



HAL
open science

Smoothness Analysis for Probabilistic Programs with Application to Optimised Variational Inference

Wonyeol Lee, Xavier Rival, Hongseok Yang

► **To cite this version:**

Wonyeol Lee, Xavier Rival, Hongseok Yang. Smoothness Analysis for Probabilistic Programs with Application to Optimised Variational Inference. Proceedings of the ACM on Programming Languages, 2023, 7, pp.335 - 366. 10.1145/3571205 . hal-03936759

HAL Id: hal-03936759

<https://hal.science/hal-03936759>

Submitted on 12 Jan 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Smoothness Analysis for Probabilistic Programs with Application to Optimised Variational Inference

WONYEOL LEE, Stanford University, USA

XAVIER RIVAL, INRIA Paris, France and ENS, CNRS, and PSL University, Paris, France

HONGSEOK YANG, KAIST, South Korea and Institute for Basic Science (IBS), South Korea

We present a static analysis for discovering differentiable or more generally smooth parts of a given probabilistic program, and show how the analysis can be used to improve the pathwise gradient estimator, one of the most popular methods for posterior inference and model learning. Our improvement increases the scope of the estimator from differentiable models to non-differentiable ones without requiring manual intervention of the user; the improved estimator automatically identifies differentiable parts of a given probabilistic program using our static analysis, and applies the pathwise gradient estimator to the identified parts while using a more general but less efficient estimator, called score estimator, for the rest of the program. Our analysis has a surprisingly subtle soundness argument, partly due to the misbehaviours of some target smoothness properties when viewed from the perspective of program analysis designers. For instance, some smoothness properties, such as partial differentiability and partial continuity, are not preserved by function composition, and this makes it difficult to analyse sequential composition soundly without heavily sacrificing precision. We formulate five assumptions on a target smoothness property, prove the soundness of our analysis under those assumptions, and show that our leading examples satisfy these assumptions. We also show that by using information from our analysis instantiated for differentiability, our improved gradient estimator satisfies an important differentiability requirement and thus computes the correct estimate on average (i.e., returns an unbiased estimate) under a regularity condition. Our experiments with representative probabilistic programs in the Pyro language show that our static analysis is capable of identifying smooth parts of those programs accurately, and making our improved pathwise gradient estimator exploit all the opportunities for high performance in those programs.

CCS Concepts: • **Software and its engineering** → **Correctness; Automated static analysis**; • **Mathematics of computing** → **Bayesian computation; Variational methods**.

Additional Key Words and Phrases: smoothness, static analysis, probabilistic programming, variational inference

ACM Reference Format:

Wonyeol Lee, Xavier Rival, and Hongseok Yang. 2023. Smoothness Analysis for Probabilistic Programs with Application to Optimised Variational Inference. *Proc. ACM Program. Lang.* 7, POPL, Article 12 (January 2023), 32 pages. <https://doi.org/10.1145/3571205>

1 INTRODUCTION

Probabilistic programs define models from machine learning and statistics, and are used to analyse datasets from a wide range of applications [Bingham et al. 2019; Carpenter et al. 2017; Ge et al. 2018; Gehr et al. 2016; Goodman et al. 2008; Gordon et al. 2014; Mansinghka et al. 2014; Minka et al. 2014; Narayanan et al. 2016; Salvatier et al. 2016; Siddharth et al. 2017; Tolpin et al. 2016; Tran et al. 2018, 2016;

Authors' addresses: Wonyeol Lee, Computer Science, Stanford University, USA, wonyeol@cs.stanford.edu; Xavier Rival, INRIA Paris, France, Département d'Informatique, ENS, CNRS, and PSL University, Paris, France, rival@di.ens.fr; Hongseok Yang, School of Computing and Kim Jaechul Graduate School of AI, KAIST, South Korea, Discrete Mathematics Group, Institute for Basic Science (IBS), South Korea, hongseok.yang@kaist.ac.kr.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2023 Copyright held by the owner/author(s).

2475-1421/2023/1-ART12

<https://doi.org/10.1145/3571205>

Wood et al. 2014]. These programs are written in languages with special runtimes, called inference engines, which can be used to answer probabilistic queries, such as posterior inference and marginal likelihood estimation, or to learn model parameters in those programs, such as weights of neural networks. Whether a probabilistic program is useful for, for instance, discovering a hidden pattern in a given dataset or making an accurate prediction largely lies in these inference engines. These engines should compute accurate approximations or good model parameters within a fixed time budget. It is, thus, not surprising that substantial research efforts have been made to develop efficient inference algorithms and their implementations (as inference engines) [Chaganty et al. 2013; Holtzen et al. 2020; Kucukelbir et al. 2015; Mansinghka et al. 2014; Nori et al. 2014; Ritchie et al. 2016a; Schulman et al. 2015; van de Meent et al. 2018; Wingate et al. 2011a,b; Wingate and Weber 2013; Zhou et al. 2020].

We are concerned with smoothness¹ properties of probabilistic programs, which have been exploited by performant posterior-inference and model-learning algorithms and engines. For instance, when probabilistic programs are differentiable (in the sense that they define differentiable unnormalised densities), their posteriors can be inferred by Hamiltonian Monte Carlo [Neal 2011], one of the best performing MCMC algorithms. Also, in that case, their posteriors and model parameters can be inferred or learnt using the pathwise gradient estimator [Kingma and Welling 2014; Rezende et al. 2014], a popular technique for estimating the gradient of a function using samples. We also point out that the need for smoothness arises in a broader context of machine learning and computer science; Lipschitz continuity is one of the desired or at least recommended properties for neural networks [Arjovsky et al. 2017; Kim et al. 2021], and also differentiability commonly features as a requirement for pieces of code inside simulation software and cyber physical systems, where differential equations are used to specify the environment [Platzer 2018].

We present a static analysis that enables optimised posterior inference and model learning for probabilistic programs. We develop a static analysis that discovers differentiable or more generally smooth parts of given probabilistic programs, and show how the analysis can be used to improve the pathwise gradient estimator. Our improvement increases the scope of the estimator from differentiable to non-differentiable models, without requiring any intervention from the user; the improved estimator automatically identifies differentiable parts of probabilistic programs using our static analysis, and applies the pathwise gradient estimator to the identified parts while using a more general but less efficient estimator, called score estimator [Ranganath et al. 2014; Williams 1992], for the rest of the programs.

Our static analysis for smoothness has a surprisingly subtle soundness argument, partly due to the misbehaviours of some target smoothness properties when viewed from the perspective of program analysis designers. For instance, some smoothness properties, such as partial differentiability and partial continuity, are not preserved by function composition, and this makes it difficult to analyse sequential composition soundly without heavily sacrificing precision. In fact, overlooking such misbehaviours has been a source of errors in published static analyses for continuity [Chaudhuri et al. 2010, 2012].² We formulate five assumptions that clearly identify what a smoothness property should satisfy in order to avoid unsound analysis. Interestingly, these assumptions also determine what the property is allowed to violate. For instance, they reveal that the smoothness property does not

¹In mathematics, “smoothness” typically refers to the *specific* property of functions: being infinitely differentiable. In this paper, we override the term to denote a *set* of properties of functions describing well-behavedness (e.g., differentiability).

²The analysis in [Chaudhuri et al. 2010] infers the continuity property for multivariate programs, but it incorrectly joins two input-variable sets if a program is continuous with respect to each set *jointly*. Such a rule would hold if separate per-input-variable continuity were considered, but it does not hold for multivariate joint continuity. Conversely, the analysis in [Chaudhuri et al. 2012] considers a *per-input-variable* definition of continuity, but incorrectly assumes that this per-input-variable continuity is preserved by function composition. These (and other) issues make the two analyses unsound in several aspects (see §A for details). We do not claim that these unsoundness issues are hard to fix. Instead, our point is that a similar issue may be introduced easily and remain undetected due to the subtlety in the soundness of a static analysis.

have to be closed under the limits of chains of smooth (partial) functions, although the closure under such limits, called admissibility, has often been used to justify proof rules about or static analysis of loops. Dispensing with the admissibility requirement broadens the scope of our program analysis non-trivially; some useful smoothness properties from mathematics fail to meet the requirement.

Our variant of the pathwise gradient estimator works by program transformation and non-standard execution. It first transforms given probabilistic programs based on the results of our static analysis. Then, our estimator executes the transformed programs according to a standard (sampling) semantics, and collects sampled values during the execution. Finally, using the collected values, the estimator executes the original programs according to a non-standard (density) semantics this time. During the execution, our estimator computes a quantity involving differentiation, which becomes the estimate of the target gradient. We prove that our estimator satisfies an important differentiability requirement and thus, under a regularity condition, it is correct: the computed estimate is unbiased, i.e., it is the target gradient when averaged over random choices made during execution.

Our static analysis and variant of the pathwise gradient estimator have been implemented for a subset of the Pyro probabilistic programming language [Bingham et al. 2019]. They have been successfully applied to the 13 representative Pyro examples, which include advanced models with deep neural networks, such as attend-infer-repeat [Eslami et al. 2016] and single-cell annotation using variational inference [Xu et al. 2021]. For each of these examples, Pyro provides a (default) selective use of the pathwise gradient estimator but without any correctness guarantee. Our analysis and improved estimator automatically reproduced those uses, and proved that in those use cases, the estimator satisfies an important differentiability requirement and it is, thus, correct (i.e., unbiased) under a regularity condition (which needs to be discharged separately).

We summarise the main contributions of the paper:

- We present a program analysis for smoothness properties such as differentiability, and explain a subtle soundness argument for the analysis. Our argument identifies five assumptions for target smoothness properties, which are violated by some well-known smoothness properties and can help detect and prevent soundness errors in static analyses for smoothness properties (§5).
- We present a gradient estimator for probabilistic programs that improves the well-known pathwise gradient estimator using our program analysis. We also prove that our estimator satisfies an important differentiability requirement and it is, thus, correct (namely unbiased) under a regularity condition (§4 and §6).
- We show that our program analysis and gradient estimator can be successfully applied to representative probabilistic programs in Pyro, and can prove that existing unproved optimisations for these programs satisfy the differentiability requirement (§7).

The appendix (i.e., §A–§I) includes omitted proofs and details, and can be found in [Lee et al. 2022b].

2 INFORMAL DESCRIPTION OF BASIC CONCEPTS AND OUR APPROACH

We start by describing informally basic concepts and the goal of our approach, which we hope helps the reader to see the big picture of our technical contributions. To simplify presentation, we use toy examples in the section. But we emphasise that our approach has been applied to representative Pyro programs that describe advanced machine-learning models with deep neural networks.

