitful avenue for future research. Finally, our formulation treats the state-wide household distribution as invariant conditional on the block-level statistics. If, however, the researcher has more information about geographic variation that the block-level statistics do not encode (for example, that a particular region has more multiracial households than another), they could in principle encode that information by using a different distribution over households  $\mathcal{D}$  for each block.

Ultimately, there is plenty of room to customize the methods we have presented here for any particular use case. Our results have shown that it is possible to efficiently generate samples from a plausible distribution over synthetic Census microdata. Our hope is that researchers will be able to tailor our approach to fit their needs, enabling more comprehensive research into privacy-preserving methods and their broader impacts on downstream analyses.

**Acknowledgements.** The authors thank Parikshit Gopalan for helpful conversations. This work was supported in part by a Google Research Scholar award and the Center for Research on Computation and Society (CRCS), Harvard.

#### REFERENCES

- K. Aardal, C. A. Hurkens, and A. K. Lenstra. Solving a system of linear Diophantine equations with lower and upper bounds on the variables. *Mathematics of operations research*, 25(3):427–442, 2000. <https://doi.org/10.1287/moor.25.3.427.12219>.

- J. M. Abowd. Staring-down the database reconstruction theorem. In *Joint statistical meetings, Vancouver, BC*, volume 234, 2018a. <https://www.census.gov/content/dam/Census/newsroom/press-kits/2018/jsm/jsm-presentation-database-reconstruction.pdf>.
- J. M. Abowd. The US Census Bureau adopts differential privacy. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 2867–2867, 2018b. <https://doi.org/10.1145/3219819.3226070>.
- J. M. Abowd. Declaration of John Abowd. State of Alabama et al. v. United States Department of Commerce et al., 2021. <https://www.documentcloud.org/documents/21018464-fair-lines-america-foundation-july-26-2021-declaration-of-john-m-abowd>.
- J. M. Abowd, R. Ashmead, R. Cumings-Menon, S. Garfinkel, M. Heineck, C. Heiss, R. Johns, D. Kifer, P. Leclerc, A. Machanavajjhala, et al. The 2020 census disclosure avoidance system TopDown algorithm. *Harvard Data Science Review*, 2, 2022. <https://doi.org/10.1162/99608f92.529e3cb9>.
- O. Achou. The number of feasible solutions to a knapsack problem. *SIAM Journal on Applied Mathematics*, 27(4):606–610, 1974. <https://doi.org/10.1137/0127050>.
- G. Albiston, T. Osman, and D. Brown. A neural network approach for population synthesis. *SIMULATION*, page 00375497241233597, 2024. <https://doi.org/10.1177/00375497241233597>.
- J. Bailie, R. Gong, and X.-L. Meng. Can swapping be differentially private? A refreshment stirred, not shaken. *NBER working paper*, 2023. <https://arxiv.org/pdf/2504.15246>.
- D. Ballas, G. Clarke, D. Dorling, H. Eyre, B. Thomas, and D. Rossiter. Simbritain: a spatial microsimulation approach to population dynamics. *Population, Space and Place*, 11(1): 13–34, 2005. <https://doi.org/10.1002/psp.351>.
- R. J. Beckman, K. A. Baggerly, and M. D. McKay. Creating synthetic baseline populations. *Transportation Research Part A: Policy and Practice*, 30(6):415–429, 1996. [https://doi.org/10.1016/0965-8564\(96\)00004-3](https://doi.org/10.1016/0965-8564(96)00004-3).
- A. G. Begeed-Dov. Lower and upper bounds for the number of lattice points in a simplex. *SIAM Journal on Applied Mathematics*, 22(1):106–108, 1972. <https://doi.org/10.1137/0122012>.
- M. Birkin and M. Clarke. Synthesis—a synthetic spatial information system for urban and regional analysis: methods and examples. *Environment and planning A*, 20(12):1645–1671, 1988. <https://doi.org/10.1068/a201645>.
- W. Blankinship. Algorithm 288: solution of simultaneous linear Diophantine equations [f4]. *Communications of the ACM*, 9(7):514, 1966. <https://doi.org/10.1145/365719.365971>.
- J. Bossek, A. Neumann, and F. Neumann. Exact counting and sampling of optima for the knapsack problem. In *International Conference on Learning and Intelligent Optimization*, pages 40–54. Springer, 2021. [https://doi.org/10.1007/978-3-030-92121-7\\_4](https://doi.org/10.1007/978-3-030-92121-7_4).
- K. Brádler. On the number of nonnegative solutions of a system of linear Diophantine equations. *arXiv preprint arXiv:1610.06387*, 2016. <https://arxiv.org/pdf/1610.06387>.
- G. H. Bradley. Algorithms for Hermite and Smith normal matrices and linear Diophantine equations. *Mathematics of Computation*, 25(116):897–907, 1971. <https://doi.org/10.1090/S0025-5718-1971-0301909-X>.
- A. V. Cabot. An enumeration algorithm for knapsack problems. *Operations Research*, 18(2): 306–311, 1970. <https://doi.org/10.1287/opre.18.2.306>.

- D. Casati, K. Müller, P. J. Fourie, A. Erath, and K. W. Axhausen. Synthetic population generation by combining a hierarchical, simulation-based approach with reweighting by generalized raking. *Transportation Research Record*, 2493(1):107–116, 2015. <https://doi.org/10.3141/2493-12>.
- K. Chapuis, P. Taillandier, and A. Drogoul. Generation of synthetic populations in social simulations: a review of methods and practices. *Journal of Artificial Societies and Social Simulation*, 25(2), 2022. <https://doi.org/10.18564/jasss.4762>.
- T.-W. J. Chou and G. E. Collins. Algorithms for the solution of systems of linear Diophantine equations. *SIAM Journal on computing*, 11(4):687–708, 1982. <https://doi.org/10.1137/0211057>.
- M. Christ, S. Radway, and S. M. Bellovin. Differential privacy and swapping: Examining de-identification’s impact on minority representation and privacy preservation in the US Census. In *2022 IEEE Symposium on Security and Privacy (SP)*, pages 457–472. IEEE, 2022. <https://doi.org/10.1109/SP46214.2022.9833668>.
- A. Cohen, M. Duchin, J. Matthews, and B. Suwal. Census TopDown: The impacts of differential privacy on redistricting. In *2nd Symposium on Foundations of Responsible Computing (FORC)*, 2021. <https://doi.org/10.4230/LIPIcs.FORC.2021.5>.
- R. Creecy. The feasibility of creating a fully synthetic decennial census micro-data file. 2009. <https://ecommons.cornell.edu/server/api/core/bitstreams/71719069-2b68-4995-a88c-92aab7d989f1/content>.
- W. E. Deming and F. F. Stephan. On a least squares adjustment of a sampled frequency table when the expected marginal totals are known. *The Annals of Mathematical Statistics*, 11(4):427–444, 1940. [10.1214/aoms/1177731829](https://doi.org/10.1214/aoms/1177731829).
- P. Diaconis and D. Stroock. Geometric bounds for eigenvalues of Markov chains. *The annals of applied probability*, pages 36–61, 1991. <https://doi.org/10.1214/aoap/1177005980>.
- T. Dick, C. Dwork, M. Kearns, T. Liu, A. Roth, G. Vietri, and Z. S. Wu. Confidence-ranked reconstruction of census microdata from published statistics. *Proceedings of the National Academy of Sciences*, 120(8):e2218605120, 2023. <https://doi.org/10.1073/pnas.2218605120>.
- I. Dinur and K. Nissim. Revealing information while preserving privacy. In *Proceedings of the twenty-second ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pages 202–210, 2003. <https://doi.org/10.1145/773153.773173>.
- C. Dwork, K. Greenewald, and M. Raghavan. Synthetic census data generation via multidimensional multiset sum. In *Proceedings of the 2024 Symposium on Foundations of Responsible Computing*, 2024. <https://arxiv.org/pdf/2404.10095>.
- M. Dyer. Approximate counting by dynamic programming. In *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pages 693–699, 2003. <https://doi.org/10.1145/780542.780643>.
- M. Dyer, A. Frieze, R. Kannan, A. Kapoor, L. Perkovic, and U. Vazirani. A mildly exponential time algorithm for approximating the number of solutions to a multidimensional knapsack problem. *Combinatorics, Probability and Computing*, 2(3):271–284, 1993. <https://doi.org/10.1017/S0963548300000675>.
- I. Z. Emiris, A. Karasoulou, and C. Tzovas. Approximating multidimensional subset sum and Minkowski decomposition of polygons. *Mathematics in Computer Science*, 11:35–48, 2017. <https://doi.org/10.1214/aoms/1177731829>.
- B. Farooq, M. Bierlaire, R. Hurtubia, and G. Flötteröd. Simulation based population synthesis. *Transportation Research Part B: Methodological*, 58:243–263, 2013. <https://doi.org/10.1016/j.trb.2013.05.008>.

- [//doi.org/10.1016/j.trb.2013.09.012](https://doi.org/10.1016/j.trb.2013.09.012).
- W. Feller. *An introduction to probability theory and its applications*, volume 1. John Wiley & Sons, 3 edition, 1968. ISBN 978-0-471-25708-0.
- P. Francis. A note on the misinterpretation of the US Census re-identification attack. In *International Conference on Privacy in Statistical Databases*, pages 299–311. Springer, 2022. [https://doi.org/10.1007/978-3-031-13945-1\\_21](https://doi.org/10.1007/978-3-031-13945-1_21).
- P. Gawrychowski, L. Markin, and O. Weimann. A faster FPTAS for #knapsack. In I. Chatzigiannakis, C. Kaklamanis, D. Marx, and D. Sannella, editors, *45th International Colloquium on Automata, Languages, and Programming, ICALP*, volume 107, pages 64:1–64:13, 2018. <https://doi.org/10.4230/LIPIcs.ICALP.2018.64>.
- P. Gopalan, A. Klivans, R. Meka, D. Štefankovic, S. Vempala, and E. Vigoda. An fptas for #knapsack and related counting problems. In *2011 IEEE 52nd Annual Symposium on Foundations of Computer Science*, pages 817–826. IEEE, 2011. <https://doi.org/10.1109/FOCS.2011.32>.
- G. Grimmett. *Percolation*. Springer, 2 edition, 1999. ISBN 978-3-540-64902-1.
- V. Guruswami. Rapidly mixing Markov chains: a comparison of techniques (a survey). *arXiv preprint arXiv:1603.01512*, 2016. <https://arxiv.org/pdf/1603.01512>.
- J. Gussenbauer, M. Templ, S. Fritzmann, and A. Kowarik. Simulation of calibrated complex synthetic population data with xgboost. *Algorithms*, 17(6):249, 2024. <https://doi.org/10.3390/a17060249>.
- K. Harland, A. Heppenstall, D. Smith, and M. H. Birkin. Creating realistic synthetic populations at varying spatial scales: A comparative critique of population synthesis techniques. *Journal of Artificial Societies and Social Simulation*, 15(1), 2012. [10.18564/jasss.1909](https://doi.org/10.18564/jasss.1909).
- M. E. Hauer and A. R. Santos-Lozada. Differential privacy in the 2020 census will distort COVID-19 rates. *Socius*, 7:2378023121994014, 2021. <https://doi.org/10.1177/2378023121994014>.
- J. Hendrix and B. F. Jones. Bounded integer linear constraint solving via lattice search. In *Proceedings of the International Workshop on Satisfiability Modulo Theories*, 2015. <https://bfj7.com/papers/blt-paper.pdf>.
- M. Hujter. Improved lower and upper bounds for the number of feasible solutions to a knapsack problem. *Optimization*, 19(6):889–894, 1988. <https://doi.org/10.1080/02331938808843401>.
- G. P. Ingargiola and J. F. Korsh. A general algorithm for one-dimensional knapsack problems. *Operations Research*, 25(5):752–759, 1977. <https://doi.org/10.1287/opre.25.5.752>.
- M. Jerrum and A. Sinclair. Conductance and the rapid mixing property for Markov chains: the approximation of permanent resolved. In *Proceedings of the twentieth annual ACM symposium on Theory of computing*, pages 235–244, 1988. <https://doi.org/10.1145/62212.62234>.
- M. Jerrum and A. Sinclair. Approximating the permanent. *SIAM journal on computing*, 18(6):1149–1178, 1989. <https://doi.org/10.1137/0218077>.
- M. R. Jerrum, L. G. Valiant, and V. V. Vazirani. Random generation of combinatorial structures from a uniform distribution. *Theoretical computer science*, 43:169–188, 1986. [https://doi.org/10.1016/0304-3975\(86\)90174-X](https://doi.org/10.1016/0304-3975(86)90174-X).
- V. Kargin. Relaxation time is monotone in temperature in the mean-field Ising model. *Statistics & probability letters*, 81(8):1094–1097, 2011. <https://doi.org/10.1016/j.spl.2011.03.001>.