Probabilistic programming and variational inference. In a probabilistic programming language (PPL), a program expresses a probabilistic model. As an example, consider the program c_m in Fig. 1, which describes a probabilistic model of the random variables z_1 and z_2 in \mathbb{R} by specifying their *unnormalised* density

$$p_{c_m}(z_1, z_2) = \mathcal{N}(z_1; 0, 5) \cdot \mathcal{N}(z_2; z_1, 3) \cdot (\mathbf{1}_{[z_2 > 0]} \cdot \mathcal{N}(0; 1, 1) + \mathbf{1}_{[z_2 \leq 0]} \cdot \mathcal{N}(0; -2, 1)),$$

$$c_m = \left(\begin{array}{l} x_1 := \text{sam}("z_1", \text{dist}_N(0, 5), \lambda y.y); \\ x_2 := \text{sam}("z_2", \text{dist}_N(x_1, 3), \lambda y.y); \\ \text{if } (x_2 > 0) \{ \text{obs}(\text{dist}_N(1, 1), 0) \} \\ \text{else} \quad \{ \text{obs}(\text{dist}_N(-2, 1), 0) \} \end{array} \right), \quad c_g = \left(\begin{array}{l} x_1 := \text{sam}("z_1", \text{dist}_N(\theta_1, 1), \lambda y.y); \\ x_2 := \text{sam}("z_2", \text{dist}_N(\theta_2, 1), \lambda y.y) \end{array} \right).$$

Fig. 1. A model c_m and a guide c_g in a PPL. Here $\text{dist}_N(a, b)$ is the distribution expression, and denotes the normal distribution with mean a and variance b .

$$c'_g = \left(\begin{array}{l} x_1 := \text{sam}("z_1", \text{dist}_N(0, 1), \lambda y.y + \theta_1); \\ x_2 := \text{sam}("z_2", \text{dist}_N(0, 1), \lambda y.y + \theta_2) \end{array} \right), \quad c''_g = \left(\begin{array}{l} x_1 := \text{sam}("z_1", \text{dist}_N(0, 1), \lambda y.y + \theta_1); \\ x_2 := \text{sam}("z_2", \text{dist}_N(\theta_2, 1), \lambda y.y) \end{array} \right).$$

Fig. 2. A fully (or selectively) reparameterised guide c'_g (or c''_g).

where $\mathcal{N}(x; a, b)$ is the probability density of a normal distribution with mean a and variance b , and $1_{[\varphi]}$ is the indicator function that returns 1 if φ holds and 0 otherwise. The first two \mathcal{N} factors in the equation come from the sample commands (sam) in c_m . They are called prior distributions, and describe prior knowledge on two random variables named z_1 and z_2 . The last factor comes from the if and observe commands (if and obs), which express that an unnamed random variable is sampled and observed to be 0 and its distribution is $\text{dist}_N(1, 1)$ or $\text{dist}_N(-2, 1)$ depending on whether z_2 is positive or not. This factor is called likelihood, and it states information about z_1 and z_2 that comes from an observed data point 0. Ignore the third arguments of the sample commands of c_m for now, which have no effect on p_{c_m} ; they will be explained later.

The purpose of writing c_m in a PPL, called *model*, is to infer its *normalised* probability density

$$\bar{p}_{c_m}(z_1, z_2) \triangleq p_{c_m}(z_1, z_2) / \int p_{c_m}(z_1, z_2) dz_1 dz_2,$$

also called normalised posterior density. Intuitively, this normalised density brings together two types of information about z_1 and z_2 , the first from their prior distributions (expressed in the first and second lines of c_m), and the second from the observed data point 0 that depends on z_1 and z_2 (the third and fourth lines of c_m). This inference task is called *posterior inference* problem. Among a wide range of approaches to the problem, we focus on the approach called *variational inference*, which forms the core of the recent combination of PPLs and deep learning.

In variational inference, we posit another program c_g , called *guide*, that is simpler than c_m and parameterised by θ . Then, we approximate the normalised density of c_m by c_g with an optimal choice of θ . For instance, consider the program c_g in Fig. 1. The program specifies the following already-normalised probability density

$$p_{c_g, \theta}(z_1, z_2) = \mathcal{N}(z_1; \theta_1, 1) \cdot \mathcal{N}(z_2; \theta_2, 1).$$

It can serve as a guide program for \bar{p}_{c_m} . To best approximate \bar{p}_{c_m} by $p_{c_g, \theta}$, variational inference aims at finding θ that minimises some notion of the discrepancy (called KL divergence) between $p_{c_g, \theta}$ and \bar{p}_{c_m} , or equivalently that maximises the objective function \mathcal{L} (called evidence lower bound):

$$\arg \max_{\theta} \mathcal{L}(\theta) \quad \text{for } \mathcal{L}(\theta) \triangleq \mathbb{E}_{p_{c_g, \theta}(z_1, z_2)} [f_{\theta}(z_1, z_2)] \quad \text{with } f_{\theta}(z_1, z_2) \triangleq \log(p_{c_m}(z_1, z_2) / p_{c_g, \theta}(z_1, z_2)).$$

A standard way to solve this optimisation problem is to apply the gradient-ascent algorithm: starting from an initial value $\theta^{(0)}$ of θ , compute $\theta^{(t)}$ iteratively by $\theta^{(t+1)} \triangleq \theta^{(t)} + \eta \cdot \nabla_{\theta} \mathcal{L}(\theta^{(t)})$, and return $\theta^{(T)}$ for a sufficiently large $T \in \mathbb{N}$. Here $\eta \in \mathbb{R}_{>0}$ denotes a learning rate.

A challenging part in the algorithm is to compute $\nabla_{\theta} \mathcal{L}(\theta)$. An exact computation of the gradient is mostly intractable due to the expectation inside \mathcal{L} , which hinders the gradient from having a closed-form formula. Hence, in practice, we rather *estimate* (not exactly compute) the gradient via a Monte Carlo method: draw a random sample (\hat{z}_1, \hat{z}_2) from some distribution q_{θ} , apply some function g_{θ} to the sample, and use the result as an estimate to the gradient, i.e.,

$$g_{\theta}(\hat{z}_1, \hat{z}_2) \approx \nabla_{\theta} \mathcal{L}(\theta) \quad \text{for a sample } (\hat{z}_1, \hat{z}_2) \text{ drawn from } q_{\theta}(\hat{z}_1, \hat{z}_2). \quad (1)$$

An important desired property of such a gradient estimator is *unbiasedness*, which states that the estimate is accurate in expectation: $\mathbb{E}_{q_\theta(z_1, z_2)}[g_\theta(z_1, z_2)] = \nabla_\theta \mathcal{L}(\theta)$. The property is necessary for the algorithm to converge to a local optimum, and is, thus, desired.

Gradient estimators for variational inference: SCE, PGE, and SPGE. A standard estimator for $\nabla_\theta \mathcal{L}(\theta)$ is the *score estimator* (SCE) [Ranganath et al. 2014; Williams 1992], which is unbiased under mild conditions. It estimates $\nabla_\theta \mathcal{L}(\theta)$ by using the recipe in Eq. (1) with $q_\theta(z_1, z_2) = p_{c_g, \theta}(z_1, z_2)$ and

$$g_\theta(z_1, z_2) = f_\theta(z_1, z_2) \cdot \nabla_\theta \log q_\theta(z_1, z_2).$$

That is, the estimator draws a sample from the guide distribution $p_{c_g, \theta}$ and applies the above g_θ to obtain a gradient estimate.³ It is applicable to a wide range of model-guide pairs while remaining unbiased, but it is known to have a large approximation error (i.e., have a large variance).

The *pathwise gradient estimator* (PGE) [Kingma and Welling 2014; Rezende et al. 2014] is another standard gradient estimator, which is known to have a smaller approximation error than the SCE and thus has been a preferred option against the SCE. The PGE requires an additional program c'_g that is a θ -independent reparameterisation of the guide c_g . A program c' is said to be θ -independent if the probability densities of the sampled random variables in c' are θ -independent. It is called a *reparameterisation* of c if c and c' sample the same set of random variables and they have the same semantics on those variables in the following sense: when there are n random variables, for any measurable $h : \mathbb{R}^n \rightarrow \mathbb{R}$, we have $\mathbb{E}_{p_c(z)}[h(v_c(z))] = \mathbb{E}_{p_{c'}(z)}[h(v_{c'}(z))]$, where $p_c : \mathbb{R}^n \rightarrow \mathbb{R}$ is the probability density of all n random variables in c , and $v_c : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is the so called value function of c , which applies the lambda functions in the third arguments of c 's sample commands to the corresponding random variables. For example, c'_g in Fig. 2 is a θ -independent reparameterisation of c_g for $\theta = (\theta_1, \theta_2)$. It has the probability density $p_{c'_g}(z_1, z_2) = \mathcal{N}(z_1; 0, 1) \cdot \mathcal{N}(z_2; 0, 1)$, and the value function $v_{c'_g, \theta}(z_1, z_2) = (z_1 + \theta_1, z_2 + \theta_2)$. Note that $p_{c'_g}$ does not depend on θ , as required by the θ -independence of c_g . We can show this c'_g is a reparameterisation of c_g in Fig. 1 by using the fact that v_{c_g} is the identity function and $y = x + a$ for x drawn from $\mathcal{N}(x; 0, 1)$ follows the distribution $\mathcal{N}(y; a, 1)$.

Given a reparameterised guide c'_g , the PGE estimates $\nabla_\theta \mathcal{L}(\theta)$ by again following the recipe in Eq. (1) this time with $q'(z_1, z_2) = p_{c'_g}(z_1, z_2)$ and

$$g'_\theta(z_1, z_2) = \nabla_\theta f_\theta(z'_1, z'_2) \quad \text{for } (z'_1, z'_2) = v_{c'_g, \theta}(z_1, z_2).$$

This estimator differs from the SCE in two aspects. First, a random sample is drawn from a reparameterised-guide distribution $p_{c'_g}$, not from $p_{c_g, \theta}$. Next, the estimation function g'_θ computes the derivative of $f_\theta(z'_1, z'_2)$ with respect to θ and (z'_1, z'_2) (not with respect to θ only), since the argument of $f_\theta(-)$ in g'_θ depends on θ via $v_{c'_g, \theta}$.⁴

While having a small approximation error, the PGE requires more than the SCE to ensure the unbiasedness. An important additional requirement for the PGE is that (i) $p_{c_m}(z_1, z_2)$ and $p_{c_g, \theta}(z_1, z_2)$ should be differentiable in θ and z_1, z_2 and (ii) $v_{c'_g, \theta}(z_1, z_2)$ be differentiable in θ for all z_1, z_2 . The requirement is imposed partly to ensure that no differentiation error arises in computing g'_θ . This differentiability requirement, however, can be easily violated if a model or a guide starts to use branches or loops. For instance, it is violated by our example in Figs. 1 and 2 as $p_{c_m}(z_1, z_2)$ is not differentiable in z_2 . This violation makes the PGE biased for the example, i.e.,

$$\mathbb{E}_{q'_\theta(z_1, z_2)}[g'_\theta(z_1, z_2)] = (\dots, \frac{1}{3}(\theta_1 - \theta_2)) \neq (\dots, \frac{1}{3}(\theta_1 - \theta_2) + \frac{1}{2}\mathcal{N}(-\theta_2; 0, 1)) = \nabla_\theta \mathcal{L}(\theta),$$

³The log term in g_θ comes from the well-known log-derivative trick: $\nabla_\theta q_\theta(z_1, z_2) = q_\theta(z_1, z_2) \cdot \nabla_\theta \log q_\theta(z_1, z_2)$.

⁴By the chain rule, we have the following for each $i \in \{1, 2\}$:

$$\frac{\partial f_\theta(z'_1, z'_2)}{\partial \theta_i} = \frac{\partial f_\theta(x_1, x_2)}{\partial \theta_i} \Big|_{(x_1, x_2, \theta) = (z'_1, z'_2, \theta)} + \left(\left(\frac{\partial f_\theta(x_1, x_2)}{\partial x_1} \Big|_{(x_1, x_2, \theta) = (z'_1, z'_2, \theta)}, \frac{\partial f_\theta(x_1, x_2)}{\partial x_2} \Big|_{(x_1, x_2, \theta) = (z'_1, z'_2, \theta)} \right), \left(\frac{\partial v_{c'_g, \theta}(y_1, y_2)}{\partial \theta_i} \Big|_{(y_1, y_2, \theta) = (z_1, z_2, \theta)} \right) \right).$$

and thus causes the gradient-ascent algorithm to converge to a suboptimal θ : applying the PGE to the example produces a suboptimal solution $\theta = (0, 0)$, whereas the optimal solution is $\theta \approx (0.95, 1.52)$.

The *selective pathwise gradient estimator* (SPGE) [Schulman et al. 2015] combines the two previous gradient estimators to alleviate their limitations: one has a large approximation error, and the other imposes a strong requirement for unbiasedness. The SPGE requires an additional program c_g'' that is a reparameterisation of the guide c_g but needs not be θ -independent (unlike the PGE). An instance of c_g'' for our example is given in Fig. 2, which changes the sample command for z_1 in c_g but keeps the one for z_2 . Note that the changed sample command for z_1 in c_g'' uses a θ -independent probability distribution. Typically, c_g'' is obtained by *selecting* a subset of the random variables in c_g and changing the sample commands for the selected variables such that their probability distributions become θ -independent; the sample commands for the unselected remain as they are. Given c_g'' , the SPGE estimates $\nabla_\theta \mathcal{L}(\theta)$ by following the recipe in Eq. (1) with $q_\theta''(z_1, z_2) = p_{c_g'', \theta}(z_1, z_2)$ and

$$g_\theta''(z_1, z_2) = \nabla_\theta f_\theta(z_1'', z_2'') + f_\theta(z_1'', z_2'') \cdot \nabla_\theta \log q_\theta''(z_1, z_2) \quad \text{for } (z_1'', z_2'') = v_{c_g'', \theta}(z_1, z_2). \quad (2)$$

Note that the estimation function g_θ'' consists of two terms, which come from that of the PGE and the SCE. The second term adjusts the PGE to correctly account for unchanged random variables (e.g., z_2 in the example of Fig. 2).

By allowing a guide that makes only some selected (not all) random variables θ -independent, the SPGE offers two advantages at the same time: it achieves a smaller approximation error than the SCE, and imposes a weaker requirement for unbiasedness than the PGE. In particular, the differentiability requirement for the SPGE is weaker than that for the PGE, which is as follows for our example in Figs. 1 and 2: (i) $p_{c_m}(z_1, z_2)$ and $p_{c_g, \theta}(z_1, z_2)$ be differentiable in θ and z_1 (but they may be non-differentiable in z_2); and (ii) $v_{c_g'', \theta}(z_1, z_2)$ and $p_{c_g'', \theta}(z_1, z_2)$ be differentiable in θ for all z_1, z_2 . This requirement holds, and as a result, the SPGE with this c_g'' is unbiased (whereas the PGE with the given c_g' is biased as seen before). You can find in §B a table summarising and comparing SCE, PGE, and SPGE.

Variable-selection problem for SPGE. To maximize the advantages offered by the SPGE, we consider the following algorithmic problem:

Definition 2.1 (SPGE Variable-Selection Problem; Informal). Assume that we are given a model c_m , a guide c_g , and a *reparameterisation plan* π , i.e., a map from sample commands to sample commands. Then, find automatically a large subset S of random variables such that if we let $\bar{c}_g^{\pi, S}$ be the result of π -transforming every sample command in c_g that defines a random variable in S , then $\bar{c}_g^{\pi, S}$ is a reparameterisation of c_g and $(c_m, c_g, \bar{c}_g^{\pi, S})$ satisfies the differentiability requirement for the SPGE. \square

An instantiation of the problem for our example is that c_m and c_g are programs in Fig. 1 and π transforms commands of the form $y := \text{sam}(n, \text{dist}_N(e', 1), \lambda y.y)$ to $y := \text{sam}(n, \text{dist}_N(0, 1), \lambda y.y + e')$, while leaving all the other sample commands as they are. In this instantiation, the condition in the problem is met by $S = \emptyset$ and $S = \{z_1\}$, and the latter option is preferred due to its size. Note that the solution $S = \{z_1\}$ yields the guide c_g'' in Fig. 2, that is, $\bar{c}_g^{\pi, S} = c_g''$.

Existing PPLs, when applying the SPGE, choose an S without checking the differentiability requirement, and this can make the requirement easily violated. For instance, given a model-guide pair, in one of its standard settings, Pyro automatically applies the SPGE with S being the set of all continuous random variables in the guide. This choice of S , however, does not guarantee the requirement is met. For our example in Fig. 1, Pyro chooses $S = \{z_1, z_2\}$, but this S violates the requirement; due to this, the SPGE becomes biased and Pyro returns a suboptimal $\theta = (0, 0)$.

In the rest of the paper, we will present our solution to the SPGE variable-selection problem. A core component of our solution is a general static analysis framework for smoothness properties such as differentiability (§5), which our solution uses to discharge the differentiability requirement

for the SPGE correctly and automatically. As we briefly mentioned in the introduction, automatically analysing the smoothness properties of a program in a sound manner is surprisingly subtle. Our analysis framework identifies five assumptions for smoothness properties, and prove that the analysis is sound if a target smoothness property satisfies these assumptions.

Our solution for the SPGE variable-selection problem (§6) runs the static analysis on given c_m and c_g , and computes a maximal set S' of random variables in which p_{c_m} and $p_{c_g, \theta}$ are differentiable. Then, it heuristically searches for a subset of S' starting from S' itself such that $\overline{c_g}^{\pi, S'}$ satisfies the differentiability requirement. For instance, for our example in Fig. 1, our differentiability analysis infers that p_{c_m} and $p_{c_g, \theta}$ are differentiable in $\{z_1\}$ and $\{z_1, z_2\}$, respectively. From this, we set $S' = \{z_1\}$, run our analysis again on $\overline{c_g}^{\pi, S'}$, and get confirmation that $\overline{c_g}^{\pi, S'}$ meets the requirement, i.e., $p_{\overline{c_g}^{\pi, S'}, \theta}$ and $v_{\overline{c_g}^{\pi, S'}, \theta}$ are differentiable in θ . Thus, this S' becomes the final result. In fact, this first-round success appeared in our experiments (§7): our implementation shows on all tested examples that the initial choice of S' is indeed valid in the above sense so that no subset search is necessary.

We point out that to mathematically develop and analyse our solution for the SPGE variable-selection problem, we formalise the SPGE in the PPL setting and formally derive a sufficient condition for its unbiasedness, which includes the differentiability requirement (§4).

3 SETUP

We use a simple imperative probabilistic programming language, which models the core of popular imperative PPLs, such as Pyro. Programs in the language describe densities, which are sometimes unnormalised (i.e., they do not integrate to 1). In this section, we describe the syntax and semantics of the language, and also variational inference for the language.

Syntax of a simple imperative PPL. Let PVar be a finite set of program variables, Str be a finite set of strings, and Fn be a set of function symbols that represent measurable maps of type $\mathbb{R}^k \rightarrow \mathbb{R}$. The language has the following syntax:

Real Expr. $e ::= x \mid r \mid op(e_1, \dots, e_k)$ Boolean Expr. $b ::= \text{true} \mid e_1 < e_2 \mid b_1 \wedge b_2 \mid \neg b$
 Name Expr. $n ::= \text{name}(\alpha, e)$ Distribution Expr. $d ::= \text{dist}_N(e, e')$
 Command $c ::= \text{skip} \mid x := e \mid c; c' \mid \text{if } b \{c\} \text{ else } \{c'\} \mid \text{while } b \{c\} \mid x := \text{sam}(n, d, \lambda y.e) \mid \text{obs}(d, r)$

Here x, r, op , and α stand for a program variable in PVar , a real number, a function symbol in Fn , and a string in Str , respectively.

The language has four kinds of expressions, which denote maps from states to values of appropriate types. All the real and boolean expressions are standard. The name expressions n denote the identifiers of drawn samples (i.e., random variables). They are built by appending an integer (obtained by the floor of a real) to a string in Str ; e.g., $\text{name}("z", 3.2)$ denotes the name $("z", 3)$.⁵ The distribution expression $\text{dist}_N(e, e')$ denotes the normal distribution with mean e and variance e' . The language supports standard commands for imperative computation, and additionally has sample and observe for probabilistic programming. The sample command $x := \text{sam}(n, d, \lambda y.e)$ creates a random variable named n by drawing a sample r from d ; then, it transforms r to $e[r/y]$ and stores the result in the program variable x . In the programs written by the user of the language, only the identity function $\lambda y.y$ appears as the third argument of the sample commands. But as we explain later, when a program is constructed from another by a gradient estimator, such as the SPGE, it may contain sample commands with non-identity function arguments. The observe command $\text{obs}(d, r)$ describes that an unnamed random variable is drawn from d and is immediately observed to have the value r . Computationally, $\text{obs}(d, r)$ calculates the probability density of d at r and updates a variable that tracks the product of these probabilities from all the observations, by multiplying the variable with the calculated density.

⁵The name construct has the second argument to easily support the sampling of (conditionally) i.i.d. random variables.

Density semantics of the PPL.⁶ We use a semantics of our language where commands are interpreted as calculators for densities, which may be unnormalised. Commands transform states, but in so doing, they compute densities of sampled random variables. More precisely, in the semantics, a command starts with an initial state that fixes not just the values of program variables but also those of all the random variables that are to be sampled during execution. When the command runs, it calculates the densities of those random variables at their given initial values, and also computes the probability density of all the observations, called *likelihood*. The product of all these densities and the likelihood becomes the so called *unnormalised posterior density*.

Let \mathbb{N} be the set of natural numbers. Fix $N \in \mathbb{N}$ with $N \geq 1$. Formally, the semantics uses the states of the following form:

$$\begin{aligned} \mu \in \text{Name} &\triangleq \{(\alpha, i) \mid \alpha \in \text{Str}, i \in \mathbb{N} \cap [0, N)\}, \\ a \in \text{AVar} &\triangleq \{\text{like}\} \cup \{pr_\mu, val_\mu, cnt_\mu \mid \mu \in \text{Name}\}, \\ u, v \in \text{Var} &\triangleq \text{Name} \uplus \text{PVar} \uplus \text{AVar}, \\ \sigma \in \text{St} &\triangleq [\text{Var} \rightarrow \mathbb{R}], \quad \text{St}[K] \triangleq [K \rightarrow \mathbb{R}] \text{ for } K \subseteq \text{Var}. \end{aligned}$$

Here $\sigma(\mu)$ for $\mu \in \text{Name}$ is the initial value of the random variable μ , which is used by the sample command and does not change during execution. For technical simplicity, the set Name has the restriction that the integer part of a name must be in $[0, N)$.⁷ The set AVar consists of four types of auxiliary variables. The auxiliary variable *like* stores the likelihood (i.e., the probability density of all the observations), and its value is initialised to 1 and changes whenever the observe command $\text{obs}(d, r)$ runs; the new value becomes the old times the density of the probability distribution d at r . The other auxiliary variables pr_μ , val_μ , and cnt_μ are associated with a random variable μ , standing for the “prior”, “value”, and “counter” of μ . They are initialised with $\mathcal{N}(\sigma(\mu); 0, 1)$ (i.e., the density of the standard normal distribution at $\sigma(\mu)$), $\sigma(\mu)$, and 0, respectively, and get updated by the sample command $x := \text{sam}(n, d, \lambda y. e)$ where n denotes μ . The command increases cnt_μ by 1, so as to record the occurrence of a sampling event for μ . Then, it looks up the given value $\sigma(\mu)$ of the random variable μ , transforms the value to $e[\sigma(\mu)/y]$, and stores the result in x and val_μ . Finally, the command computes the density of the distribution d at the looked-up value $\sigma(\mu)$, and updates pr_μ with this density. The unnormalised posterior density (i.e., the joint density of all the random variables and observations) is then obtained by multiplying at the end of program execution the values of *like* and pr_μ for all $\mu \in \text{Name}$.

The formal semantics of expressions is standard, and has the following types:

$$\llbracket e \rrbracket : \text{St} \rightarrow \mathbb{R}, \quad \llbracket b \rrbracket : \text{St} \rightarrow \mathbb{B}, \quad \llbracket n \rrbracket : \text{St} \rightarrow \text{Name}, \quad \llbracket d \rrbracket : \text{St} \rightarrow \mathbb{D}.$$

Here \mathbb{B} is the set of booleans, i.e., *true* and *false*, and \mathbb{D} is the set of positive probability-density functions on \mathbb{R} , i.e., a subset of $[\mathbb{R} \rightarrow (0, \infty)]$ whose elements are measurable functions that integrate to 1. The semantics is defined for a minor extension of the set of expressions where non-program variables are allowed to appear, such as $(\mu + x)$. The interpretation of expressions is mostly standard. We

⁶Our semantics is an instance (or variant) of existing density semantics (e.g., [Lee et al. 2020]), and is different from sampling semantics (e.g., [Staton et al. 2016]). Although the density semantics and the sampling semantics have different presentations, they are closely related and equivalent in a formal sense (see, e.g., [Lee et al. 2020]). We use the density semantics instead of the sampling semantics, because the gradient estimator (§4) of our interest performs computation on (unnormalised) densities and it is easier for a program analysis (§5) to work with the density semantics than the sampling semantics.