- K. K. Kayibi, S. Pirzada, and C. Rutherford. Mixing time of Markov chain of the knapsack problem. *arXiv preprint arXiv:1803.06914*, 2018. <https://arxiv.org/pdf/1803.06914>.
- C. T. Kenny, S. Kuriwaki, C. McCartan, E. Rosenman, T. Simko, and K. Imai. The use of differential privacy for census data and its impact on redistricting: The case of the 2020 U.S. Census. *Science Advances*, 7, 2021. <https://doi.org/10.1126/sciadv.abk3283>.
- C. T. Kenny, C. McCartan, S. Kuriwaki, T. Simko, and K. Imai. Evaluating bias and noise induced by the US Census Bureau’s privacy protection methods. *Science Advances*, 10(18), 2024. <https://doi.org/10.1126/sciadv.adl2524>.
- A. Kulik and H. Shachnai. There is no EPTAS for two-dimensional knapsack. *Information Processing Letters*, 110(16):707–710, 2010. <https://doi.org/10.1016/j.ipl.2010.05.031>.
- T. A. Lambe. Bounds on the number of feasible solutions to a knapsack problem. *SIAM Journal on Applied Mathematics*, 26(2):302–305, 1974. <https://doi.org/10.1137/0126027>.
- T. A. Lambe. Upper bound on the number of nonnegative integer solutions to a linear equation. *SIAM Journal on Applied Mathematics*, 32(1):215–219, 1977. <https://doi.org/10.1137/0132018>.
- E. L. Lawler. Fast approximation algorithms for knapsack problems. *Mathematics of Operations Research*, 4:339–356, 1979. <https://doi.org/10.1287/moor.4.4.339>.
- G. F. Lawler and A. D. Sokal. Bounds on the  $L^2$  spectrum for Markov chains and Markov processes: a generalization of Cheeger’s inequality. *Transactions of the American mathematical society*, 309(2):557–580, 1988. <https://doi.org/10.2307/2000925>.
- F. Lazebnik. On systems of linear Diophantine equations. *Mathematics Magazine*, 69(4):261–266, 1996. <https://doi.org/10.2307/2690528>.
- D. A. Levin, Y. Peres, and E. L. Wilmer. *Markov chains and mixing times*. American Mathematical Soc., 2017. ISBN 978-1-4704-2962-1.
- R. Lovelace and D. Ballas. ‘Truncate, replicate, sample’: A method for creating integer weights for spatial microsimulation. *Computers, Environment and Urban Systems*, 41:1–11, 2013. <https://doi.org/10.1016/j.compenvurbsys.2013.03.004>.
- L. Ma and S. Srinivasan. Synthetic population generation with multilevel controls: A fitness-based synthesis approach and validations. *Computer-Aided Civil and Infrastructure Engineering*, 30(2):135–150, 2015. <https://doi.org/10.1111/mice.12085>.
- M. J. Magazine and M.-S. Chern. A note on approximation schemes for multidimensional knapsack problems. *Mathematics of Operations Research*, 9(2):244–247, 1984. <https://doi.org/10.1287/moor.9.2.244>.
- R. Mahmoudvand, H. Hassani, A. Farzaneh, and G. Howell. The exact number of nonnegative integer solutions for a linear Diophantine inequality. *IAENG International Journal of Applied Mathematics*, 40(1):1–5, 2010. [https://www.iaeng.org/IJAM/issues\\_v40/issue\\_1/IJAM\\_40\\_1\\_01.pdf](https://www.iaeng.org/IJAM/issues_v40/issue_1/IJAM_40_1_01.pdf).
- S. Martello and P. Toth. Algorithms for knapsack problems. *North-Holland Mathematics Studies*, 132:213–257, 1987. [https://doi.org/10.1016/S0304-0208\(08\)73237-7](https://doi.org/10.1016/S0304-0208(08)73237-7).
- B. Morris and A. Sinclair. Random walks on truncated cubes and sampling 0-1 knapsack solutions. *SIAM journal on computing*, 34(1):195–226, 2004. <https://doi.org/10.1137/S0097539702411915>.
- J. T. Mueller and A. R. Santos-Lozada. The 2020 US Census differential privacy method introduces disproportionate discrepancies for rural and non-white populations. *Population Research and Policy Review*, 41(4):1417–1430, 2022. <https://doi.org/10.1007/s11113-022-09698-3>.

- K. Müller and K. W. Axhausen. Population synthesis for microsimulation: State of the art. *Arbeitsberichte Verkehrs-und Raumplanung*, 638, 2010. <https://doi.org/10.3929/ethz-a-006127782>.
- Ş. Nacu. Glauber dynamics on the cycle is monotone. *Probability theory and related fields*, 127:177–185, 2003. <https://doi.org/10.1007/s00440-003-0279-x>.
- OEIS Foundation Inc. The On-Line Encyclopedia of Integer Sequences, 2023. Published electronically at <http://oeis.org>.
- M. W. Padberg. A remark on “An inequality for the number of lattice points in a simplex”. *SIAM Journal on Applied Mathematics*, 20(4):638–641, 1971. <https://doi.org/10.1137/0120063>.
- S. Petti and A. Flaxman. Differential privacy in the 2020 US census: What will it do? Quantifying the accuracy/privacy tradeoff. *Gates open research*, 3, 2019. <https://doi.org/10.12688/gatesopenres.13089.2>.
- D. Pisinger. An exact algorithm for large multiple knapsack problems. *European Journal of Operational Research*, 114(3):528–541, 1999. [https://doi.org/10.1016/S0377-2217\(98\)00120-9](https://doi.org/10.1016/S0377-2217(98)00120-9).
- U. Porod. *Dynamics of Markov Chains for Undergraduates*. Preprint, 2024. <https://www.math.northwestern.edu/documents/book-markov-chains.pdf>.
- J. Puchinger, G. R. Raidl, and U. Pferschy. The multidimensional knapsack problem: Structure and algorithms. *INFORMS Journal on Computing*, 22(2):250–265, 2010. <https://doi.org/10.1287/ijoc.1090.0344>.
- R. Rizzi and A. I. Tomescu. Faster FPTASes for counting and random generation of knapsack solutions. In *European Symposium on Algorithms*, pages 762–773. Springer, 2014. [https://doi.org/10.1007/978-3-662-44777-2\\_63](https://doi.org/10.1007/978-3-662-44777-2_63).
- S. Ruggles, C. Fitch, D. Magnuson, and J. Schroeder. Differential privacy and census data: Implications for social and economic research. In *AEA papers and proceedings*, volume 109, pages 403–408, 2019. <https://doi.org/10.1257/pandp.20191107>.
- S. Ruggles, S. Flood, M. Sobek, D. Backman, A. Chen, G. Cooper, S. Richards, R. Rodgers, and M. Schouweiler. IPUMS USA: Version 15.0 [dataset], 2024.
- H. M. Salkin and C. A. De Kluyver. The knapsack problem: a survey. *Naval Research Logistics Quarterly*, 22(1):127–144, 1975. <https://doi.org/10.1002/nav.3800220110>.
- A. Sánchez-Roselly Navarro. *Linear Diophantine equations and applications*. PhD thesis, Universidad de Granada, 2016. <https://digibug.ugr.es/bitstream/handle/10481/40613/24868875.pdf>.
- A. R. Santos-Lozada, J. T. Howard, and A. M. Verdery. How differential privacy will affect our understanding of health disparities in the United States. *Proceedings of the National Academy of Sciences*, 117(24):13405–13412, 2020. <https://doi.org/10.1073/pnas.2003714117>.
- A. Sinclair and M. Jerrum. Approximate counting, uniform generation and rapidly mixing markov chains. *Information and Computation*, 82(1):93–133, 1989. [https://doi.org/10.1016/0890-5401\(89\)90067-9](https://doi.org/10.1016/0890-5401(89)90067-9).
- P. Sousi. Mixing times of Markov chains. *Preprint*, 2020. <https://personal.math.ubc.ca/~jhermon/Mixing/mixing-notes.pdf>.
- R. P. Stanley. *Enumerative combinatorics*, volume 2. Cambridge university press, 1999. <https://doi.org/10.1017/CB09780511609589>.
- R. Steed, T. Liu, Z. S. Wu, and A. Acquisti. Policy impacts of statistical uncertainty and privacy. *Science*, 377(6609):928–931, 2022. <https://doi.org/10.1126/science>.

- abq4481.
- L. Sun and A. Erath. A Bayesian network approach for population synthesis. *Transportation Research Part C: Emerging Technologies*, 61:49–62, 2015. <https://doi.org/10.1016/j.trc.2015.10.010>.
- L. Sun, A. Erath, and M. Cai. A hierarchical mixture modeling framework for population synthesis. *Transportation Research Part B: Methodological*, 114:199–212, 2018. <https://doi.org/10.1016/j.trb.2018.06.002>.
- U.S. Census Bureau. 2020 census privacy-protected microdata file (ppmf): 2020 census of population and housing, technical documentation. Technical documentation, U.S. Census Bureau, Washington, DC, September 2024. <https://www.census.gov>. TD/DHC2020.
- D. Voas and P. Williamson. An evaluation of the combinatorial optimisation approach to the creation of synthetic microdata. *International Journal of Population Geography*, 6(5):349–366, 2000. [https://doi.org/10.1002/1099-1220\(200009/10\)6:5%3C349::AID-IJPG196%3E3.0.CO;2-5](https://doi.org/10.1002/1099-1220(200009/10)6:5%3C349::AID-IJPG196%3E3.0.CO;2-5).
- H. Wang, S. Sudalairaj, J. Henning, K. Greenewald, and A. Srivastava. Post-processing private synthetic data for improving utility on selected measures. *Advances in Neural Information Processing Systems*, 36, 2024. [10.5555/3666122.3668924](https://arxiv.org/abs/2405.15555).
- A. Whitworth. synthACS: spatial microsimulation modeling with synthetic American Community Survey data. *Journal of Statistical Software*, 104:1–30, 2022. <https://doi.org/10.18637/jss.v104.i07>.
- P. Williamson, M. Birkin, and P. H. Rees. The estimation of population microdata by using data from small area statistics and samples of anonymised records. *Environment and Planning A*, 30(5):785–816, 1998. <https://doi.org/10.1068/a300785>.
- R. L. Winkler, J. L. Butler, K. J. Curtis, and D. Egan-Robertson. Differential privacy and the accuracy of county-level net migration estimates. *Population Research and Policy Review*, pages 1–19, 2021. <https://doi.org/10.1007/s11113-021-09664-5>.
- G. Wu, A. Heppenstall, P. Meier, R. Purshouse, and N. Lomax. A synthetic population dataset for estimating small area health and socio-economic outcomes in Great Britain. *Scientific Data*, 9(1):19, 2022. <https://doi.org/10.1038/s41597-022-01124-9>.

**Organization of the appendix.** Appendix A contains detailed information on our encoding of Census data into  $d$ -dimensional vectors. We present a rejection sampling algorithm that is equivalent to our algorithm under our choice of  $\pi$  in Appendix B. In Appendices C and D, we provide omitted proofs and details from Sections 4.1 and 4.2 respectively. We present additional empirical results for both the simple and reduced chains in Appendix E. We formally show that **MMS** is NP-hard in Appendix F. In Appendix G, we construct pathological synthetic examples under which the reduced chain is either not irreducible or has exponentially large mixing time. Finally, in Appendix H, we describe our final algorithm and theoretically characterize its performance.

#### APPENDIX A. DATA ENCODING DETAILS

Census data are encoded with 7 race categories (white, Black, American Indian and Native Alaskan, Asian, Hawaiian and Pacific Islander, other, and two or more) and two ethnicity categories (Hispanic or non-Hispanic). We encode each distinct household  $\mathbf{v}_i$  and block-level statistic  $\mathbf{c}$  as a 135-dimensional integer vector. The encoding is as follows, annotated with the dimensionality of each component:

- number of individuals of each (race, ethnicity) combination (**14**).
- number of adults (defined as 18 or older) of each race (**7**).
- number of Hispanic adults (**1**).
- (householder race)  $\times$  (whether or not the household is a family)  $\times$  (household size in  $\{1, 2, 3, 4, 5, 6, 7+\}$ ) (**98**). Note that these form a partition, so every household has a 1 in exactly one of these 98 components.
- for households with a Hispanic householder, (whether or not the household is a family)  $\times$  (household size in  $\{1, 2, 3, 4, 5, 6, 7+\}$ ) (**14**).<sup>24</sup>
- number of households (a constant 1 for each household) (**1**).