⁷This restriction is often respected by probabilistic programs in practice, since they commonly sample random variables whose number is uniformly bounded over all traces; note, however, that it is not always respected (e.g., by programs from Bayesian nonparametrics). The uniform bound N can often be found by a simple static analysis. This restriction along with the finiteness of PVar and Str implies the finiteness of Var , and this makes our technical development simpler since $\sigma \in \text{St}$ becomes a function on a finite-dimensional space. Lifting the restriction would make the technical development more complicated, since this would require St to be isomorphic to $[\mathbb{R}^\infty \rightarrow \mathbb{R}]$ or $\biguplus_k [\mathbb{R}^k \rightarrow \mathbb{R}]$ and the former (or latter) choice of St makes defining differentiability (or formalising our program analysis) nontrivial; we leave it as future work.

show only the case for the name expressions $n \equiv \text{name}(\alpha, e) : \llbracket \text{name}(\alpha, e) \rrbracket \sigma \triangleq \text{create_name}(\alpha, \llbracket e \rrbracket \sigma)$, where $\text{create_name} : \text{Str} \times \mathbb{R} \rightarrow \text{Name}$ is an operator converting a string-real pair to a name, defined by $\text{create_name}(\alpha, r) = (\alpha, \min\{\max\{\lfloor r \rfloor, 0\}, N - 1\})$.⁸

Note that the types of the semantics of expressions imply that the evaluation of an expression always produces a value of the right type. In particular, it never generates an error. For instance, when an argument of an operator op or a distribution constructor dist_N is outside its intended domain as in $\log(-3)$ and $\text{dist}_N(0, -2)$, or when the integer part of a name expression is outside $[0, N)$ as in $\text{name}("z", -1)$, our semantics does not generate an error. Instead, it returns some pre-chosen default value of the right type. This slightly unusual way of handling errors is also adopted in our semantics of commands to be presented shortly, and it lets us avoid the complexity caused by error handling when we formalise variational inference and develop our program analysis for smoothness properties.

The formal semantics of commands is also mostly standard with the handling of errors via default values, although its interpretation of sample and observe commands deserves special attention. Let \perp be an element not in St , and define St_\perp to be the usual lifting of St with \perp . That is, St_\perp is a partially-ordered set $\text{St} \sqcup \{\perp\}$ with the following order: for $\xi, \xi' \in \text{St}_\perp$, we have $\xi \sqsubseteq \xi'$ if and only if $\xi = \perp$ or $\xi = \xi'$. We write the standard lifting of a function $f : \text{St} \rightarrow \text{St}_\perp$ by $f^\dagger : \text{St}_\perp \rightarrow \text{St}_\perp$ (i.e., $f^\dagger(\xi) \triangleq \text{if } (\xi = \perp) \text{ then } \xi \text{ else } f(\xi)$). The semantics of a command c is a map $\llbracket c \rrbracket : \text{St} \rightarrow \text{St}_\perp$, and is defined inductively as shown below:

$$\begin{aligned} \llbracket \text{skip} \rrbracket \sigma &\triangleq \sigma, \\ \llbracket x := e \rrbracket \sigma &\triangleq \sigma[x \mapsto \llbracket e \rrbracket \sigma], \\ \llbracket c; c' \rrbracket \sigma &\triangleq \llbracket c' \rrbracket^\dagger(\llbracket c \rrbracket \sigma), \\ \llbracket \text{if } b \{c\} \text{ else } \{c'\} \rrbracket \sigma &\triangleq \text{if } (\llbracket b \rrbracket \sigma = \text{true}) \text{ then } \llbracket c \rrbracket \sigma \text{ else } \llbracket c' \rrbracket \sigma, \\ \llbracket \text{while } b \{c\} \rrbracket \sigma &\triangleq (\text{fix } F)(\sigma) \quad \text{where } F(f)(\sigma) \triangleq \text{if } (\llbracket b \rrbracket \sigma = \text{true}) \text{ then } f^\dagger(\llbracket c \rrbracket \sigma) \text{ else } \sigma, \\ \llbracket x := \text{sam}(n, d, \lambda y. e') \rrbracket \sigma &\triangleq \sigma[x \mapsto r, \text{val}_\mu \mapsto r, \text{pr}_\mu \mapsto \llbracket d \rrbracket \sigma(\sigma(\mu)), \text{cnt}_\mu \mapsto \sigma(\text{cnt}_\mu) + 1] \\ &\quad \text{where } \mu \triangleq \llbracket n \rrbracket \sigma \text{ and } r \triangleq \llbracket e'[\mu/y] \rrbracket \sigma, \\ \llbracket \text{obs}(d, r) \rrbracket \sigma &\triangleq \sigma[\text{like} \mapsto \sigma(\text{like}) \cdot \llbracket d \rrbracket \sigma(r)]. \end{aligned}$$

The interpretation uses the least fixed-point operator fix for continuous maps F on the function space $[\text{St} \rightarrow \text{St}_\perp]$, where the function space is ordered pointwise and continuity means the one with respect to this order. It also uses the notation $e'[e''/y]$ to denote the substitution of y in e' by e'' . According to this interpretation, $x := \text{sam}(n, d, \lambda y. e')$ increments the cnt_μ variable for the name $n = \mu$ so that the variable, which has 0 initially, records the number of times that the random variable with the same name n is sampled during execution.

Having some cnt_μ variable increased by 2 or larger at some point of execution is not an intended behaviour of a command c . That is, if c is a well-designed command, every random variable with a fixed name should be sampled at most once during the execution of c . This intended behaviour of commands plays an important role in our results, and we refer to it using the following terminology.

Definition 3.1. An always-terminating command c does not have a double-sampling error if for any $\sigma \in \text{St}$, we have $\llbracket c \rrbracket \sigma(\text{cnt}_\mu) - \sigma(\text{cnt}_\mu) \leq 1$ for all $\mu \in \text{Name}$. \square

Example 3.2 (Density semantics). Consider the following state $\sigma \in \text{St} : \sigma \triangleq [x \mapsto 0, y \mapsto 0, ("a", 0) \mapsto 2, ("b", 0) \mapsto 4, \text{cnt}_{("a", 0)} \mapsto 0, \text{cnt}_{("b", 0)} \mapsto 0, \dots]$, where $x, y \in \text{PVar}$ denote program variables and $("a", 0), ("b", 0) \in \text{Name}$ denote random variables. Note that σ consists of three parts: the PVar part says that the values of x and y are both 0; the Name part says that the values of $("a", 0)$ and $("b", 0)$ are 2 and 4; and the AVar part says that $("a", 0)$ and $("b", 0)$ have not been sampled yet.

⁸There are multiple valid ways to convert a string-real pair to a name (i.e., to define create_name); we choose just one of them.

Next, consider a command $c \equiv (x := \text{sam}(\text{name}("a", 0), \text{dist}_N(-3, 1), \lambda z.z); y := \text{sam}(\text{name}("b", 0), \text{dist}_N(5, 1), \lambda z.z))$. Given the command c and the input state σ , our density semantics computes the following output state $\sigma' \in \text{St}$: $\sigma' \triangleq \llbracket c \rrbracket \sigma = [x \mapsto 2, y \mapsto 4, ("a", 0) \mapsto 2, ("b", 0) \mapsto 4, \text{cnt}("a", 0) \mapsto 1, \text{cnt}("b", 0) \mapsto 1, \text{pr}("a", 0) \mapsto \mathcal{N}(2; -3, 1), \text{pr}("b", 0) \mapsto \mathcal{N}(4; 5, 1), \dots]$. The input/output states σ and σ' illustrate two aspects of our semantics. First, the semantics uses the Name part of an input state to determine the sampled values of sample commands: x and y in σ' take the values of $\sigma(("a", 0)) = 2$ and $\sigma(("b", 0)) = 4$. Second, the semantics records, in the AVar part of an output state, the densities of sample commands: $\text{pr}("a", 0)$ and $\text{pr}("b", 0)$ in σ' store the densities of the two sample commands in c evaluated at 2 and 4. \square

Variational inference. We consider the most common form of *variational inference* for Pyro-like PPLs where we are asked to learn a good approximation of the posterior of a given model, i.e., the conditional distribution of the model given a dataset. Typically, a parameterised approximate posterior is given in variational inference, and learning corresponds to finding good values of those parameters. A popular approach is to measure the quality of parameter values by the so called evidence lower bound (ELBO), and to optimise ELBO.

To translate what we have described so far to our context, we need to explain a general recipe for generating a density p_c for a command c , which is in general unnormalised (i.e., does not integrate to 1). The recipe specifies p_c as follows: for each $\sigma_\theta \in \text{St}[\theta]$ with $\theta \subseteq \text{PVar}$, $p_{c, \sigma_\theta} : \text{St}[\text{Name}] \rightarrow [0, \infty)$ is defined by

$$p_{c, \sigma_\theta}(\sigma_n) \triangleq \begin{cases} \llbracket c \rrbracket \sigma(\text{like}) \cdot \prod_{\mu \in \text{Name}} \llbracket c \rrbracket \sigma(\text{pr}_\mu) & \text{if } \llbracket c \rrbracket \sigma \in \text{St} \text{ and } \llbracket c \rrbracket \sigma(\text{cnt}_\mu) \leq 1 \text{ for all } \mu \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

where $\sigma = \sigma_0 \oplus \sigma_\theta \oplus \sigma_n \in \text{St}$, and the \oplus operator combines two real-valued maps with disjoint domains in the standard way. Also, $\sigma_0 \in \text{St}[(\text{PVar} \setminus \theta) \uplus \text{AVar}]$ maps *like* to 1, and pr_μ to $\mathcal{N}(\sigma_n(\mu); 0, 1)$ and val_μ to $\sigma_n(\mu)$ for every $\mu \in \text{Name}$, and all other variables to 0. Here $\text{St}[\text{Name}]$ is understood as a measurable space constructed by taking the product of the $|\text{Name}|$ copies of the measurable space \mathbb{R} and the integral is taken over the uniform measure on $\text{St}[\text{Name}]$ (i.e., the product of the $|\text{Name}|$ copies of the Lebesgue measure on \mathbb{R}).

In variational inference in our PPL context, we are given two commands c_m and c_g , called *model* and *guide*. We assume that (i) these commands always terminate and do not have a double-sampling error, (ii) some variables $\theta = \{\theta_1, \dots, \theta_k\} \subseteq \text{PVar}$ that only appear in c_g , not in c_m , are identified as parameters to be optimised, and (iii) the density p_{c_g, σ_θ} of the guide c_g integrates to 1 and defines a probability distribution.⁹ Given the model-guide pair (c_m, c_g) , a popular approach for variational inference is to solve the following optimisation problem approximately,

$$\underset{\sigma_\theta}{\text{argmax}} \mathbb{E}_{p_{c_g, \sigma_\theta}(\sigma_n)} [\log(p_{c_m}(\sigma_n)/p_{c_g, \sigma_\theta}(\sigma_n))], \quad (4)$$

when the expectation is well-defined for all σ_θ . The objective of this optimisation is the ELBO that we mentioned earlier. Here p_{c_m} means p_{c_m, σ'_θ} for some/any σ'_θ ; the choice of σ'_θ does not matter since c_m does not access the parameters θ and so $p_{c_m, \sigma'_\theta} = p_{c_m, \sigma''_\theta}$ for all $\sigma'_\theta, \sigma''_\theta \in \text{St}[\theta]$.

Often variational inference is applied when the model c_m is parameterised as well. In those cases, it asks for finding good parameters of the model c_m as well as those of the guide c_g . So, an algorithm for variational inference this time simultaneously learns a good model for given observations and a good approximate posterior for the learnt model. This more general form of variational inference

⁹In practice, one more assumption is required: the set of random variables sampled from the model should be the same as that from the guide. This assumption can be checked automatically, e.g., by [Lee et al. 2020; Lew et al. 2020]. In this work, however, this assumption is always satisfied: all random variables in Name are sampled by a sample command or at the beginning (via initialisation).

can be easily accommodated in our setup. We just need to drop the condition that the parameters may not appear in c_m , and to use p_{c_m, σ_θ} instead of p_{c_m} in the optimisation objective in Eq. (4):

$$\operatorname{argmax}_{\sigma_\theta} \mathcal{L}(\sigma_\theta) \text{ for } \mathcal{L}(\sigma_\theta) \triangleq \mathbb{E}_{p_{c_g, \sigma_\theta}(\sigma_n)} [\log(p_{c_m, \sigma_\theta}(\sigma_n)/p_{c_g, \sigma_\theta}(\sigma_n))]. \quad (5)$$

The rest of the paper focuses on this general form of variational inference (often called model learning).

4 SELECTIVE PATHWISE GRADIENT ESTIMATOR

We consider a gradient-based algorithm for the optimisation problem in Eq. (5). The algorithm finds a good σ_θ by repeatedly estimating the gradient of the optimisation objective at the current σ_θ ,

$$\operatorname{grad_est}(\sigma_\theta) \approx \nabla_{\theta} \mathbb{E}_{p_{c_g, \sigma_\theta}(\sigma_n)} [\log(p_{c_m, \sigma_\theta}(\sigma_n)/p_{c_g, \sigma_\theta}(\sigma_n))],$$

and updating σ_θ with the estimate under a learning rate $\eta > 0$, that is, $\sigma_\theta \leftarrow \sigma_\theta + \eta \cdot \operatorname{grad_est}(\sigma_\theta)$. Note that the core of the algorithm lies in the computation of $\operatorname{grad_est}(\sigma_\theta)$.

In this section, we describe a particular algorithm for the gradient computation, called *selective pathwise gradient estimator (SPGE)*, which is often regarded as the algorithm of choice and corresponds to the inference algorithm developed for stochastic computation graphs [Schulman et al. 2015] and implemented for Pyro. Our description of the SPGE takes the often-ignored aspect of customising the SPGE algorithm for PPLs seriously, and it is accompanied with a novel formal analysis of the customisation (§4.1 and §4.2). Our analysis clearly identifies information about probabilistic programs that is useful for this customised SPGE algorithm, and prepares the stage for our program analysis for smoothness properties in §5 (§4.2 and §4.3).

4.1 Program Transformation

We start by describing a program transformation that changes some sample commands in a given probabilistic program. This transformation is used crucially by the SPGE.

The key component of the transformation is a partial function π called *reparameterisation plan*, which has the type $\text{NameEx} \times \text{DistEx} \times \text{LamEx} \rightarrow \text{DistEx} \times \text{LamEx}$. Here NameEx , DistEx , and LamEx denote the sets of name expressions, distribution expressions, and lambda expressions of the form $\lambda y.e$, respectively. The plan π specifies how we transform sample commands. Concretely, assume that we are given $x := \text{sam}(n, d, \lambda y.e)$. We check whether $\pi(n, d, \lambda y.e)$ is defined or not. If not, we keep the original sample command. Otherwise, if $\pi(n, d, \lambda y.e)$ is $(d', \lambda y'.e')$, we replace the command with $x := \text{sam}(n, d', \lambda y'.e')$.

A natural extension of this intended transformation of π leads to the following program transformation for a general command c , denoted by \bar{c}^π :

$$\begin{aligned} \overline{\text{skip}}^\pi &\triangleq \text{skip}, \\ \overline{x := e}^\pi &\triangleq x := e, \\ \overline{c; c'}^\pi &\triangleq \bar{c}^\pi; \bar{c}'^\pi, \\ \overline{\text{if } b \{c\} \text{ else } \{c'\}}^\pi &\triangleq \text{if } b \{\bar{c}^\pi\} \text{ else } \{\bar{c}'^\pi\}, \\ \overline{\text{while } b \{c\}}^\pi &\triangleq \text{while } b \{\bar{c}^\pi\}, \\ \overline{x := \text{sam}(n, d, l)}^\pi &\triangleq \begin{cases} x := \text{sam}(n, d', l') & \text{if } \exists (d', l'). \pi(n, d, l) = (d', l') \\ x := \text{sam}(n, d, l) & \text{otherwise,} \end{cases} \\ \overline{\text{obs}(d, r)}^\pi &\triangleq \text{obs}(d, r). \end{aligned}$$

The transformation recursively traverses c , and applies π to all the sample commands in c . Note that for any π , there exists a total function π' such that $\bar{c}^\pi = \bar{c}^{\pi'}$ for all c ; the π' coincides with π in the domain of π , and outside of this domain, it is the identity function. But such π' loses information about the domain of π , which plays a crucial role in our formalisation of the SPGE.

We are primarily interested in semantics-preserving instances of $\bar{\tau}^\pi$. The next definition helps us to identify such instances.

Definition 4.1. A reparameterisation plan π is *valid* if for all $n \in \text{NameEx}$, $d, d' \in \text{DistEx}$, and $(\lambda y.e), (\lambda y'.e') \in \text{LamEx}$ such that $\pi(n, d, \lambda y.e) = (d', \lambda y'.e')$, the following condition holds: for all states $\sigma \in \text{St}$ and measurable subsets $A \subseteq \mathbb{R}$,

$$\int \mathbf{1}_{[\![e[r/y]]\!] \sigma \in A} \cdot \llbracket d \rrbracket \sigma(r) dr = \int \mathbf{1}_{[\![e'[r/y']]\!] \sigma \in A} \cdot \llbracket d' \rrbracket \sigma(r) dr. \quad \square (6)$$

The condition says that the distribution obtained by sampling from d and applying $\lambda y.e$ is the same as that obtained by sampling from d' and applying $\lambda y'.e'$. An example of a widely-used valid reparameterisation plan maps its input as follows, whenever defined: $\pi_0(n, \text{dist}_\mathbb{N}(e_1, e_2), \lambda y.e_3) = (\text{dist}_\mathbb{N}(0, 1), \lambda y.e_3[(y \times \sqrt{e_2 + e_1})/y])$, where we assume y does not appear in e_1 and e_2 , the substitution in π_0 expresses the composition of two functions $\lambda y.e_3$ and $\lambda y.(y \times \sqrt{e_2 + e_1})$, and $\sqrt{-}$ denotes a square-root operator that handles non-positive arguments in the same way as $\text{dist}_\mathbb{N}(e, -)$ does: if $\llbracket e_2 \rrbracket \sigma \leq 0$ and $\llbracket \text{dist}_\mathbb{N}(e_1, e_2) \rrbracket \sigma = \lambda r. \mathcal{N}(r; \llbracket e_1 \rrbracket \sigma, r_2)$ for some $r_2 > 0$, then $\llbracket \sqrt{e_2} \rrbracket \sigma = \sqrt{r_2}$. The above plan satisfies the condition in Eq. (6), because $y \times \sqrt{r_2 + r_1}$ with a sample y from $\mathcal{N}(0, 1)$ is distributed by $\mathcal{N}(r_1, r_2)$.

We now show that $\bar{\tau}^\pi$ with a valid π preserves semantics. For a command c and $\sigma_\theta \in \text{St}[\theta]$, define the *value function* $v_{c, \sigma_\theta} : \text{St}[\text{Name}] \rightarrow \text{St}[\text{Name}]$ as follows:

$$v_{c, \sigma_\theta}(\sigma_n)(\mu) \triangleq \text{let } \sigma \triangleq \sigma_0 \oplus \sigma_\theta \oplus \sigma_n \text{ in } \begin{cases} \llbracket c \rrbracket \sigma(\text{val}_\mu) & \text{if } \llbracket c \rrbracket \sigma \in \text{St} \text{ and } \llbracket c \rrbracket \sigma(\text{cnt}_{\mu'}) \leq 1 \text{ for all } \mu' \\ 0 & \text{otherwise} \end{cases}$$

where $\sigma_0 \in \text{St}[(\text{PVar} \setminus \theta) \uplus \text{AVar}]$ maps *like* to 1, and pr_μ to $\mathcal{N}(\sigma_n(\mu); 0, 1)$ and val_μ to $\sigma_n(\mu)$ for every $\mu \in \text{Name}$, and it also maps all the other variables to 0. The value function basically applies the lambda functions in c 's sample commands to the corresponding random variables. The next theorem proves that if π is valid, the program transformation $\bar{\tau}^\pi$ preserves the semantics in the sense that the integral of a function h remains the same under c and \bar{c}^π for any c . Note that the two integrals in the theorem are connected via the value functions of c and \bar{c}^π .

THEOREM 4.2. *Let π be a valid reparameterisation plan, and c be a command. Then, for all $\sigma_\theta \in \text{St}[\theta]$ and all measurable $h : \text{St}[\text{Name}] \rightarrow \mathbb{R}$, we have*

$$\int d\sigma_n \left(p_{c, \sigma_\theta}(\sigma_n) \cdot h(v_{c, \sigma_\theta}(\sigma_n)) \right) = \int d\sigma_n \left(p_{\bar{c}^\pi, \sigma_\theta}(\sigma_n) \cdot h(v_{\bar{c}^\pi, \sigma_\theta}(\sigma_n)) \right)$$

where the left integral is defined if and only if the right integral is defined.

Remark 4.3. One immediate yet important consequence of the theorem is that if p_{c, σ_θ} is a probability density, so is $p_{\bar{c}^\pi, \sigma_\theta}$. This consequence will be used in §4.2 and the proof of Theorem 4.5 later. \square

4.2 Gradient Estimator via Program Transformation

Let c be a command that always terminates and does not have a double-sampling error, and let $\sigma_\theta \in \text{St}[\theta]$. We define the *partial density function* $p_{c, \sigma_\theta}^{(S)}$ of c over a subset $S \subseteq \text{Name}$ as

$$p_{c, \sigma_\theta}^{(S)} : \text{St}[\text{Name}] \rightarrow (0, \infty), \quad p_{c, \sigma_\theta}^{(S)}(\sigma_n) \triangleq \prod_{\mu \in S} \llbracket c \rrbracket (\sigma_0 \oplus \sigma_\theta \oplus \sigma_n)(pr_\mu),$$

where σ_0 is set as in the definition of p_{c, σ_θ} in Eq. (3). The partial density $p_{c, \sigma_\theta}^{(S)}$ is essentially the full density p_{c, σ_θ} in Eq. (3) with the omission of the factors not mentioned in S . Intuitively, it computes the density of the random variables in S conditioned on the random variables outside of S .

The SPGE computes an approximate gradient of the objective \mathcal{L} in Eq. (5) using the program transformation in the previous subsection. Its inputs are a model c_m , a guide c_g , parameters θ to optimise, and a reparameterisation plan π , where

- c_m , c_g , and $\overline{c_g}^\pi$ always terminate and do not have a double-sampling error, and
 - c_g defines the normalised probability density p_{c_g, σ_θ} for all $\sigma_\theta \in \text{St}[\theta]$.
- (7)

Given these inputs, the SPGE computes an approximate gradient in three steps. First, it defines the set $rv(\pi) \subseteq \text{Name}$ of random variables to be reparameterised:

$$rv(\pi) \triangleq \{(\alpha, i) \in \text{Name} \mid (\text{name}(\alpha, _), _, _) \in \text{dom}(\pi)\}$$

where $_$ means some existentially quantified (meta) variable. Second, the SPGE transforms the guide c_g to $\overline{c_g}^\pi$, and draws a sample $\hat{\sigma}_n$ from $p_{\overline{c_g}^\pi, \sigma_\theta}$.¹⁰ Drawing a sample $\hat{\sigma}_n$ makes sense here since $p_{\overline{c_g}^\pi, \sigma_\theta}$ is a probability density (i.e., it normalises to 1) by Remark 4.3. Another important point is that drawing $\hat{\sigma}_n$ can be done simply by executing $\overline{c_g}^\pi$ in the standard sampling semantics (not in our density semantics), where each sample command is interpreted as a random draw, not as a density calculator. Third, the SPGE computes the following approximation of $\nabla_\theta \mathcal{L}(\sigma_\theta)$ and returns it as a result:

$$\begin{aligned} \text{grad_est}(\sigma_\theta; \hat{\sigma}_n) \triangleq & (\nabla_\theta \log p_{c_g, \sigma_\theta}^{(\text{Name} \setminus rv(\pi))}(\sigma'_n)) \cdot \log(p_{c_m, \sigma_\theta}(\sigma'_n) / p_{c_g, \sigma_\theta}(\sigma'_n)) \\ & - \nabla_\theta \log p_{c_g, \sigma_\theta}^{(rv(\pi))}(\sigma'_n) + \nabla_\theta \log p_{c_m, \sigma_\theta}(\sigma'_n), \quad \text{for } \sigma'_n \triangleq v_{\overline{c_g}^\pi, \sigma_\theta}(\hat{\sigma}_n). \end{aligned} \quad (8)$$

Recall that if a command c always terminates, both the partial density $p_{c, \sigma_\theta}^{(S)}(\sigma_n)$ and the full density $p_{c, \sigma_\theta}(\sigma_n)$ can be computed simply by executing c in our semantics and calculating the defining formulas of both densities from the final state of the execution. Thus, all the terms in grad_est can be computed by executing c_g and c_m according to our density semantics or differentiating the results of these executions via, for instance, automatic differentiation as done in Pyro. Note that grad_est applies two non-trivial optimisations, when compared with the (naive) SPGE explained in Eq. (2): its first term involves a partial density of c_g instead of the full density of $\overline{c_g}^\pi$, and its second term involves (again) a partial density of c_g instead of the full density of c_g .