Each row of the PUMS sample contains all of this information. We choose this particular encoding because this information can also be found in block-level aggregate data from SF1.<sup>25</sup> Each block-level vector  $\mathbf{c}$  is the sum of these characteristics of all households in the block. These totals can be found in a combination of SF1 tables P3, P28A–H, and P16A–H. Some of these tables do not account for *group quarters* (as opposed to households), which exist in around 1% of blocks in both AL and NV. We exclude these blocks from our analysis. In principle, our methods would also apply when reconstructing group quarters. Moreover, there are blocks for which no solution to our optimization problem exists (i.e., no combination of PUMS households matches the reported block totals). This can happen because the PUMS data are incomplete, since a “rare” household might be needed to match the block-level constraints. We exclude these blocks as well (about 5% in AL and 14% in NV). In our publicly available code, when we encounter such blocks, we simply reduce the dimensionality of the vectors and apply the same techniques to solve them. For example, instead of matching the number of adults of each race and ethnicity, we simply match the total number of adults in the block.

There are  $n = 2,745$  distinct households (encoded as 135-dimensional vectors) in the AL PUMS and  $n = 4,094$  distinct households in the NV PUMS. The data we use are

<sup>24</sup>Because ethnicity is encoded as a binary variable and the remaining two attributes are completely determined by the (race)  $\times$  (is family)  $\times$  (household size) variables, including both (Householder is Hispanic) and (Householder is not Hispanic) is redundant here.

<sup>25</sup>SF1 data can be accessed at <https://www.nhgis.org/>.

subject to disclosure avoidance systems before their release. They do not represent “ground truth”; nevertheless, we treat them as the source of the distribution we seek to sample from. Downstream analyses should be careful not to rely on the exact locations of very rare households, since these are unlikely to be accurate.

## APPENDIX B. REJECTION SAMPLING FORMULATION

Here, we present a rejection sampling algorithm (Algorithm 1) that samples from  $\mathcal{X}$  according to the distribution given in (2.1). This algorithm is far too inefficient to be practical; it is intended simply to illustrate and provide motivation for the problem specification here.

---

### Algorithm 1 REJECTIONSAMPLING( $\mathbf{V}, \mathbf{c}, \mathcal{D}$ )

---

```

repeat                                     ▷ Repeat until a sample is accepted
   $\mathbf{x} \leftarrow \mathbf{0} \in \mathbb{Z}_{\geq 0}^n$            ▷ Empty multiset
  while  $\mathbf{V}\mathbf{x} \preceq \mathbf{c}$  do                 ▷  $\preceq$  denotes element-wise  $\leq$ .
    Sample  $\mathbf{v}_i$  from  $\mathcal{D}$ 
     $\mathbf{x}[i] \leftarrow \mathbf{x}[i] + 1$ 
  end while
until  $\mathbf{V}\mathbf{x} = \mathbf{c}$ 
return  $\mathbf{x}$ 

```

---

## APPENDIX C. OMITTED PROOFS AND DETAILS FROM SECTION 4.1

---

### Algorithm 2 SIMPLE( $\mathbf{V}, \mathbf{c}, f, \gamma, t$ )

---

```

repeat
   $\mathbf{x} \leftarrow \mathbf{0} \in \mathbb{Z}_{\geq 0}^n$ 
  for  $t$  iterations do
    with probability 0.5
      continue
    end
     $i \leftarrow$  a random integer in  $[n]$  such that  $\mathbf{v}_i \preceq \mathbf{c}$    ▷ Ignore ineligible households
     $\mathbf{r} \leftarrow \mathbf{c} - \mathbf{V}\mathbf{x}_{i \leftarrow 0}$ 
     $g_{\max} \leftarrow \max \{g \in \mathbb{Z}_{\geq 0} : g\mathbf{v}_i \preceq \mathbf{r}\}$ 
     $\mathcal{S} \leftarrow \{\mathbf{x}_{i \leftarrow g} : 0 \leq g \leq g_{\max}\}$ 
     $\mathbf{x} \leftarrow$  a random sample from  $\tilde{\pi}_\gamma \mid \mathbf{x} \in \mathcal{S}$  (i.e.,  $\propto f(\mathbf{x}) \exp(-\gamma \|\mathbf{V}\mathbf{x} - \mathbf{c}\|)$  on  $\mathcal{S}$ )
  end for
until  $\mathbf{V}\mathbf{x} = \mathbf{c}$ 
return  $\mathbf{x}$ 

```

---

### C.1. Formal algorithm description for the simple chain.

**C.2. Theoretical results for the simple chain.** First, we prove Lemma 1. We restate each result before proving it.

**Lemma 1.** Let  $\sigma$  and  $\sigma'$  be distributions defined on  $\mathcal{Y}$ . Let  $\sigma_{\mathcal{X}}$  and  $\sigma'_{\mathcal{X}}$  be their respective conditional distributions on  $\mathcal{X} \subset \mathcal{Y}$ , i.e., for  $\mathbf{x} \in \mathcal{X}$ ,  $\sigma_{\mathcal{X}}(\mathbf{x}) = \sigma(\mathbf{x})/\sigma(\mathcal{X})$ . If  $d_{\text{TV}}(\sigma, \sigma') \leq \varepsilon$ , then  $d_{\text{TV}}(\sigma_{\mathcal{X}}, \sigma'_{\mathcal{X}}) \leq 3\varepsilon/(2\sigma(\mathcal{X}))$ . For any  $\varepsilon$ , there exist instances for which this is tight to within a constant factor.

*Proof.* First, we show that bounded TVD implies bounded TVD on conditional distributions, inflated by a factor of  $1/\sigma(\mathcal{X})$ .

$$\begin{aligned}
 d_{\text{TV}}(\sigma, \sigma') &\leq \varepsilon \\
 \|\sigma - \sigma'\|_1 &\leq 2\varepsilon \\
 \sum_{\mathbf{x} \in \mathcal{X}} |\sigma(\mathbf{x}) - \sigma'(\mathbf{x})| &\leq 2\varepsilon \\
 \sum_{\mathbf{x} \in \mathcal{X}} \left| \frac{\sigma(\mathbf{x})}{\sigma(\mathcal{X})} - \frac{\sigma'(\mathbf{x})}{\sigma(\mathcal{X})} \right| &\leq \frac{2\varepsilon}{\sigma(\mathcal{X})} \\
 \sum_{\mathbf{x} \in \mathcal{X}} \left| \frac{\sigma(\mathbf{x})}{\sigma(\mathcal{X})} - \frac{\sigma'(\mathbf{x})}{\sigma'(\mathcal{X})} + \frac{\sigma'(\mathbf{x})}{\sigma'(\mathcal{X})} - \frac{\sigma'(\mathbf{x})}{\sigma(\mathcal{X})} \right| &\leq \frac{2\varepsilon}{\sigma(\mathcal{X})} \\
 \sum_{\mathbf{x} \in \mathcal{X}} \left| \frac{\sigma(\mathbf{x})}{\sigma(\mathcal{X})} - \frac{\sigma'(\mathbf{x})}{\sigma'(\mathcal{X})} \right| - \left| \frac{\sigma'(\mathbf{x})}{\sigma'(\mathcal{X})} - \frac{\sigma'(\mathbf{x})}{\sigma(\mathcal{X})} \right| &\leq \frac{2\varepsilon}{\sigma(\mathcal{X})} \quad (\text{triangle inequality}) \\
 \sum_{\mathbf{x} \in \mathcal{X}} \left| \sigma_{\mathcal{X}}(\mathbf{x}) - \sigma'_{\mathcal{X}}(\mathbf{x}) \right| - \sigma'(\mathbf{x}) \left| \frac{1}{\sigma'(\mathcal{X})} - \frac{1}{\sigma(\mathcal{X})} \right| &\leq \frac{2\varepsilon}{\sigma(\mathcal{X})} \\
 \|\sigma_{\mathcal{X}} - \sigma'_{\mathcal{X}}\|_1 - \sigma'(\mathcal{X}) \left| \frac{1}{\sigma'(\mathcal{X})} - \frac{1}{\sigma(\mathcal{X})} \right| &\leq \frac{2\varepsilon}{\sigma(\mathcal{X})} \\
 \|\sigma_{\mathcal{X}} - \sigma'_{\mathcal{X}}\|_1 &\leq \frac{2\varepsilon}{\sigma(\mathcal{X})} + \sigma'(\mathcal{X}) \frac{|\sigma(\mathcal{X}) - \sigma'(\mathcal{X})|}{\sigma'(\mathcal{X})\sigma(\mathcal{X})} \\
 \|\sigma_{\mathcal{X}} - \sigma'_{\mathcal{X}}\|_1 &\leq \frac{2\varepsilon}{\sigma(\mathcal{X})} + \frac{|\sigma(\mathcal{X}) - \sigma'(\mathcal{X})|}{\sigma(\mathcal{X})} \\
 \|\sigma_{\mathcal{X}} - \sigma'_{\mathcal{X}}\|_1 &\leq \frac{3\varepsilon}{\sigma(\mathcal{X})} \quad (d_{\text{TV}}(\sigma, \sigma') \leq \varepsilon) \\
 d_{\text{TV}}(\sigma_{\mathcal{X}}, \sigma'_{\mathcal{X}}) &\leq \frac{3\varepsilon}{2\sigma(\mathcal{X})}
 \end{aligned}$$

To show that this is tight to within a constant factor, consider an instance where  $\mathcal{X} = \{\mathbf{x}_0, \mathbf{x}_1\}$  and assume without loss of generality  $\sigma(\mathbf{x}_0) \leq \sigma(\mathbf{x}_1)$ . For a given  $\varepsilon$  define  $\sigma'$  as follows:

$$\begin{aligned}
 \sigma'(\mathbf{x}_0) &= \sigma(\mathbf{x}_0) + \varepsilon \\
 \sigma'(\mathbf{x}_1) &= \max(0, \sigma(\mathbf{x}_1) - \varepsilon) \\
 \sigma'(\mathbf{x}) &= \begin{cases} \sigma(\mathbf{x}) & \sigma(\mathbf{x}_1) \geq \varepsilon \\ \sigma(\mathbf{x}) \left(1 - \frac{\varepsilon - \sigma(\mathbf{x}_1)}{1 - \sigma(\mathcal{X})}\right) & \sigma(\mathbf{x}_1) < \varepsilon \end{cases} \quad (\mathbf{x} \notin \mathcal{X})
 \end{aligned}$$

First, observe that if  $\sigma(\mathbf{x}_1) < \varepsilon$ , then  $\sigma'(\mathbf{x}_1) = 0$  and  $d_{\text{TV}}(\sigma_{\mathcal{X}}, \sigma'_{\mathcal{X}}) \geq \frac{1}{2}$ . This is within a factor of 2 of our upper bound since trivially  $d_{\text{TV}}(\sigma_{\mathcal{X}}, \sigma'_{\mathcal{X}}) \leq 1$ . If  $\sigma(\mathbf{x}_1) \geq \varepsilon$ , then

$$\begin{aligned}
d_{\text{TV}}(\sigma_{\mathcal{X}}, \sigma'_{\mathcal{X}}) &= \frac{1}{2} \|\sigma_{\mathcal{X}} - \sigma'_{\mathcal{X}}\|_1 \\
&= \frac{1}{2} \sum_{\mathbf{x} \in \mathcal{X}} |\sigma_{\mathcal{X}}(\mathbf{x}) - \sigma'_{\mathcal{X}}(\mathbf{x})| \\
&= \frac{1}{2} \sum_{\mathbf{x} \in \mathcal{X}} \left| \frac{\sigma(\mathbf{x})}{\sigma(\mathcal{X})} - \frac{\sigma'(\mathbf{x})}{\sigma'(\mathcal{X})} \right| \\
&= \frac{1}{2} \sum_{\mathbf{x} \in \mathcal{X}} \left| \frac{\sigma(\mathbf{x})}{\sigma(\mathcal{X})} - \frac{\sigma'(\mathbf{x})}{\sigma(\mathcal{X})} \right| && (\sigma'(\mathcal{X}) = \sigma(\mathcal{X})) \\
&= \frac{1}{2\sigma(\mathcal{X})} \sum_{\mathbf{x} \in \mathcal{X}} |\sigma(\mathbf{x}) - \sigma'(\mathbf{x})| \\
&= \frac{1}{2\sigma(\mathcal{X})} (|\sigma(\mathbf{x}_0) - \sigma'(\mathbf{x}_0)| + |\sigma(\mathbf{x}_1) - \sigma'(\mathbf{x}_1)|) \\
&= \frac{\varepsilon}{\sigma(\mathcal{X})}.
\end{aligned}$$

□

We next provide a theoretical explanation for the poor performance of the simple chain here, which we will use to motivate the reduced chain in Section 4.2. In particular, we characterize how  $\gamma$  impacts  $N_\gamma$ . We show how varying  $\gamma$  creates a trade-off between raising  $p_\gamma^*$  and lowering  $\tau_\gamma^*(\varepsilon)$ . First, we show that  $p_\gamma^*$  increases with  $\gamma$ , ranging from 0 to 1 as  $\gamma$  ranges from  $-\infty$  to  $\infty$ .