Is the SPGE correct in any sense? The answer depends on its inputs. If the inputs satisfy the requirements that we will explain soon, the result of the SPGE is precisely $\nabla_\theta \mathcal{L}(\sigma_\theta)$ on average, that is, $\nabla_\theta \mathcal{L}(\sigma_\theta) = \mathbb{E}[\text{grad_est}(\sigma_\theta; \hat{\sigma}_n)]$, where the expectation is taken over the sample $\hat{\sigma}_n$ used by the SPGE. This property is called *unbiasedness*, and it plays the crucial role for ensuring that parameters updated iteratively with estimated gradients converge to a local optimum.

Let us now spell out the requirements on the inputs of the SPGE. To do so, we need to introduce one further concept for the reparameterisation plans π .

Definition 4.4. A reparameterisation plan π is *simple* if for all $(n, d, \lambda y.e)$ and $(n', d', \lambda y'.e')$ in $\text{NameEx} \times \text{DistEx} \times \text{LamEx}$ such that n and n' have the same string part, we have $(n, d, \lambda y.e) \in \text{dom}(\pi) \iff (n', d', \lambda y'.e') \in \text{dom}(\pi)$. \square

The simplicity is one of the requirements that the SPGE imposes on π . It ensures the following property of the set $rv(\pi)$, which the SPGE relies on when computing grad_est by Eq. (8): $rv(\pi)$ (and $\text{Name} \setminus rv(\pi)$) over-approximates the set of random variables that, if sampled, are (and are not) reparameterised by $\overline{\cdot}^\pi$. Specifically, it forbids π from using any syntax-specific information of the arguments of a sample command when it decides whether to transform the command or not. All the requirements of the SPGE, including the simplicity just explained, are summarised in the next theorem.

THEOREM 4.5. *Let c_m, c_g , and π be the inputs to the SPGE (i.e., they satisfy the assumptions in Eq. (7)). Suppose that $\mathcal{L}(\sigma_\theta)$ and $\nabla_\theta \mathcal{L}(\sigma_\theta)$ are well-defined for every $\sigma_\theta \in \text{St}[\theta]$. Further, assume that every*

¹⁰In practice, the SPGE often draws a fixed number of independent samples $\hat{\sigma}_n^{(1)}, \dots, \hat{\sigma}_n^{(M)}$ from $p_{\overline{c_g}^\pi, \sigma_\theta}$ and computes $\frac{1}{M} \sum_{i=1}^M \text{grad_est}(\sigma_\theta; \hat{\sigma}_n^{(i)})$ as an estimate of $\nabla_\theta \mathcal{L}(\sigma_\theta)$. The presented results hold for this more general case as well.

sample command in c_g has $\lambda y.y$ as its third argument, and c_g does not have observe commands. Then, for all $\sigma_\theta \in \text{St}[\theta]$,

$$\nabla_\theta \mathcal{L}(\sigma_\theta) = \mathbb{E}_{p_{\overline{c}_g^\pi, \sigma_\theta}(\hat{\sigma}_n)} [\text{grad_est}(\sigma_\theta; \hat{\sigma}_n)] \quad (9)$$

if π satisfies the following requirements:

(R1) π is valid and simple.

(R2) The below functions from $\text{St}[\theta] \times \text{St}[\text{Name}]$ to $(0, \infty)$ are differentiable in $\theta \cup \text{rv}(\pi)$ jointly:

$$(\sigma_\theta, \sigma_n) \mapsto p_{c_m, \sigma_\theta}(\sigma_n), \quad (\sigma_\theta, \sigma_n) \mapsto p_{c_g, \sigma_\theta}^{\langle \text{rv}(\pi) \rangle}(\sigma_n), \quad (\sigma_\theta, \sigma_n) \mapsto p_{c_g, \sigma_\theta}^{\langle \text{Name} \setminus \text{rv}(\pi) \rangle}(\sigma_n).$$

(R3) For all $\sigma_n \in \text{St}[\text{Name}]$, the below functions on $\text{St}[\theta]$ are differentiable in θ jointly:

$$\sigma_\theta \mapsto v_{\overline{c}_g^\pi, \sigma_\theta}(\sigma_n), \quad \sigma_\theta \mapsto p_{\overline{c}_g^\pi, \sigma_\theta}^{\langle \text{rv}(\pi) \rangle}(\sigma_n), \quad \sigma_\theta \mapsto p_{\overline{c}_g^\pi, \sigma_\theta}^{\langle \text{Name} \setminus \text{rv}(\pi) \rangle}(\sigma_n).$$

(R4) For all $\sigma_\theta \in \text{St}[\theta]$ and $\sigma_n \in \text{St}[\text{Name}]$, we have $\nabla_\theta p_{\overline{c}_g^\pi, \sigma_\theta}^{\langle \text{rv}(\pi) \rangle}(\sigma_n) = 0$.

(R5) The below equations hold for all $\sigma_\theta \in \text{St}[\theta]$:

$$\begin{aligned} \nabla_\theta \int d\sigma_n \left(p_{\overline{c}_g^\pi, \sigma_\theta}(\sigma_n) \right) &= \int d\sigma_n \nabla_\theta \left(p_{\overline{c}_g^\pi, \sigma_\theta}(\sigma_n) \right), \\ \nabla_\theta \int d\sigma_n \left(p_{\overline{c}_g^\pi, \sigma_\theta}(\sigma_n) \cdot \log \frac{p_{c_m, \sigma_\theta}(\sigma'_n)}{p_{c_g, \sigma_\theta}(\sigma'_n)} \right) &= \int d\sigma_n \nabla_\theta \left(p_{\overline{c}_g^\pi, \sigma_\theta}(\sigma_n) \cdot \log \frac{p_{c_m, \sigma_\theta}(\sigma'_n)}{p_{c_g, \sigma_\theta}(\sigma'_n)} \right). \end{aligned}$$

In the second equation, we write σ'_n for $v_{\overline{c}_g^\pi, \sigma_\theta}(\sigma_n)$.

To be clear, $f : \text{St}[K] \rightarrow \mathbb{R}^n$ for $K \subseteq \text{Var}$ is said to be *differentiable in $K' \subseteq K$ jointly* if for any $\tau \in \text{St}[K \setminus K']$, $f|_{[\tau]} : \text{St}[K'] \rightarrow \mathbb{R}^n$ is (jointly) differentiable, where $f|_{[\tau]}(\sigma) \triangleq f(\sigma \oplus \tau)$.

In this work, we focus on the requirements (R2) and (R3) about smoothness. They require that five density functions of c_m , c_g , and \overline{c}_g^π , and the value function of \overline{c}_g^π be differentiable in certain variables. We will develop a program-analysis framework to check these differentiability requirements soundly and automatically (§5), and will describe an algorithm to the SPGE variable-selection problem, using the developed analysis framework (§6). The remaining requirements (R1), (R4), and (R5) are of less interest in this work. (R1) and (R4) can be guaranteed by simple syntactic checks and our way of constructing π (Lemma D.1). (R5) follows from (R2) and (R3) under a regularity condition on densities (Theorem D.2) which is usually satisfied in practice; in some cases, however, the condition might not hold, and we leave it as future work to automatically discharge the condition (which is about the integrability of local Lipschitz constants of densities) or more generally (R5).¹¹

We point out that Pyro uses the SPGE in their inference engine, but without checking the above requirements. In particular, its default option simply uses the π that transforms all the continuous random variables in a guide, and this can easily violate the requirements and make the SPGE biased.

4.3 Local Lipschitzness for Relaxed Requirements

In Theorem 4.5, we considered the requirements (R2) and (R3) about the differentiability of density and value functions, as a sufficient condition for the unbiasedness of the SPGE. They are, however, sometimes too strong to hold in practice due to the use of popular non-differentiable functions (e.g., ReLU). As we will see in §7, the requirements are indeed violated by some representative Pyro programs even though the conclusion of Theorem 4.5 holds for those programs (i.e., the estimated gradients by the SPGE for those programs are unbiased).

To validate the unbiasedness of the SPGE for more examples in practice, we consider the following relaxation of the requirements (R2) and (R3), which changes differentiability to local Lipschitzness:

(R2') The functions in (R2) are locally Lipschitz in $\theta \cup \text{rv}(\pi)$ jointly.

(R3') For every $\sigma_n \in \text{St}[\text{Name}]$, the functions in (R3) are locally Lipschitz in θ jointly.

¹¹Some sufficient conditions for a regularity condition similar to the one considered here have been studied for certain classes of model-guide pairs, e.g., in [Lee et al. 2020].

Here $f : V \rightarrow \mathbb{R}^m$ for $V \subseteq \mathbb{R}^n$ is *Lipschitz* if there is $C > 0$ such that $\|f(x) - f(x')\|_2 \leq C\|x - x'\|_2$ for all $x, x' \in V$; and f is *locally Lipschitz* if for all $x \in \mathbb{R}^n$, there is an open neighborhood $U \subseteq \mathbb{R}^n$ of x such that $f|_U : U \rightarrow \mathbb{R}^m$ is Lipschitz. Further, $f : \text{St}[K] \rightarrow \mathbb{R}^m$ for $K \subseteq \text{Var}$ is *locally Lipschitz in $K' \subseteq K$ jointly* if for any $\tau \in \text{St}[K \setminus K']$, $f|_{[\tau]} : \text{St}[K'] \rightarrow \mathbb{R}^m$ is locally Lipschitz, where $f|_{[\tau]}(\sigma) \triangleq f(\sigma \oplus \tau)$. Although differentiability does not imply local Lipschitzness, *continuous* differentiability does. Since most differentiable functions used in practice are continuously differentiable, asking for (R2') and (R3') amounts to relaxing the requirements of (R2) and (R3) in practice.

We choose local Lipschitzness as an alternative to differentiability in (R2) and (R3) for two main reasons. First, local Lipschitzness is satisfied by most functions used in practice, which can even be non-differentiable (e.g., ReLU). This would allow us to validate the unbiasedness of the SPGE for more programs, even when they use non-differentiable functions. Second, using local Lipschitzness in (R2) and (R3) does not break the results given in §4.2, even though local Lipschitzness is more practically permissive than differentiability (as explained above). In particular, we have a counterpart of Theorem 4.5 that uses (R2') and (R3') instead of (R2) and (R3):

THEOREM 4.6. *Consider the setup of Theorem 4.5. Then, for all $\sigma_\theta \in \text{St}[\theta]$, Eq. (9) holds if π satisfies the requirements (R1), (R2'), (R3'), (R4), and (R5).*

We can obtain Theorem 4.6 thanks to three properties of local Lipschitzness: locally Lipschitz functions are closed under function composition, have well-defined gradients almost everywhere, and satisfy the chain rule almost everywhere in restricted settings (Lemma E.1). Although the latter two properties are weaker than the corresponding properties of differentiability (i.e., differentiable functions always have well-definedness gradients and satisfy the chain rule), they are still strong enough to prove the theorem.