**Lemma 2.**  $p_\gamma^*$  is monotonically increasing in  $\gamma$ , and

$$\lim_{\gamma \rightarrow -\infty} p_\gamma^* = 0 \quad \text{and} \quad \lim_{\gamma \rightarrow \infty} p_\gamma^* = 1.$$

*Proof.* First, we show that  $p_\gamma^*$  is monotonically increasing in  $\gamma$ . By definition,

$$\begin{aligned}
\frac{d}{d\gamma} p_\gamma^* &= \frac{d}{d\gamma} \frac{\tilde{\pi}_\gamma(\mathcal{X})}{\tilde{\pi}_\gamma(\mathcal{Y})} \\
&= \frac{d}{d\gamma} \frac{\sum_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}) \exp(-\gamma \|\mathbf{V}\mathbf{x} - \mathbf{c}\|_1)}{\sum_{\mathbf{x} \in \mathcal{Y}} f(\mathbf{x}) \exp(-\gamma \|\mathbf{V}\mathbf{x} - \mathbf{c}\|_1)} \\
&= \frac{d}{d\gamma} \frac{\sum_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})}{\sum_{\mathbf{x} \in \mathcal{Y}} f(\mathbf{x}) \exp(-\gamma \|\mathbf{V}\mathbf{x} - \mathbf{c}\|_1)} && (\mathbf{V}\mathbf{x} - \mathbf{c} = 0 \text{ for } x \in \mathcal{X}) \\
&= \left( \sum_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}) \right) \frac{d}{d\gamma} \frac{1}{\sum_{\mathbf{x} \in \mathcal{Y}} f(\mathbf{x}) \exp(-\gamma \|\mathbf{V}\mathbf{x} - \mathbf{c}\|_1)}.
\end{aligned}$$

It suffices to show

$$\sum_{\mathbf{x} \in \mathcal{Y}} \frac{d}{d\gamma} f(\mathbf{x}) \exp(-\gamma \|\mathbf{c} - \mathbf{V}\mathbf{x}\|_1) = \sum_{\mathbf{x} \in \mathcal{Y}} -\|\mathbf{c} - \mathbf{V}\mathbf{x}\|_1 f(\mathbf{x}) \exp(-\gamma \|\mathbf{c} - \mathbf{V}\mathbf{x}\|_1) < 0.$$

Next, we show that  $p_\gamma^*$  has the claimed limits. For any  $\mathbf{x}' \in \mathcal{Y} \setminus \mathcal{X}$  and  $\gamma < 0$ ,

$$\exp(-\gamma \|\mathbf{c} - \mathbf{V}\mathbf{x}'\|_1) f(\mathbf{x}') \geq \exp(-\gamma) f(\mathbf{x}').$$

Therefore,

$$\begin{aligned} \lim_{\gamma \rightarrow -\infty} p_\gamma^* &= \lim_{\gamma \rightarrow -\infty} \frac{\tilde{\pi}_\gamma(\mathcal{X})}{\tilde{\pi}_\gamma(\mathcal{Y})} \\ &= \lim_{\gamma \rightarrow -\infty} \frac{\sum_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}) \exp(-\gamma \|\mathbf{c} - \mathbf{V}\mathbf{x}\|_1)}{\sum_{\mathbf{x} \in \mathcal{Y}} f(\mathbf{x}) \exp(-\gamma \|\mathbf{c} - \mathbf{V}\mathbf{x}\|_1)} \\ &\leq \lim_{\gamma \rightarrow -\infty} \frac{\sum_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}) \exp(-\gamma \|\mathbf{c} - \mathbf{V}\mathbf{x}\|_1)}{\sum_{\mathbf{x} \in \mathcal{Y} \setminus \mathcal{X}} f(\mathbf{x}) \exp(-\gamma \|\mathbf{c} - \mathbf{V}\mathbf{x}\|_1)} \\ &\leq \lim_{\gamma \rightarrow -\infty} \frac{\sum_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})}{\exp(-\gamma) \sum_{\mathbf{x}' \in \mathcal{Y} \setminus \mathcal{X}} f(\mathbf{x}')} \\ &= 0. \end{aligned}$$

Similarly, for  $\mathbf{x}' \in \mathcal{Y} \setminus \mathcal{X}$  and  $\gamma > 0$ ,

$$\exp(-\gamma \|\mathbf{c} - \mathbf{V}\mathbf{x}'\|_1) f(\mathbf{x}') \leq \exp(-\gamma) f(\mathbf{x}').$$

Therefore,

$$\begin{aligned} \lim_{\gamma \rightarrow \infty} p_\gamma^* &= \lim_{\gamma \rightarrow \infty} \frac{\tilde{\pi}_\gamma(\mathcal{X})}{\tilde{\pi}_\gamma(\mathcal{Y})} \\ &= \lim_{\gamma \rightarrow \infty} \frac{\sum_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}) \exp(-\gamma \|\mathbf{c} - \mathbf{V}\mathbf{x}\|_1)}{\sum_{\mathbf{x} \in \mathcal{Y}} f(\mathbf{x}) \exp(-\gamma \|\mathbf{c} - \mathbf{V}\mathbf{x}\|_1)} \\ &= \lim_{\gamma \rightarrow \infty} \frac{\sum_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})}{\sum_{\mathbf{x} \in \mathcal{Y} \setminus \mathcal{X}} f(\mathbf{x}) \exp(-\gamma \|\mathbf{c} - \mathbf{V}\mathbf{x}\|_1) + \sum_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})} \\ &\geq \lim_{\gamma \rightarrow \infty} \frac{\sum_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})}{\exp(-\gamma) \sum_{\mathbf{x} \in \mathcal{Y} \setminus \mathcal{X}} f(\mathbf{x}) + \sum_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})} \\ &= 1. \end{aligned}$$

□

**Lemma 3.** For the stationary distribution  $\tilde{\pi}_\gamma$  as defined in (4.4),  $\tau_{\text{rel}, \gamma} = \Omega(\exp(\gamma \min_i \|\mathbf{v}_i\|_1))$ .

*Proof.* We will provide a lower and an upper bound on the conductance of  $M_\gamma$  and use Cheeger's inequality to lower-bound the mixing time. The conductance (see, e.g., Guruswami, 2016) of  $M_\gamma$  is defined as

$$\Phi(M_\gamma) = \min_{\substack{\mathcal{S} \subset \mathcal{Y} \\ 0 < \tilde{\pi}_\gamma(\mathcal{S}) \leq 1/2}} \frac{Q_\gamma(\mathcal{S}, \overline{\mathcal{S}})}{\tilde{\pi}_\gamma(\mathcal{S})}, \quad (\text{C.1})$$

where  $Q_\gamma(\mathbf{x}, \mathbf{x}') = \tilde{\pi}_\gamma(\mathbf{x}) P_\gamma(\mathbf{x}, \mathbf{x}')$  and  $Q_\gamma(\mathcal{S}, \overline{\mathcal{S}}) = \sum_{\mathbf{x} \in \mathcal{S}, \mathbf{x}' \notin \mathcal{S}} Q_\gamma(\mathbf{x}, \mathbf{x}')$ . We will show that in particular,  $Q_\gamma(\mathcal{X}, \overline{\mathcal{X}}) / \tilde{\pi}_\gamma(\mathcal{X})$  is small, providing an upper bound for  $\Phi(M_\gamma)$ .

To do so, let  $p_i = \Pr_{\mathcal{D}}[\mathbf{v}_i]$  and  $\ell = \|\mathbf{v}_i\|_1$ . Observe that for any  $\mathbf{x} \in \mathcal{X}$  and  $i \in [n]$ , if  $\gamma$  is sufficiently large,

$$\begin{aligned}
\frac{\tilde{\pi}_\gamma(\mathbf{x})}{\tilde{\pi}_\gamma(\Delta(\mathbf{x}, i))} &= \frac{\tilde{\pi}_\gamma(\mathbf{x})}{\sum_{g=0}^{\mathbf{x}[i]} \tilde{\pi}_\gamma(\mathbf{x}_{i \leftarrow g})} \\
&= \frac{f(\mathbf{x})}{\sum_{g=0}^{\mathbf{x}[i]} f(\mathbf{x}_{i \leftarrow g}) \exp(-\gamma \|\mathbf{c} - \mathbf{V}\mathbf{x}_{i \leftarrow g}\|_1)} \\
&= \frac{f(\mathbf{x})}{\sum_{g=0}^{\mathbf{x}[i]} f(\mathbf{x}_{i \leftarrow (\mathbf{x}[i] - g)}) \exp(-\gamma \|\mathbf{c} - \mathbf{V}\mathbf{x}_{i \leftarrow (\mathbf{x}[i] - g)}\|_1)} \\
&= \frac{f(\mathbf{x})}{\sum_{g=0}^{\mathbf{x}[i]} f(\mathbf{x}_{i \leftarrow (\mathbf{x}[i] - g)}) \exp(-\gamma g \|\mathbf{v}_i\|_1)} \\
&= \frac{1}{1 + \sum_{g=1}^{\mathbf{x}[i]} \frac{f(\mathbf{x}_{i \leftarrow (\mathbf{x}[i] - g)})}{f(\mathbf{x})} \exp(-\gamma g \ell)} \\
&\geq 1 - \sum_{g=1}^{\mathbf{x}[i]} \frac{f(\mathbf{x}_{i \leftarrow (\mathbf{x}[i] - g)})}{f(\mathbf{x})} \exp(-\gamma g \ell) && \left(\frac{1}{1+a} \geq 1 - a \text{ for } a \geq 0\right) \\
&= 1 - \sum_{g=1}^{\mathbf{x}[i]} \frac{\binom{k-g}{\mathbf{x}[i]-g}}{\binom{k}{\mathbf{x}[i]} p_i^g} \exp(-\gamma g \ell) \\
&\geq 1 - \sum_{g=1}^{\mathbf{x}[i]} p_i^{-g} \exp(-\gamma g \ell) \\
&= 1 - \sum_{g=1}^{\mathbf{x}[i]} (p_i \exp(\gamma \ell))^{-g} \\
&\geq 1 - \sum_{g=1}^{\infty} (p_i \exp(\gamma \ell))^{-g} \\
&= 1 - \frac{p_i^{-1} \exp(-\gamma \ell)}{1 - p_i^{-1} \exp(-\gamma \ell)} && \text{(for } p_i^{-1} \exp(-\gamma \ell) < 1) \\
&\geq 1 - 2(\min_i p_i)^{-1} \exp(-\gamma \ell) && \text{(for } p_i^{-1} \exp(-\gamma \ell) < 1/2) \\
&= 1 - R \exp(-\gamma \ell) && \text{(C.2)}
\end{aligned}$$

for  $R \triangleq 2(\min_i p_i)^{-1}$ . With this,

$$\begin{aligned}
 \frac{Q_\gamma(\mathcal{X}, \bar{\mathcal{X}})}{\tilde{\pi}_\gamma(\mathcal{X})} &= \frac{\sum_{\mathbf{x} \in \mathcal{X}, \mathbf{x}' \in \bar{\mathcal{X}}_{\mathbf{x}}} \tilde{\pi}_\gamma(\mathbf{x}) P_\gamma(\mathbf{x}, \mathbf{x}')}{\tilde{\pi}_\gamma(\mathcal{X})} \\
 &= \frac{\sum_{\mathbf{x} \in \mathcal{X}} \tilde{\pi}_\gamma(\mathbf{x}) \sum_{\mathbf{x}' \in \bar{\mathcal{X}}_{\mathbf{x}}} P_\gamma(\mathbf{x}, \mathbf{x}')}{\tilde{\pi}_\gamma(\mathcal{X})} \\
 &= \frac{\sum_{\mathbf{x} \in \mathcal{X}} \tilde{\pi}_\gamma(\mathbf{x}) \sum_{i \in [n]} \sum_{\mathbf{x}' \in \bar{\mathcal{X}}_{\mathbf{x}} \cap \Delta(\mathbf{x}, i)} \frac{\tilde{\pi}_\gamma(\mathbf{x}')}{2n \tilde{\pi}_\gamma(\Delta(\mathbf{x}, i))}}{\tilde{\pi}_\gamma(\mathcal{X})} \\
 &= \frac{\sum_{\mathbf{x} \in \mathcal{X}} \tilde{\pi}_\gamma(\mathbf{x}) \sum_{i \in [n]} \frac{\tilde{\pi}_\gamma(\Delta(\mathbf{x}, i)) - \tilde{\pi}_\gamma(\mathbf{x})}{2n \tilde{\pi}_\gamma(\Delta(\mathbf{x}, i))}}{\tilde{\pi}_\gamma(\mathcal{X})} \\
 &= \frac{\sum_{\mathbf{x} \in \mathcal{X}} \tilde{\pi}_\gamma(\mathbf{x}) \frac{1}{2n} \sum_{i \in [n]} \left(1 - \frac{\tilde{\pi}_\gamma(\mathbf{x})}{\tilde{\pi}_\gamma(\Delta(\mathbf{x}, i))}\right)}{\tilde{\pi}_\gamma(\mathcal{X})} \\
 &\leq \frac{\sum_{\mathbf{x} \in \mathcal{X}} \tilde{\pi}_\gamma(\mathbf{x}) \frac{1}{2n} \sum_{i \in [n]} R \exp(-\gamma \ell)}{\tilde{\pi}_\gamma(\mathcal{X})} && \text{(by (C.2))} \\
 &= \frac{\sum_{\mathbf{x} \in \mathcal{X}} \tilde{\pi}_\gamma(\mathbf{x}) \frac{1}{2} R \exp(-\gamma \ell)}{\tilde{\pi}_\gamma(\mathcal{X})} \\
 &= \frac{\tilde{\pi}_\gamma(\mathcal{X}) R \exp(-\gamma \ell)}{2 \tilde{\pi}_\gamma(\mathcal{X})} \\
 &= \frac{R \exp(-\gamma \ell)}{2}.
 \end{aligned}$$

As a result,

$$\Phi(M_\gamma) \leq \frac{Q_\gamma(\mathcal{X}, \bar{\mathcal{X}})}{\tilde{\pi}_\gamma(\mathcal{X})} \leq \frac{R \exp(-\gamma \ell)}{2}. \tag{C.3}$$

Next, we lower-bound  $\lambda_2(P_\gamma)$  with Cheeger's inequality (Jerrum and Sinclair, 1988; Lawler and Sokal, 1988):

$$1 - 2\Phi(M_\gamma) \leq \lambda_2 \tag{C.4}$$

which yields

$$\lambda_2 \geq 1 - \exp(-\ell\gamma)R. \tag{C.5}$$

Finally, we bound the relaxation time:

$$\begin{aligned}
 \tau_{\text{rel}, \gamma} &= \frac{1}{1 - \lambda_2} \\
 &\geq \frac{1}{1 - (1 - \exp(-\ell\gamma)R)} \\
 &= \frac{1}{\exp(-\ell\gamma)R} \\
 &= \Omega(\ell\gamma).
 \end{aligned}$$

□

Taken together, Lemmas 2 and 3 describe the trade-off in our choice of  $\gamma$ : for small values of  $\gamma$ , samples from  $\tilde{\pi}_\gamma$  rarely fall in  $\mathcal{X}$ , making rejection sampling inefficient. For large values of  $\gamma$ , the mixing time of the simple chain increases exponentially, requiring many iterations of the simple chain to generate each sample. These results tell us that the optimal choice of  $\gamma$  is finite.

**Lemma 4.** For every instance, there is some finite  $\gamma$  that minimizes  $N_\gamma$ .

*Proof.* Using (4.8) and lemma 2,

$$\lim_{\gamma \rightarrow -\infty} N_\gamma \geq \lim_{\gamma \rightarrow -\infty} \frac{N_\gamma}{\gamma} = \lim_{\gamma \rightarrow -\infty} \frac{(\tau_{\text{rel},\gamma} - 1) \log\left(\frac{3}{4\varepsilon p_\gamma^*}\right)}{p_\gamma^* \left(1 + \frac{2\varepsilon}{3}\right)} = \infty.$$

Using (4.8) and lemma 3,

$$\lim_{\gamma \rightarrow \infty} N_\gamma \geq \lim_{\gamma \rightarrow \infty} \frac{N_\gamma}{\gamma} = \lim_{\gamma \rightarrow \infty} \frac{(\tau_{\text{rel},\gamma} - 1) \log\left(\frac{3}{4\varepsilon p_\gamma^*}\right)}{p_\gamma^* \left(1 + \frac{2\varepsilon}{3}\right)} = \lim_{\gamma \rightarrow \infty} \Omega(\exp(\ell\gamma)) = \infty.$$

Since  $N_\gamma$  is finite for any finite  $\gamma$ , it must be minimized by some finite  $\gamma$ .  $\square$

#### APPENDIX D. OMITTED PROOFS AND DETAILS FROM SECTION 4.2

---

##### Algorithm 3 REDUCED( $\mathbf{V}, \mathbf{c}, f, k, t, \mathbf{x}_0$ )

---

```

 $\mathbf{x} \leftarrow \mathbf{x}_0$ 
for  $t$  iterations do
  with probability 0.5
    continue
  end
   $\mathbf{z} \leftarrow$  random  $k$  items from  $\mathbf{x}$  without replacement
  with probability  $1 - 1/\prod_i \binom{\mathbf{x}^{[i]}}{\mathbf{z}^{[i]}}$ 
    continue
  end
   $\mathcal{Z}(\mathbf{z}) \leftarrow \{\mathbf{z}' : \mathbf{V}\mathbf{z} = \mathbf{V}\mathbf{z}'\}$   $\triangleright$  Using ILP; can be cached for efficiency
   $\mathcal{S} \leftarrow \emptyset$ 
  for  $\mathbf{z}' \in \mathcal{Z}(\mathbf{z})$  do
     $\mathbf{x}' \leftarrow \mathbf{x} - \mathbf{z} + \mathbf{z}'$ 
    Add  $\mathbf{x}'$  to  $\mathcal{S}$ 
  end for
   $\mathbf{x} \leftarrow$  a random  $\mathbf{x} \in \mathcal{S}$  according to  $\pi \mid \mathbf{x} \in \mathcal{S}$  (i.e., proportional to  $f(\cdot)$  on  $\mathcal{S}$ )
end for
return  $\mathbf{x}$ 

```

---

##### D.1. An algorithm based on the reduced chain.

**Choosing an initial state  $\mathbf{x}_0$ .** In order to guarantee that the initial state for the reduced chain  $\mathbf{x}_0$  satisfies  $\pi(\mathbf{x}_0) \geq 1/|\mathcal{X}|$ , we use the fact that ILP solvers like Gurobi support maximizing linear objectives subject to a linear constraint. While our chosen  $f(\cdot)$  from (2.1) is not linear in  $\mathbf{x}$ , we can instead use the following linear approximation to  $\log f$ :

$$L(\mathbf{x}) = \sum_{i=1}^n \mathbf{x}[i] \log \Pr_{\mathcal{D}}[\mathbf{v}_i]. \tag{D.1}$$

Note that to obtain  $L$ , we dropped the multinomial coefficient in (2.1). To obtain an initial  $\mathbf{x}_0$  with large  $\pi(\mathbf{x}_0)$ , we use Gurobi to enumerate the  $N$  solutions with the largest values of  $L$  and, of those, choose the one that maximizes  $f$ . Gurobi additionally provides a bound on the largest objective value of any solution not yet enumerated, which we can translate into a bound on  $f$ . This allows us to guarantee that our chosen initial  $\mathbf{x}_0$  satisfies  $\pi(\mathbf{x}_0) \geq 1/|\mathcal{X}|$ .

Finally, we remark on hyperparameters  $k$  and  $t$ . As discussed in Section 4.2.2,  $k = 3$  appears to be a good choice. For  $t$ , experiments with Alabama and Nevada suggest  $t \approx 10^5$  suffices; however, given that AL and NV have different characteristics, it is possible that other states will require larger values of  $t$ .

**D.2. Omitted proofs from Section 4.2.3.** We restate each result from Section 4.2.3 before proving it.

**Lemma 5.**  $\log |\mathcal{X}| < \log |\mathcal{Y}| \leq m(1 + \log(n + 1))$ .

*Proof.* We take inspiration from a long line of work that bounds the number of feasible solutions to a knapsack problem (Begeed-Dov, 1972; Padberg, 1971; Achou, 1974; Lambe, 1974, 1977; Hujter, 1988; Mahmoudvand et al., 2010). The first inequality is trivially true since  $\mathcal{X} \subset \mathcal{Y}$ . We use the following combinatorial identity (see, e.g., Feller (1968, eq. (5.2))). Let  $A_{r,n}$  be the number of distinct tuples  $\mathbf{r} \in \mathbb{Z}_{\geq 0}^n$  such that  $\sum_{i=1}^n \mathbf{r}[i] = r$ . Then,

$$A_{r,n} = \binom{n+r-1}{r}. \tag{D.2}$$

Let  $j$  be the index such that  $\mathbf{v}_i[j] = 1$  for all  $i$ , which exists by assumption since  $\mathbf{c}$  encodes the number of elements in each solution. Observe that every  $\mathbf{x} \in \mathcal{Y}$  satisfies

$$\begin{aligned} \mathbf{V}\mathbf{x} &\leq \mathbf{c} \\ (\mathbf{V}\mathbf{x})[j] &\leq \mathbf{c}[j] \\ \sum_{i=1}^n \mathbf{x}[i] \mathbf{v}_i[j] &\leq m \\ \sum_{i=1}^n \mathbf{x}[i] &\leq m \end{aligned}$$

This means that each  $\mathbf{x} \in \mathcal{Y}$  satisfies  $\sum_{i=1}^n \mathbf{x}[i] = \ell$  for some  $\ell \in \{0, \dots, m\}$ . Thus, by (D.2),

$$\begin{aligned}
|\mathcal{Y}| &\leq \sum_{\ell=0}^m A_{\ell,n} \\
&= \sum_{\ell=0}^m \binom{n+\ell-1}{\ell} \\
&= \sum_{\ell=0}^m \binom{n+\ell-1}{n-1} \\
&= \binom{n+m}{n} && \text{(hockey stick identity)} \\
&= \binom{n+m}{m}.
\end{aligned}$$

Therefore,

$$\begin{aligned}
\log |\mathcal{Y}| &\leq \log \binom{n+m}{m} \\
&\leq \log \left( \frac{e(n+m)}{m} \right)^m \\
&= m \log \left( \frac{e(n+m)}{m} \right) \\
&\leq m \left( 1 + \log \left( \frac{n}{m} + 1 \right) \right) \\
&\leq m (1 + \log(n+1)).
\end{aligned}$$

The second line is a standard upper bound on binomial coefficients and follows from the fact that  $m! \geq \frac{m^m}{e^{m-1}}$ .  $\square$

**Lemma 6.** If  $\mathcal{V} = \mathcal{A}$ , then  $M_k$  is irreducible for  $k \geq 2$ .

*Proof.* We will show by induction that for any  $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$ ,  $\mathbf{x}$  can be transformed to  $\mathbf{x}'$  through a sequence of transitions under  $P_{k=2}$ . Because we're considering the case where  $\mathcal{V}$  is the intersection between a lattice and a hyperrectangle, it suffices to consider each dimension independently. Thus, we proceed under the assumption that  $d = 1$ . Let  $\mathbf{y}$  and  $\mathbf{y}'$  respectively be the matrix representations of  $\mathbf{x}$  and  $\mathbf{x}'$  obtained by listing all of their elements  $\mathbf{v}_i$  in order, with repeats for multiplicity. Thus,  $\mathbf{y}, \mathbf{y}' \in \mathbb{Z}_{\geq 0}^m$ . Note that this is a one-to-one mapping.

We proceed by induction on  $s \in \{1, \dots, m\}$ . Assume  $\mathbf{y}[i] = \mathbf{y}'[i]$  for all  $i < s$ . We will show the existence of a path from  $\mathbf{x}$  to some  $\mathbf{x}^+$  such that the corresponding  $\mathbf{y}^+[i] = \mathbf{y}'[i]$  for all  $i \leq s$ . If  $\mathbf{y}[s] = \mathbf{y}'[s]$ , we are done; otherwise, assume without loss of generality that  $\mathbf{y}[s] < \mathbf{y}'[s]$ .