As in §4.2, we focus on the requirements (R2') and (R3') and less on (R1), (R4), and (R5). Note that (R1) and (R4) can be checked syntactically in the same way as discussed in §4.2. On the other hand, (R5) follows from (R2') and (R3') now under two (instead of one) regularity conditions on densities (Theorem E.2), where the first condition is the one mentioned in §4.2 and the second (new) condition is about some form of almost-everywhere differentiability of densities. Although we believe that the two regularity conditions would usually hold in practice, they can be violated in some cases that involve locally Lipschitz, non-differentiable functions; hence, it would be worthwhile to devise an automatic way of checking the two conditions or more generally (R5), which we leave as future work.

Because local Lipschitzness property has wider coverage than differentiability in practice while ensuring that the results in §4.2 remain valid, our implementation and experiments consider the option of using (R2') and (R3') as well as that of using (R2) and (R3); see §5.3–§7 for details.

5 PROGRAM ANALYSIS FOR SMOOTHNESS

Recall that our goal is to develop an algorithm for the SPGE variable-selection problem in Definition 2.1, which asks for finding a large set S of random variables with a certain property when given a model c_m , a guide c_g , and a reparameterisation plan π_0 . When rephrased using the terminologies that we covered so far, finding such an S amounts to finding a restriction π of the given π_0 such that (c_m, c_g, π) satisfies the requirements in Theorem 4.5 (or 4.6). Thus, the key for developing a desired algorithm for the problem lies in constructing an automatic method for proving that the requirements in Theorem 4.5 (or 4.6), in particular, the smoothness requirements (R2) and (R3) (or (R2') and (R3')) are met. In this section, we propose a program analysis for smoothness properties, which can help find π that meets the smoothness requirements, and which, together with the optimiser in the next section, leads to an algorithm for solving the SPGE variable-selection problem.

We first define a parametric abstraction for smoothness properties (§5.1). We then describe a program analysis based on this abstraction and prove the soundness of the analysis (§5.2). We finally

instantiate the analysis to differentiability and local Lipschitzness (§5.3). The results in this section are not limited to PPLs, but are applicable to general imperative programming languages.

5.1 Parametric Abstraction for Smoothness Properties

At a high level, our parametric abstraction for smoothness properties is built out of two components. The first is a predicate over commands that expresses a target smoothness property but in a conditional form. The predicate is parameterised by two sets of variables, K for the input variables and L for the output variables. Intuitively, the predicate holds for a command if conditioning the input variables outside of K to any fixed values and varying only the ones in K makes the command a smooth function on the output variables in L . Our program analysis tracks a conditional smoothness property formalised by this predicate, and in so doing, it identifies a smooth part of a given command, even when the command fails to be so with respect to some variables.

The second component is also a predicate over commands, but it deals with dependency, instead of smoothness. It is again parameterised by K and L , and expresses that to compute the output variables in L , a command accesses only the input variables in K . Our program analysis tracks dependency formalised by this predicate, so as to achieve high precision, especially when handling sequential composition. To see this, imagine that we want to check the differentiability of a sequence $c; c'$. A natural approach is to use the chain rule. If the dependency-tracking part of our analysis is missing, in order to establish that the output on a variable v by the sequence is differentiable in an input variable u , the analysis should show the output on v by the second command c' is differentiable in *all* variables, and the output on any variable by the first command c is differentiable in u . This requirement on c and c' is too strong. Often, c' uses only a small number of variables to compute v , and it is sufficient to require that just on those used variables, c should be differentiable in u . Similarly, c commonly updates only a small number of variables using u , and it is enough to require that just in those u -dependent variables, the second command c' is differentiable when computing the output v . The dependency-tracking part lets our analysis carry out such reasoning and achieve better precision. Formally, this means our analysis uses a version of reduced product [Cousot and Cousot 1979] between dependency analysis and the analysis that tracks the target smoothness property.

We now formally describe each of these components as well as their combination.

5.1.1 Family of Smoothness Predicates. Our program analysis assumes that a target smoothness property is specified in terms of a family of predicates,

$$\phi = (\phi_{K,L} : K, L \subseteq \text{Var}),$$

where $\phi_{K,L}$ is a set of partial functions from $\text{St}[K]$ to $\text{St}[L]$ (i.e., $\phi_{K,L} \subseteq [\text{St}[K] \rightarrow \text{St}[L]]$).

Example 5.1 (Differentiability). In the instantiation of our program analysis for differentiability, we use the family $\phi^{(d)}$ where for all $K, L \subseteq \text{Var}$, a partial function $f : \text{St}[K] \rightarrow \text{St}[L]$ belongs to $\phi_{K,L}^{(d)}$ if and only if (i) $\text{dom}(f)$ is open and (ii) f is (jointly) differentiable in its domain. \square

At first, one may wonder why we use a family of $\phi_{K,L}$ predicates instead of a single predicate ϕ_0 over $[\text{St} \rightarrow \text{St}_\perp]$. The reason is that, as mentioned above, the analysis aims at a conditional variant of the traditional notion of smoothness. For instance, instead of checking that a function $f : \text{St} \rightarrow \text{St}_\perp$ is differentiable on $\text{St} \setminus f^{-1}(\{\perp\})$, the analysis proves differentiability conditioned on certain variables being fixed: if we fix the input variables in $\text{Var} \setminus K$ and vary just those in K in the initial state, and look at the output variables in L only, then the function f becomes differentiable, although it might not be so when all input/output variables are considered. To express this, we need the whole family of $\phi_{K,L}$ predicates.

This notion of conditional differentiability is similar to, but not the same as, so called partial differentiability. Partial differentiability in K says that, for every $v \in K$, if we fix all the input variables except v , including those in $K \setminus \{v\}$, and consider the output variables in L only, f becomes

differentiable.¹² As we will show in Remark 5.12, the set of partially-differentiable functions is not closed under a certain operator, but we need the closure to ensure that our program analysis is sound. Our conditional differentiability does not suffer from this issue.

5.1.2 Smoothness Abstraction. Based on the family ϕ , we build a predicate Φ that captures the smoothness of commands. The Φ constrains functions from St to St_\perp , unlike $\phi_{K,L}$. For $K, L \subseteq \text{Var}$ with $K \supseteq L$, define $\pi_{K,L}$ to be the projection from $\text{St}[K]$ to $\text{St}[L]$.

Definition 5.2. The *smoothness abstraction* Φ is the predicate over a function $f \in [\text{St} \rightarrow \text{St}_\perp]$ and variable sets $K, L \subseteq \text{Var}$. It is satisfied by (f, K, L) if for all $\tau \in \text{St}[\text{Var} \setminus K]$, the predicate $\phi_{K,L}$ holds for the following partial function $g : \text{St}[K] \rightarrow \text{St}[L]$: $\text{dom}(g) \triangleq \{\sigma \in \text{St}[K] \mid f(\sigma \oplus \tau) \in \text{St}\}$ and $g(\sigma) \triangleq (\pi_{\text{Var},L} \circ f)(\sigma \oplus \tau)$ for $\sigma \in \text{dom}(g)$. We denote the satisfaction of Φ by

$$\models \Phi(f, K, L). \quad \square$$

Note that the function g is constructed from f by fixing the $\text{Var} \setminus K$ part of the input state to τ , and looking at only the L part of the output. This construction is precisely the one used in the informal definition of conditional differentiability described above, and its use reflects the fact that our program analysis attempts to prove a conditional smoothness property.

5.1.3 Dependency Abstraction. Our abstract domain has a component for tracking dependency between input-output variables. Dependency here means that a given input variable is used for computing a given output variable. We define a predicate Δ that has a similar format as Φ . Intuitively, $\Delta(f, K, L)$ holds if and only if the L part of the output of f depends at most on the K part of the input to f . To define Δ formally, for $K \subseteq \text{Var}$, let \sim_K be the following equivalence relation over states: $\sigma \sim_K \sigma'$ if and only if $\sigma(v) = \sigma'(v)$ for all $v \in K$.

Definition 5.3. The *dependency abstraction* Δ is the predicate on $f \in [\text{St} \rightarrow \text{St}_\perp]$ and $K, L \subseteq \text{Var}$ that holds if for all $\sigma, \sigma' \in \text{St}$ with $\sigma \sim_K \sigma'$, we have $(f(\sigma) \in \text{St} \iff f(\sigma') \in \text{St})$ and $(f(\sigma) \in \text{St} \implies f(\sigma) \sim_L f(\sigma'))$. We denote the satisfaction of Δ by

$$\models \Delta(f, K, L). \quad \square$$

5.1.4 Combined Abstraction. We bring together the two abstractions that we just defined, and construct the final abstract domain \mathcal{D}^\sharp used by our program analysis.

Intuitively, each element of \mathcal{D}^\sharp is a predicate on a function $f \in [\text{St} \rightarrow \text{St}_\perp]$ expressed by the conjunction of the following form: $\bigwedge_{i=1}^m \Phi(f, K_i, L_i) \wedge \bigwedge_{j=1}^n \Delta(f, K'_j, L'_j)$. A direct but naive way of implementing this intuition is to let \mathcal{D}^\sharp be the collection of all the constraints of this form, but it permits too many constraints and leads to a costly program analysis. We take a more economical alternative that further restricts the allowed form of the constraints. The alternative requires that the conjunction from above should be constructed out of two mappings p and d from output variables to input variable sets, and a set V of input variables. The p component describes smoothness, and the d and V components dependency. They together encode the constraint

$$\bigwedge_{v \in \text{Var}} \Phi(f, p(v), \{v\}) \wedge \bigwedge_{u \in \text{Var}} \Delta(f, d(u), \{u\}) \wedge \Delta(f, V, \emptyset).$$

Thus, a function $f \in [\text{St} \rightarrow \text{St}_\perp]$ satisfies the constraint encoded by p , d , and V if (i) for every output variable v , when we fix the values of all the input variables outside of $p(v)$, the (partial) function $\sigma \mapsto f(\sigma)(v)$ is smooth (e.g., differentiable); (ii) for every output variable u , the (partial) function $\sigma \mapsto f(\sigma)(u)$ does not access any variable outside of $d(u)$ to compute the value of u ; and (iii) the values of input variables in V determine whether f returns \perp or not.

¹²Conditional differentiability extends partial differentiability in the sense that the latter can be expressed as a conjunction of the former (but not vice versa): f is partially differentiable in K if and only if f is conditionally differentiable in $\{v\}$ for all $v \in K$.

Definition 5.4. The *abstract domain* \mathcal{D}^\sharp consists of triples $(p, d, V) \in [\text{Var} \rightarrow \mathcal{P}(\text{Var})]^2 \times \mathcal{P}(\text{Var})$, called *abstract state*, such that $p(v) \supseteq d(v)^c$ and $d(v) \supseteq V$ for all $v \in \text{Var}$, where $-^c$ is the standard operation for set complement. That is,

$$\mathcal{D}^\sharp \triangleq \{(p, d, V) \in [\text{Var} \rightarrow \mathcal{P}(\text{Var})]^2 \times \mathcal{P}(\text{Var}) \mid p(v) \supseteq d(v)^c \text{ and } d(v) \supseteq V \text{ for all } v \in \text{Var}\}.$$

We order abstract states as follows: $(p, d, V) \sqsubseteq (p', d', V')$ if and only if $V \subseteq V'$ and for all $v \in \text{Var}$, $p(v) \supseteq p'(v)$ and $d(v) \subseteq d'(v)$. These abstract states are concretised by $\gamma : \mathcal{D}^\sharp \rightarrow \mathcal{P}([\text{St} \rightarrow \text{St}_\perp])$:

$$f \in \gamma(p, d, V) \iff \models \Delta(f, V, \emptyset), \models \Phi(f, p(v), \{v\}), \text{ and } \models \Delta(f, d(v), \{v\}) \text{ for all } v \in \text{Var}. \quad \square(10)$$

Note that the definition of \mathcal{D}^\sharp contains two conditions. The first condition $p(v) \supseteq d(v)^c$ comes from our assumption that if a function does not depend on a variable u , it is smooth in u . This and other assumptions of the analysis will be explained shortly in §5.2.2. The other condition $d(v) \supseteq V$ originates from the relationship that if $\Delta(f, K, \{v\})$ holds, so does $\Delta(f, K, \emptyset)$.