Let  $j$  be an index such that  $j > s$  and  $\mathbf{y}[j] > \mathbf{y}'[j]$ . This must exist since  $\|\mathbf{y}\|_1 = \mathbf{V}\mathbf{x} = \mathbf{V}\mathbf{x}' = \|\mathbf{y}'\|_1$ ,  $\mathbf{y}[s]' > \mathbf{y}[s]$ , and  $\mathbf{y}[i] = \mathbf{y}'[i]$  for all  $i < s$ . We will construct a valid solution  $\mathbf{x}''$  by removing  $(\mathbf{y}[s], \mathbf{y}[j])$  from  $\mathbf{x}$  and replacing them with  $(\mathbf{y}[s] + 1, \mathbf{y}[j] - 1)$ . By convexity, these two new elements are both in  $\mathcal{V}$ . Repeating this argument produces a path from  $\mathbf{x}$  to  $\mathbf{x}^+$  where  $\mathbf{x}^+[i] = \mathbf{x}'[i]$  for all  $i \leq s$  as desired.  $\square$

**Lemma 7.** Let  $\{\mathcal{C}_i\}_{i=1}^c$  be the connected components of the undirected graph induced by  $P_k$ . Given initial state  $\mathbb{X}_0 \sim \pi_0$ , Algorithm 3 converges to  $\pi^+(\pi_0) \triangleq \lim_{t \rightarrow \infty} \pi_0 P_k^t$  satisfying

$$\pi^+(\pi_0) \in \arg \min_{\pi': \pi'(\mathcal{C}_i) = \pi_0(\mathcal{C}_i) \ \forall i \in [c]} d_{\text{TV}}(\pi, \pi').$$

In other words,  $\pi^+(\pi_0)$  is as close as possible to  $\pi$  subject to the constraint that it puts the same probability mass in each connected component that  $\pi_0$  does.

*Proof.* For clarity of notation, we will write  $\pi^+$  instead of  $\pi^+(\pi_0)$ . Let  $\mathcal{C}(\mathbf{x}_0)$  be the partition to which  $\mathbf{x}_0$  belongs, i.e.,  $\mathbf{x}_0 \in \mathcal{C}(\mathbf{x}_0)$ . For a connected component  $\mathcal{C}$ , consider the Markov chain  $M_k(\mathcal{C})$  defined over the states in  $\mathcal{C}$ .  $M_k(\mathcal{C}(\mathbf{x}_0))$  is lazy and irreducible and  $\pi$  satisfies the detailed balance equations. As a result, the stationary distribution of  $M_k(\mathcal{C}(\mathbf{x}_0))$  is  $\pi \mid \mathbf{x} \in \mathcal{C}(\mathbf{x}_0)$ . Because  $\mathbf{x}_0$  is sampled from  $\pi_0$ , we can write  $\pi^+$  as a linear combination of conditional distributions:

$$\pi^+ = \sum_{\mathbf{x} \in \mathcal{X}} \pi_0(\mathbf{x}) \cdot (\pi \mid \mathbf{x} \in \mathcal{C}(\mathbf{x})) = \sum_{i \in [c]} \pi_0(\mathcal{C}_i) \cdot (\pi \mid \mathbf{x} \in \mathcal{C}_i). \quad (\text{D.3})$$

Therefore,

$$\begin{aligned} d_{\text{TV}}(\pi^+, \pi) &= \frac{1}{2} \sum_{\mathbf{x} \in \mathcal{X}} |\pi^+(\mathbf{x}) - \pi(\mathbf{x})| \\ &= \frac{1}{2} \sum_{i \in [c]} \sum_{\mathbf{x} \in \mathcal{C}_i} |\pi^+(\mathbf{x}) - \pi(\mathbf{x})| \\ &= \frac{1}{2} \sum_{i \in [c]} \sum_{\mathbf{x} \in \mathcal{C}_i} |\pi_0(\mathcal{C}_i) \cdot (\pi \mid \mathbf{x} \in \mathcal{C}_i)(\mathbf{x}) - \pi(\mathbf{x})| \quad (\text{by (D.3)}) \\ &= \frac{1}{2} \sum_{i \in [c]} \sum_{\mathbf{x} \in \mathcal{C}_i} \left| \frac{\pi(\mathbf{x}) \pi_0(\mathcal{C}_i)}{\pi(\mathcal{C}_i)} - \pi(\mathbf{x}) \right| \\ &= \frac{1}{2} \sum_{i \in [c]} \sum_{\mathbf{x} \in \mathcal{C}_i} \pi(\mathbf{x}) \left| \frac{\pi_0(\mathcal{C}_i)}{\pi(\mathcal{C}_i)} - 1 \right| \\ &= \frac{1}{2} \sum_{i \in [c]} \pi(\mathcal{C}_i) \left| \frac{\pi_0(\mathcal{C}_i)}{\pi(\mathcal{C}_i)} - 1 \right| \\ &= \frac{1}{2} \sum_{i \in [c]} |\pi_0(\mathcal{C}_i) - \pi(\mathcal{C}_i)|. \quad (\text{D.4}) \end{aligned}$$

For any distribution  $\pi'$  such that  $\pi'(\mathcal{C}_i) = \pi_0(\mathcal{C}_i)$  for all  $i \in [c]$ ,

$$\begin{aligned}
d_{\text{TV}}(\pi', \pi) &= \frac{1}{2} \sum_{\mathbf{x} \in \mathcal{X}} |\pi'(\mathbf{x}) - \pi(\mathbf{x})| \\
&= \frac{1}{2} \sum_{i \in [c]} \sum_{\mathbf{x} \in \mathcal{C}_i} |\pi'(\mathbf{x}) - \pi(\mathbf{x})| \\
&\geq \frac{1}{2} \sum_{i \in [c]} \left| \sum_{\mathbf{x} \in \mathcal{C}_i} \pi'(\mathbf{x}) - \pi(\mathbf{x}) \right| \\
&= \frac{1}{2} \sum_{i \in [c]} |\pi'(\mathcal{C}_i) - \pi(\mathcal{C}_i)| \\
&= \frac{1}{2} \sum_{i \in [c]} |\pi_0(\mathcal{C}_i) - \pi(\mathcal{C}_i)| && \text{(by assumption)} \\
&= d_{\text{TV}}(\pi^+, \pi) && \text{(by (D.4))}
\end{aligned}$$

as claimed.  $\square$

## APPENDIX E. ADDITIONAL EMPIRICAL RESULTS

**Computation time.** To compare the computation time per iteration of the simple and reduced chains, we select a random 20 blocks from Alabama and perform  $10^5$  MCMC iterations of each method. We choose blocks such that:

- $|\mathcal{X}| > 1$  (some blocks have no solution with our formulation, which we ignore; in practice, we can solve these by matching a subset of statistics. See Appendix A for details.),
- $m > 3$  (since we will be running the reduce chain with  $k = 3$ , and the reduced chain is trivial when  $m \leq k$ ).

For each block, we run the simple chain with  $\gamma = 0.8$  and the reduced chain with  $k = 3$  for  $10^5$  iterations. Note that  $\gamma$  is unlikely to affect computation time; it primarily influences mixing time. We find that  $10^5$  iterations requires on average 10.22 seconds (std. dev. 2.71) and 8.23 seconds (std. dev. 6.29) for the simple and reduced chains respectively. Experiments are run on an Apple M1 MacBook Pro with 32GB memory. For reference, enumerating 5000 solutions via ILP for 20 randomly selected Alabama blocks (with at least 5000 solutions each) takes 4.58 seconds (std. dev. 0.95) on the same hardware. As a result, running the reduced chain on blocks with  $|\mathcal{X}_b| \geq 5000$  does not dramatically increase computation time.

As a final point of reference, running Algorithm 4 on a computing cluster with  $N = 5,000$ ,  $k = 3$ , and  $t = 10^5$  takes around 21 days (Alabama) and 37 days (Nevada) of 12-threaded CPU time to sample a dataset for the entire state. This is highly parallelizable since our algorithm operates independently on each block.




FIGURE 6. Optimal choice of  $\gamma$  vs.  $m$ . Marker size is proportional to the number of such blocks.




FIGURE 7. Optimal choice of  $k$  vs.  $m$ . Marker size is proportional to the number of such blocks.

**Optimal choices of  $\gamma$  and  $k$ .** Figure 6 shows the relationship between  $m$  and the  $\gamma^*$  that minimizes  $\underline{N}_{\gamma^*}$ . There are a few blocks for which  $\gamma = 1.2$  appears to be optimal, but this is often because low-precision estimates for  $\lambda_2(P_{\gamma=1.2})$  lead to underestimates of  $\underline{N}_{\gamma=1.2}$ . Computing  $\lambda_2(P_\gamma)$  is computationally very expensive for larger values of  $\gamma$ , which is why we don't attempt to evaluate  $\gamma > 1.2$ . Figure 7 provides analogous results for the reduced chain, finding that  $k = 4$  appears to be optimal for smaller blocks, but as blocks get larger, smaller values of  $k$  are better.

**Disconnected graphs.** Finally, we consider instances where the reduced chain is not irreducible for  $k = 2$ . Figure 8 shows one such example, where we draw the (undirected) graph given by one-step transitions using  $P_{k=2}$ . Each node in the graph is some  $\mathbf{x} \in \mathcal{X}$ . The graph has 3 connected components, and crucially, increasing to  $k = 3$  adds many crossing




FIGURE 8. Connected components in the state space of the reduced chain with  $k = 2$  for a particular block in Alabama (Madison County, tract 2501, block 2053). Dashed purple lines show additional crossing edges present with  $k = 3$ , making the resulting graph connected.

edges (shown with purple dashes). Other examples we examined are qualitatively similar, though the density of additional connections with  $k = 3$  varies.

#### APPENDIX F. HARDNESS RESULTS

Existing hardness results for multidimensional knapsack and subset sum do not strictly apply here since (1) we are interested in the decision version, and (2) our solutions can be multisets. However, standard 3SAT reductions to subset sum can be easily adapted to our multidimensional multiset case.

**Claim 1.** **MMS** is NP-hard.

*Proof.* Given a 3SAT instance with variables  $y_1, \dots, y_n$  and clauses  $c_1, \dots, c_m$ , we can construct a **MMS** instance with  $2n + 3m$  elements and dimension  $n + m$  using a slight variant of the standard 3SAT reduction to subset sum. We use the following encoding. For

$i \in [n]$ ,

$$\begin{aligned} \mathbf{v}_{2(i-1)}[j] &= \begin{cases} 1 & j \leq n \wedge i = j \\ 1 & j > n \wedge c_{j-n} \text{ contains } y_i \\ 0 & \text{otherwise} \end{cases} \\ \mathbf{v}_{2i-1}[j] &= \begin{cases} 1 & j \leq n \wedge i = j \\ 1 & j > n \wedge c_{j-n} \text{ contains } \bar{y}_i \\ 0 & \text{otherwise} \end{cases} \\ \mathbf{v}_{2i}[j] &= \begin{cases} 4 & j = 2n - 1 + i \\ 0 & \text{otherwise} \end{cases} \\ \mathbf{v}_{2i+1}[j] &= \begin{cases} 5 & j = 2n + i \\ 0 & \text{otherwise} \end{cases} \\ \mathbf{v}_{2i+2}[j] &= \begin{cases} 6 & j = 2n + 1 + i \\ 0 & \text{otherwise} \end{cases} \end{aligned}$$

Intuitively, the  $\mathbf{v}_i$  for  $i \leq 2n$  represent the variables, with one variable for  $y_i$  and one for  $\bar{y}_i$ . The remaining  $3m$  vectors are slack vectors for the clauses, corresponding to the fact that clause  $j$  can be satisfied by 1, 2, or 3 of its variables. Let the counts vector be

$$\mathbf{c}[j] = \begin{cases} 1 & j \leq n \\ 7 & j > n \end{cases}.$$

We will show that this **MMS** instance is feasible if and only if the original 3SAT instance was feasible.

**3SAT solution**  $\implies$  **MMS solution**. Let  $Y_i^*$  be a 3SAT solution. Let  $z_j$  be the number of ways in which  $c_j$  is satisfied by  $Y_i^*$ . Note that  $z_j \in \{1, 2, 3\}$ .

$$\mathbf{x}[j] = \begin{cases} 1 & j \leq 2n \wedge j \text{ is odd} \wedge Y_{(j+1)/2}^* = \text{true} \\ 1 & j \leq 2n \wedge j \text{ is even} \wedge Y_{j/2}^* = \text{false} \\ 1 & j > 2n \wedge j \pmod 3 = 1 \wedge z_{(j+2)/3} = 3 \\ 1 & j > 2n \wedge j \pmod 3 = 2 \wedge z_{(j+1)/3} = 2 \\ 1 & j > 2n \wedge j \pmod 3 = 3 \wedge z_{j/3} = 1 \\ 0 & \text{otherwise} \end{cases}$$

Observe that  $\mathbf{V}\mathbf{x} = \mathbf{c}$  because:

- For  $j \leq n$ ,  $(\mathbf{V}\mathbf{x})[j] = \mathbf{c}[j]$ . This is because exactly one of  $\mathbf{x}[2i-1], \mathbf{x}[2i]$  can be 1 for  $i \leq n$  by construction.
- For  $j > n$ ,  $(\mathbf{V}\mathbf{x})[j] = \mathbf{c}[j]$ . This is because for each clause  $\ell$ , the first  $2n$  columns of  $\mathbf{V}$  contribute exactly  $z_\ell$  to  $\mathbf{c}[\ell+2n]$ . The remaining  $3m$  columns contribute  $7 - z_\ell$ .

**MMS solution**  $\implies$  **3SAT solution**. Given a solution  $\mathbf{x}$ , we assign variables as follows.  $y_i$  is set to true if  $\mathbf{x}[2i-1] = 1$  and false if not. We will show that for  $j > n$ ,  $\mathbf{V}\mathbf{x}[j] = \mathbf{c}[j] \implies$  clause  $c_{j-n}$  is satisfied. To see this, note that no multiset consisting of only the last  $3m$  columns of  $\mathbf{V}$  can sum to 7, since they only contain the values 4, 5, and 6. Thus,  $\mathbf{x}$  must select at least one column from the first  $2n$  that has a nonzero entry in position  $j$ . The only such columns correspond to variables satisfying  $c_{j-n}$  (either  $\mathbf{v}_{2j-1}$  or  $\mathbf{v}_{2j}$ ).  $\square$

## APPENDIX G. PATHOLOGICAL EXAMPLES FOR THE REDUCED CHAIN

**Example 2** (Disconnected reduced chain). In the following example, the reduced chain is disconnected for  $k = 2$ .

$$\mathbf{V} = \begin{pmatrix} 3 & 0 & 0 & 1 \\ 0 & 3 & 0 & 1 \\ 0 & 0 & 3 & 1 \\ 1 & 1 & 1 & 1 \end{pmatrix} \quad \mathbf{c} = \begin{pmatrix} 3 \\ 3 \\ 3 \\ 3 \end{pmatrix}$$

Solutions:

$$\mathcal{X} = \left\{ \begin{pmatrix} 1 \\ 1 \\ 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 0 \\ 0 \\ 3 \end{pmatrix} \right\}$$

To move between the two solutions in  $\mathcal{X}$ , we must remove and replace 3 elements at a time, meaning the reduced chain is disconnected for  $k = 2$ . This example can be generalized to be disconnected for any  $k < m$ .

**Example 3** (Connected reduced chain with high mixing time).

Here, we provide an instance  $(\mathbf{V}, \mathbf{c}, \pi)$  such that the reduced chain with  $k = 2$  is connected but has exponentially large mixing time. Let  $\pi$  assign equal probability to all solutions, i.e.,  $\pi(\mathbf{x}) = 1/|\mathcal{X}|$ . Let  $\mathcal{I}_\ell$  be the set of  $\ell$ -dimensional unit basis vectors, i.e., each  $\mathbf{e}_i \in \mathcal{I}_\ell$  has a 1 in the  $i$ th position and 0 elsewhere. Let  $\mathbf{M}_\ell \in \mathbb{Z}_{\geq 0}^{\ell \times (\ell + \binom{\ell}{2})}$  be the matrix formed by taking as columns all possible sums of two vectors (with replacement) from  $\mathcal{I}_\ell$  and multiplying them by  $\ell$ . For example,

$$\mathbf{M}_3 = \begin{pmatrix} 6 & 0 & 0 & 3 & 3 & 0 \\ 0 & 6 & 0 & 3 & 0 & 3 \\ 0 & 0 & 6 & 0 & 3 & 3 \end{pmatrix}.$$

Let  $\mathbf{S}_\ell \in \mathbb{Z}_{\geq 0}^{\ell \times (\ell-2)}$  be the matrix where column  $i \in \{1, \dots, \ell-2\}$  has as its first  $i+1$  entries  $2\ell$  and 0 elsewhere. For example,

$$\mathbf{S}_3 = \begin{pmatrix} 6 \\ 6 \\ 0 \end{pmatrix} \quad \mathbf{S}_4 = \begin{pmatrix} 8 & 8 \\ 8 & 8 \\ 0 & 8 \\ 0 & 0 \end{pmatrix}.$$

Let  $\mathbf{T}_\ell \in \mathbb{Z}_{\geq 0}^{\ell \times (2\ell-1)}$  be the matrix where the  $i$ th column has  $i+1$  everywhere. For example,

$$\mathbf{T}_3 = \begin{pmatrix} 2 & 3 & 4 & 5 & 6 \\ 2 & 3 & 4 & 5 & 6 \\ 2 & 3 & 4 & 5 & 6 \end{pmatrix}.$$



**The chain is irreducible.** Observe that  $\mathcal{X}_L$  is connected to  $\mathcal{X}_R$  by swaps of size  $k = 2$ : intuitively, starting at a state in  $\mathcal{X}_L$ , we can successively replace elements of  $\mathcal{V}_L$  with elements from  $\mathbf{S}_\ell$  and the vector  $[0 \dots 0 1 \dots 1]$  until we reach the state with 1 copy of  $[2\ell \dots 2\ell 1 \dots 1]$  and  $2\ell - 1$  copies of  $[0 \dots 0 1 \dots 1]$ . From there, we can use elements from  $\mathbf{T}_\ell$  to reach the solution that consists of  $2\ell$  copies of  $[1 \dots 1]$ . Symmetrically, this state is reachable from  $\mathcal{X}_R$ , meaning the two are connected. More generally, from any solution, by construction the solution consisting of  $2\ell$  copies of  $[1 \dots 1]$  is reachable.

To bound conductance, we will show that:

- (1) only one state  $\mathbf{x}^* \in \mathcal{X}_L$  has edges to states in  $\mathcal{X} \setminus \mathcal{X}_L$ ,
- (2)  $\pi(\mathbf{x}^*)/\pi(\mathcal{X}_L) = 1/|\mathcal{X}_L|$  is small.

**Only one state has crossing edges.** Let  $\mathbf{x}^* \in \mathcal{X}_L$  have an edge to some  $\mathbf{x}' \notin \mathcal{X}_L$ . Treating  $\mathbf{x}^*$  as a multiset, we will reason about the number of copies of each vector  $\mathbf{v}$  it contains. We will show that  $\mathbf{x}^*$  has 1 copy of  $[0 \dots 0 2\ell 1 \dots 1]$ , 1 copy of  $[2\ell \dots 2\ell 0 1 \dots 1]$ , and  $2\ell - 2$  copies of  $[0 \dots 0 1 \dots 1]$ .

Let  $\mathbf{v}_a$  and  $\mathbf{v}_b$  be the two elements that can be removed from  $\mathbf{x}^*$  and replaced by  $\mathbf{v}_c$  and  $\mathbf{v}_d$  to yield some  $\mathbf{x}' \notin \mathcal{X}_L$ . Let  $\mathbf{s} = \mathbf{v}_a + \mathbf{v}_b$ . Observe that the last  $\ell$  entries of  $\mathbf{s}$  are each 2 because  $\mathbf{x}^*$  only uses vectors in  $\mathcal{V}_L$ , which all have 1's in their last  $\ell$  entries. Further, observe that each  $\mathbf{v} \notin \mathcal{V}_L$  has its first  $\ell$  entries nonzero and identical. Without loss of generality, let  $\mathbf{v}_c \notin \mathcal{V}_L$ . In order for  $\mathbf{v}_a + \mathbf{v}_b = \mathbf{s} = \mathbf{v}_c + \mathbf{v}_d$ , it must be the case that each of the first  $\ell$  entries of  $\mathbf{s}$  is nonzero, since each of the first  $\ell$  entries of  $\mathbf{v}_c$  is nonzero. For sufficiently large  $\ell$  (i.e.,  $\ell > 4$ ), only three combinations of elements  $\mathbf{v}_a, \mathbf{v}_b \in \mathcal{V}_L$  have this property. Denoting the first  $\ell$  entries of a vector  $\mathbf{v}$  as  $\mathbf{v}[:\ell]$ , these three possible combinations are:

$$\begin{aligned} \mathbf{v}_a[:\ell] &= \begin{pmatrix} 0 \\ \vdots \\ 0 \\ \ell \\ \ell \end{pmatrix}, & \mathbf{v}_b[:\ell] &= \begin{pmatrix} 2\ell \\ \vdots \\ 2\ell \\ 0 \\ 0 \end{pmatrix} \\ \mathbf{v}_a[:\ell] &= \begin{pmatrix} 0 \\ \vdots \\ 0 \\ \ell \\ \ell \end{pmatrix}, & \mathbf{v}_b[:\ell] &= \begin{pmatrix} 2\ell \\ \vdots \\ 2\ell \\ 2\ell \\ 0 \end{pmatrix} \\ \mathbf{v}_a[:\ell] &= \begin{pmatrix} 0 \\ \vdots \\ 0 \\ 0 \\ 2\ell \end{pmatrix}, & \mathbf{v}_b[:\ell] &= \begin{pmatrix} 2\ell \\ \vdots \\ 2\ell \\ 2\ell \\ 0 \end{pmatrix} \end{aligned}$$

This yields possible sums for the first  $\ell$  entries:

$$(\mathbf{v}_a + \mathbf{v}_b)[:\ell] = \mathbf{s}[:\ell] \in \left\{ \begin{pmatrix} 2\ell \\ \vdots \\ 2\ell \\ \ell \\ \ell \end{pmatrix}, \begin{pmatrix} 2\ell \\ \vdots \\ 2\ell \\ 3\ell \\ \ell \end{pmatrix}, \begin{pmatrix} 2\ell \\ \vdots \\ 2\ell \\ 2\ell \\ 2\ell \end{pmatrix} \right\}$$

For the first 2 possibilities, it must be the case that the first  $\ell$  entries of  $\mathbf{v}_c$  are identical (because  $\mathbf{v}_c \notin \mathcal{V}_L$ ) and at most  $\ell$  (because otherwise  $\mathbf{v}_d = \mathbf{s} - \mathbf{v}_c$  would contain negative entries). But we can verify that by construction, no  $\mathbf{v}_d$  exists that yields the desired sum. Thus, the only possible sum is  $\mathbf{s} = [2\ell \dots 2\ell \ 2 \dots 2]$ . Given this choice of  $\mathbf{s}$ , the only solution that includes the required  $\mathbf{v}_a$  and  $\mathbf{v}_b$  must include these two elements and  $2\ell - 2$  copies of  $[0 \dots 0 \ 1 \dots 1]$ . This is  $\mathbf{x}^*$  as claimed.

**$\mathcal{X}_L$  contains many solutions.** Next, we lower-bound  $|\mathcal{X}_L|$ . A fairly straightforward argument shows  $|\mathcal{X}_L| \geq 2^{\lfloor \ell/2 \rfloor}$ . To see this, note that for each pair of indices  $i \neq j \leq \ell$ , there are two ways to make those entries  $2\ell$  and the rest of the first  $\ell$  entries 0:

$$\begin{pmatrix} 2\ell \\ 0 \end{pmatrix} + \begin{pmatrix} 0 \\ 2\ell \end{pmatrix} \qquad \begin{pmatrix} \ell \\ \ell \end{pmatrix} + \begin{pmatrix} \ell \\ \ell \end{pmatrix}$$

Because there are  $\lfloor \ell/2 \rfloor$  disjoint pairs of entries in the first  $\ell$ , there are at least  $2^{\lfloor \ell/2 \rfloor}$  distinct solutions.

A more sophisticated argument shows that  $|\mathcal{X}_L| = \Omega\left(\left(\frac{\ell}{e}\right)^\ell\right)$ . Let  $\mathcal{X}'_L$  be the set of solutions that only use the first  $\ell + \binom{\ell}{2}$  columns of  $\mathbf{V}_\ell$ . Since  $\mathcal{X}'_L \subseteq \mathcal{X}_L$ ,  $|\mathcal{X}'_L| \leq |\mathcal{X}_L|$ . Observe that each of the first  $\ell + \binom{\ell}{2}$  columns of  $\mathbf{V}_\ell$  can be indexed by  $(i, j)$  for  $i \leq j$ , where  $\mathbf{v}_{(i,j)}$  is  $\ell \times (\mathbf{e}_i + \mathbf{e}_j)$  for unit basis vectors  $\mathbf{e}_i$  and  $\mathbf{e}_j$ . (A unit basis vector  $\mathbf{e}_i$  has a 1 in position  $i$  and 0 elsewhere.) Any  $\mathbf{x} \in \mathcal{X}'_L$  can thus be expressed as a sequence of these ordered pairs.

Further, any  $\mathbf{x} \in \mathcal{X}'_L$  must have exactly two copies of  $\ell \times \mathbf{e}_i$  for each  $i \in [\ell]$ . For some solution  $\mathbf{x} \in \mathcal{X}'_L$ , observe that beginning at an arbitrary  $\mathbf{v}_{(i,j)}$  in the solution, we can trace out a cycle  $i \rightarrow j, j \rightarrow k, \dots, z \rightarrow i$  simply by finding the item that provides the matching pair for each  $\mathbf{e}_j$ . Thus, each  $\mathbf{x} \in \mathcal{X}'_L$  uniquely corresponds to a set of (undirected) cycles over  $\ell$  elements. (For example, the appearance of  $\mathbf{v}_{(i,i)}$  corresponds to a self-loop, and  $\mathbf{v}_{(i,j)}$  appearing twice corresponds to a cycle of length 2.) Thus,  $|\mathcal{X}'_L|$  is exactly the number of distinct such collections of cycles.

Defining  $a(\ell)$  to be the number of distinct collections of cycles over  $\ell$  elements, Stanley (1999, Example 5.2.9) shows that  $a(\ell)$  satisfies

$$\sum_{\ell \geq 0} a(\ell) \frac{x^\ell}{\ell!} = (1-x)^{-1/2} \exp\left(\frac{x}{2} + \frac{x^2}{4}\right).$$

The sequence  $a(0), a(1), \dots$  appears as OEIS sequence [A002135](#) (OEIS Foundation Inc., 2023), and satisfies the recurrence  $a(\ell) = \ell \cdot a(\ell - 1) - \binom{\ell-1}{2} a(\ell - 3)$ . Asymptotically (c.f. Pietro Majer (OEIS Foundation Inc., 2023)),

$$a(\ell) \sim \sqrt{2} \exp\left(\frac{3}{4}\right) \left(\frac{\ell}{e}\right)^\ell \left(1 + O\left(\frac{1}{\ell}\right)\right).$$

Thus,  $|\mathcal{X}_L| \geq |\mathcal{X}'_L| = a(\ell) = \Omega\left(\left(\frac{\ell}{e}\right)^\ell\right)$ .

**Conductance and Cheeger's inequality.** With this, we are ready to bound the chain's conductance and thus its relaxation time. We need the following definitions. For our Markov chain  $M_\ell$  with transition matrix  $P_\ell$ ,

$$\begin{aligned}
Q_\ell(\mathbf{x}, \mathbf{x}') &\triangleq \pi(\mathbf{x}')P_\ell(\mathbf{x}, \mathbf{x}') \\
Q_\ell(\mathcal{S}, \overline{\mathcal{S}}) &\triangleq \sum_{\mathbf{x} \in \mathcal{S}, \mathbf{x}' \notin \mathcal{S}} Q_\ell(\mathbf{x}, \mathbf{x}') \\
\Phi(\mathcal{S}, \overline{\mathcal{S}}) &\triangleq \frac{Q_\ell(\mathcal{S}, \overline{\mathcal{S}})}{\pi(\mathcal{S})} \\
\Phi(M_\ell) &\triangleq \min_{\mathcal{S}: 0 < \pi(\mathcal{S}) \leq 1/2} \Phi(\mathcal{S}, \overline{\mathcal{S}}) \\
&\leq \Phi(\mathcal{X}_L, \overline{\mathcal{X}_L}) \\
&= \frac{\sum_{\mathbf{x} \in \mathcal{X}_L, \mathbf{x}' \notin \mathcal{X}_L} \pi(\mathbf{x})P_\ell(\mathbf{x}, \mathbf{x}')}{\pi(\mathcal{X}_L)} \\
&\leq \frac{\sum_{\mathbf{x} \in \mathcal{X}_L, \mathbf{x}' \notin \mathcal{X}_L} \pi(\mathbf{x})}{\pi(\mathcal{X}_L)} && (P_\ell(\cdot, \cdot) \leq 1) \\
&= \frac{\pi(\mathbf{x}^*)}{\pi(\mathcal{X}_L)} && (\mathbf{x}^* \text{ is the only state with a crossing edge}) \\
&= \frac{1}{|\mathcal{X}_L|} && (\pi(\mathbf{x}) \propto 1) \\
&= O\left(\left(\frac{e}{\ell}\right)^\ell\right)
\end{aligned}$$

By Cheeger's inequality (Jerrum and Sinclair, 1988; Lawler and Sokal, 1988),

$$\lambda_2(P_\ell) \geq 1 - 2\Phi(M_\ell).$$

This means that

$$\tau_{\text{rel}}(P_\ell) = \frac{1}{1 - \lambda_2(P_\ell)} \geq \frac{1}{2\Phi(M_\ell)} = \Omega\left(\left(\frac{\ell}{e}\right)^\ell\right).$$

□

## APPENDIX H. FINAL ALGORITHM DESCRIPTION

Recall that our upper bounds on  $N_k$  relied on the fact that  $\pi(\mathbf{x}_0) \geq 1/|\mathcal{X}|$ . This may no longer be true when  $\mathbf{x}_0$  is sampled from  $\pi \mid \mathbf{x} \in \mathcal{S}$ . Here, we provide a bound that only depends on  $\pi(\mathcal{S})$ , not  $|\mathcal{X}|$ . As long as  $\pi(\mathcal{S}) \geq 1/|\mathcal{X}|$ , our new bound is stronger.

**Lemma 8.** If  $M_k$  is irreducible, then Algorithm 4 produces an  $\varepsilon$ -approximate sample from  $\pi$  after at most

$$\tau_{\text{mix}}(\varepsilon; P_k, \mathbb{X}) \leq \tau_{\text{rel},k} \log\left(\frac{1}{2\varepsilon\sqrt{\pi(\mathcal{S})}}\right)$$

---

**Algorithm 4** ILPANDMCMC( $\mathbf{V}, \mathbf{c}, f, N, k, t$ )
 

---

$\mathcal{S} \leftarrow$  subset of  $\mathcal{X}$  of size at most  $N$   $\triangleright$  Using ILP, with objective  $L(\cdot)$  given by (D.1)  
 $\mathbf{x} \leftarrow$  a random  $\mathbf{x} \in \mathcal{S}$  according to  $\pi \mid \mathbf{x} \in \mathcal{S}$   
**if**  $|\mathcal{S}| < N$  **then**  $\triangleright$  In this case,  $(\pi \mid \mathbf{x} \in \mathcal{S}) = \pi$   
     **return**  $\mathbf{x}$   
**else**  
     **return** REDUCED( $\mathbf{V}, \mathbf{c}, f, k, t, \mathbf{x}$ )  
**end if**

---

Markov chain iterations.

*Proof.* Let  $\mathbb{X} \sim \pi_0$  where  $\pi_0 = \pi \mid \mathbf{x} \in \mathcal{S}$ . Let  $P_k^t(\mathbb{X}, \cdot)$  be the distribution of over the reduced chain beginning with  $\mathbb{X}$  after  $t$  steps. Following [Sousi \(2020, Theorem 3.4\)](#), for a reversible Markov chain with transition matrix  $P_k$  and relaxation time  $\tau_{\text{rel}}$ ,

$$\|P_k^t(\mathbb{X}, \cdot) - \pi\|_{2,\pi} \leq \exp(-t/\tau_{\text{rel}}) \|\pi_0 - \pi\|_{2,\pi}, \quad (\text{H.1})$$

where given a stationary distribution  $\pi$ , we define

$$\|\sigma - \pi\|_{2,\pi} \triangleq \left( \sum_{\mathbf{x} \in \mathcal{X}} \left( \frac{\sigma(\mathbf{x})}{\pi(\mathbf{x})} - 1 \right)^2 \pi(\mathbf{x}) \right)^{1/2}.$$

For initial distribution  $\pi_0 = \pi \mid \mathbf{x} \in \mathcal{S}$ ,

$$\begin{aligned}
 \|\pi_0 - \pi\|_{2,\pi}^2 &= \sum_{\mathbf{x} \in \mathcal{X}} \left( \frac{\pi_0(\mathbf{x})}{\pi(\mathbf{x})} - 1 \right)^2 \pi(\mathbf{x}) \\
 &= \sum_{\mathbf{x} \in \mathcal{X}} \left( \frac{\pi_0(\mathbf{x})^2}{\pi(\mathbf{x})^2} - \frac{2\pi_0(\mathbf{x})}{\pi(\mathbf{x})} + 1 \right) \pi(\mathbf{x}) \\
 &= \sum_{\mathbf{x} \in \mathcal{X}} \frac{\pi_0(\mathbf{x})^2}{\pi(\mathbf{x})} - 2\pi_0(\mathbf{x}) + \pi(\mathbf{x}) \\
 &= \left( \sum_{\mathbf{x} \in \mathcal{X}} \frac{\pi_0(\mathbf{x})^2}{\pi(\mathbf{x})} \right) - 1 \\
 &= \left( \sum_{\mathbf{x} \in \mathcal{S}} \frac{\pi_0(\mathbf{x})^2}{\pi(\mathbf{x})} + \sum_{\mathbf{x} \notin \mathcal{S}} \frac{\pi_0(\mathbf{x})^2}{\pi(\mathbf{x})} \right) - 1 \\
 &= \left( \sum_{\mathbf{x} \in \mathcal{S}} \frac{(\pi(\mathbf{x})/\pi(\mathcal{S}))^2}{\pi(\mathbf{x})} \right) - 1 \\
 &= \left( \sum_{\mathbf{x} \in \mathcal{S}} \frac{\pi(\mathbf{x})}{\pi(\mathcal{S})^2} \right) - 1 \\
 &= \frac{1}{\pi(\mathcal{S})} - 1 \\
 &\leq \frac{1}{\pi(\mathcal{S})}.
 \end{aligned}$$

Thus, by (H.1),

$$\|P_k^t(\mathbb{X}, \cdot) - \pi\|_{2,\pi} \leq \exp(-t/\tau_{\text{rel}}) \frac{1}{\sqrt{\pi(\mathcal{S})}}. \quad (\text{H.2})$$

Next, observe that for any  $\sigma$ ,

$$\begin{aligned} \|\sigma - \pi\|_1^2 &= \left( \sum_{\mathbf{x} \in \mathcal{X}} |\sigma(\mathbf{x}) - \pi(\mathbf{x})| \right)^2 \\ &= \left( \sum_{\mathbf{x} \in \mathcal{X}} \left| \frac{\sigma(\mathbf{x})}{\pi(\mathbf{x})} - 1 \right| \pi(\mathbf{x}) \right)^2 \\ &\leq \sum_{\mathbf{x} \in \mathcal{X}} \left( \frac{\sigma(\mathbf{x})}{\pi(\mathbf{x})} - 1 \right)^2 \pi(\mathbf{x}) \quad (\text{Jensen's inequality}) \\ &= \|\sigma - \pi\|_{2,\pi}^2. \end{aligned}$$

Combining this with (H.2),

$$2d_{\text{TV}}(P_k^t(\mathbb{X}, \cdot), \pi) = \|P_k^t(\mathbb{X}, \cdot) - \pi\|_1 \leq \|P_k^t(\mathbb{X}, \cdot) - \pi\|_{2,\pi} \leq \exp(-t/\tau_{\text{rel}}) \frac{1}{\sqrt{\pi(\mathcal{S})}}.$$

To achieve the desired bound on  $d_{\text{TV}}(P_k^t(\mathbb{X}, \cdot), \pi)$ , it suffices to set

$$\begin{aligned} \frac{1}{2} \exp(-t/\tau_{\text{rel}}) \frac{1}{\sqrt{\pi(\mathcal{S})}} &\leq \varepsilon \\ \exp(-t/\tau_{\text{rel}}) &\leq 2\varepsilon \sqrt{\pi(\mathcal{S})} \\ t &\geq \tau_{\text{rel}} \log \left( \frac{1}{2\varepsilon \sqrt{\pi(\mathcal{S})}} \right). \end{aligned}$$

□