Example 5.5 (Differentiability). Consider the setup of Example 5.1 and the program $c \equiv (y := x * x; \text{if } (x \geq 0) \{s := 1\} \text{ else } \{s := -1\})$. Let (p, d, V) be the smallest abstract state that describes the program. In this program, s is not differentiable in x , but y is. So, $p(s) = \text{Var} \setminus \{x\} \supseteq \{y, s\}$ and $p(y) = \text{Var} \supseteq \{x, y, s\}$. Note that $p(s)$ contains the input variables s and y because by not depending on those input variables, the output s is differentiable in those variables. For the dependency part, we have $d(s) = d(y) = \{x\}$ and $V = \emptyset$. \square

5.2 Parametric Static Program Analysis

Our program analysis is based on abstract interpretation [Cousot and Cousot 1977], and computes an approximation of the concrete semantics $\llbracket c \rrbracket$ of a given command c using the abstract domain \mathcal{D}^\sharp . We formalise this computation by the *abstract semantics* of c , which defines $\llbracket c \rrbracket^\sharp \in \mathcal{D}^\sharp$ by induction on the structure of commands, and over-approximates $\llbracket c \rrbracket$ in the sense of the concretization γ in Eq. (10).

5.2.1 Analysis Definition. Fig. 3 shows the abstract semantics of $\llbracket c \rrbracket^\sharp$. The overall structure of the semantics follows the standard compositional semantics of an imperative language. For instance, the abstract semantics of sequential composition is defined in terms of those of constituent commands, and the semantics of a loop is the least fixed point of a monotone operator over \mathcal{D}^\sharp . However, the specifics of the semantics include non-standard details, and we spell them out by going through the defining clauses of $\llbracket c \rrbracket^\sharp$.

The definition of $\llbracket \text{skip} \rrbracket^\sharp$ formalises the effect of `skip` on smoothness and dependency. The definition says that `skip` computes each output variable v in a smooth manner in all input variables, and in so doing, it creates the dependency between the variable v to itself at the input state. The V part of $\llbracket \text{skip} \rrbracket^\sharp$ is the empty set since `skip` always terminates.

The next case is $x := e$. Its abstract semantics records the smoothness and dependency information of the updated variable x by analysing the expression e . For the smoothness part, the semantics invokes the subroutine $\langle\langle e \rangle\rangle^\sharp$ that computes an under-approximation of the set of variables in which the expression e is smooth (and thus over-approximates the smoothness property of e).¹³ For the dependency part, $\llbracket x := e \rrbracket^\sharp$ computes the set of all the free variables of e so as to get an over-approximation of all variables that may affect the value of e . For variables other than x , $\llbracket x := e \rrbracket^\sharp$ behaves like $\llbracket \text{skip} \rrbracket^\sharp$.

The abstract semantics of a sequence $c; c'$ composes those of the sub-commands c and c' . It uses the liftings $f_\cup, f_\cap : \mathcal{P}(\text{Var}) \rightarrow \mathcal{P}(\text{Var})$ of functions f of type $\text{Var} \rightarrow \mathcal{P}(\text{Var})$, which are defined as follows: $f_\cup(V) \triangleq \bigcup_{v \in V} f(v)$ and $f_\cap(V) \triangleq \bigcap_{v \in V} f(v)$. The abstract semantics $\llbracket c; c' \rrbracket^\sharp$ constructs

¹³The subroutine $\langle\langle e \rangle\rangle^\sharp \subseteq \text{Var}$ is defined inductively on e , and differs for different target smoothness properties. For instance, if the target property is differentiability, we have $\langle\langle r \rangle\rangle^\sharp \triangleq \langle\langle x \rangle\rangle^\sharp \triangleq \text{Var}$, $\langle\langle e_1 + e_2 \rangle\rangle^\sharp \triangleq \langle\langle e_1 \rangle\rangle^\sharp \cap \langle\langle e_2 \rangle\rangle^\sharp$, $\langle\langle \text{ReLU}(e) \rangle\rangle^\sharp \triangleq \text{Var} \cap \text{fv}(e)^c$, etc. On the other hand, if the target property is local Lipschitzness, the subroutine changes for some cases: e.g., $\langle\langle \text{ReLU}(e) \rangle\rangle^\sharp \triangleq \langle\langle e \rangle\rangle^\sharp$.

$$\begin{aligned}
\llbracket \text{skip} \rrbracket^\# &\triangleq (\lambda v. \text{Var}, \lambda v. \{v\}, \emptyset), \\
\llbracket x := e \rrbracket^\# &\triangleq ((\lambda v. v \equiv x ? \langle e \rangle^\# : \text{Var}), (\lambda v. v \equiv x ? \text{fv}(e) : \{v\}), \emptyset), \\
\llbracket c; c' \rrbracket^\# &\triangleq \text{let } (p, d, V) \triangleq \llbracket c \rrbracket^\# \text{ and } (p', d', V') \triangleq \llbracket c' \rrbracket^\# \text{ in } (p'', d'', V'') \\
&\text{where } p''(v) \triangleq (V \cup p_\cap(d'(v))^c \cup d_\cup(p'(v)^c))^c, \\
&\quad d''(v) \triangleq V \cup d_\cup(d'(v)), \text{ and } V'' \triangleq V \cup d_\cup(V'), \\
\llbracket \text{if } b \{c\} \text{ else } \{c'\} \rrbracket^\# &\triangleq \text{let } (p, d, V) \triangleq \llbracket c \rrbracket^\# \text{ and } (p', d', V') \triangleq \llbracket c' \rrbracket^\# \text{ in } (p'', d'', V'') \\
&\text{where } p''(v) \triangleq \text{fv}(b)^c \cap p(v) \cap p'(v), \\
&\quad d''(v) \triangleq \text{fv}(b) \cup d(v) \cup d'(v), \text{ and } V'' \triangleq \text{fv}(b) \cup V \cup V', \\
\llbracket \text{while } b \{c\} \rrbracket^\# &\triangleq \text{let } (p, d, V) \triangleq \llbracket c \rrbracket^\# \text{ in } \text{fix } F^\# \\
&\text{where } F^\#(p_0, d_0, V_0) \triangleq (p', d', V'), \\
&\quad p'(v) \triangleq \text{fv}(b)^c \cap (V \cup p_\cap(d_0(v))^c \cup d_\cup(p_0(v)^c))^c, \\
&\quad d'(v) \triangleq \text{fv}(b) \cup (V \cup d_\cup(d_0(v))) \cup \{v\}, \text{ and } V' \triangleq \text{fv}(b) \cup (V \cup d_\cup(V_0)), \\
\llbracket \text{obs}(\text{dist}_N(e_1, e_2), r) \rrbracket^\# &\triangleq (p, d, \emptyset) \\
&\text{where } p(v) \triangleq (v \equiv \text{like}) ? \langle \text{like} \times \text{pdf}_N(r; e_1, e_2) \rangle^\# : \text{Var}, \\
&\quad \text{and } d(v) \triangleq (v \equiv \text{like}) ? \{\text{like}\} \cup \text{fv}(e_1) \cup \text{fv}(e_2) : \{v\}, \\
\llbracket x := \text{sam}(n, \text{dist}_N(e_1, e_2), \lambda y. e') \rrbracket^\# &\triangleq (p, d, \emptyset) \quad \text{for } n = \text{name}(\alpha, r) \text{ with } r \in \mathbb{R} \\
&\text{where } \mu \triangleq \text{create_name}(\alpha, r), \\
p(v) &\triangleq \begin{cases} \langle e'[\mu/y] \rangle^\# & \text{if } v \in \{x, \text{val}_\mu\} \\ \langle \text{pdf}_N(\mu; e_1, e_2) \rangle^\# & \text{if } v \equiv \text{pr}_\mu \\ \langle \text{cnt}_\mu + 1 \rangle^\# & \text{if } v \equiv \text{cnt}_\mu \\ \text{Var} & \text{otherwise,} \end{cases} \\
\text{and } d(v) &\triangleq \begin{cases} \text{fv}(e'[\mu/y]) & \text{if } v \in \{x, \text{val}_\mu\} \\ \{\mu\} \cup \text{fv}(e_1) \cup \text{fv}(e_2) & \text{if } v \equiv \text{pr}_\mu \\ \{v\} & \text{otherwise,} \end{cases} \\
\llbracket x := \text{sam}(n, \text{dist}_N(e_1, e_2), \lambda y. e') \rrbracket^\# &\triangleq (p, d, \emptyset) \quad \text{for } n = \text{name}(\alpha, e) \text{ with } e \notin \mathbb{R} \\
&\text{where } p(v) \triangleq \begin{cases} \text{fv}(e)^c \cap \bigcap_{\mu=(\alpha, _)} \langle e'[\mu/y] \rangle^\# & \text{if } v \equiv x \\ \text{fv}(e)^c \cap \langle e'[\mu/y] \rangle^\# & \text{if } v \equiv \text{val}_\mu \text{ for } \mu = (\alpha, _) \\ \text{fv}(e)^c \cap \langle \text{pdf}_N(\mu; e_1, e_2) \rangle^\# & \text{if } v \equiv \text{pr}_\mu \text{ for } \mu = (\alpha, _) \\ \text{fv}(e)^c \cap \langle \text{cnt}_\mu + 1 \rangle^\# & \text{if } v \equiv \text{cnt}_\mu \text{ for } \mu = (\alpha, _) \\ \text{Var} & \text{otherwise,} \end{cases} \\
\text{and } d(v) &\triangleq \begin{cases} \text{fv}(e) \cup \bigcup_{\mu=(\alpha, _)} \text{fv}(e'[\mu/y]) & \text{if } v \equiv x \\ \text{fv}(e) \cup \{\text{val}_\mu\} \cup \text{fv}(e'[\mu/y]) & \text{if } v \equiv \text{val}_\mu \text{ for } \mu = (\alpha, _) \\ \text{fv}(e) \cup \{\text{pr}_\mu, \mu\} \cup \text{fv}(e_1) \cup \text{fv}(e_2) & \text{if } v \equiv \text{pr}_\mu \text{ for } \mu = (\alpha, _) \\ \text{fv}(e) \cup \{\text{cnt}_\mu\} & \text{if } v \equiv \text{cnt}_\mu \text{ for } \mu = (\alpha, _) \\ \{v\} & \text{otherwise.} \end{cases}
\end{aligned}$$

Fig. 3. Abstract semantics of commands defining $\llbracket c \rrbracket^\# \in \mathcal{D}^\#$.

the dependency part d'' by composing d from $\llbracket c \rrbracket^\#$ and d' from $\llbracket c' \rrbracket^\#$ after lifting the former. Note the inclusion of the set V in the definition of d'' . This is to account for the case that $d'(v)$ in the definition is the empty set; in that case, $d_\cup(d'(v))$ is empty as well and does not have any information about termination. The smoothness part p'' of $\llbracket c; c' \rrbracket^\#$ is more involved, and implements the intuition described briefly in §5.1. In order to conclude that input variables in V_0 together smoothly affect an output variable v in the computation of $c; c'$, the p'' considers the intermediate state after the first command c , and forms two groups of variables at that intermediate state: $d'(v)$ and $p'(v)^c$. Note that

the desired smoothness property for the input variables in V_0 and the output variable v may fail if the first command c uses some variable $u_0 \in V_0$ non-smoothly to update a variable u'_0 in $d'(v)$, or it uses some variable $u_1 \in V_0$ to compute the value of a variable $u'_1 \in p'(v)^c$. In the former case, the non-smoothness of c causes an issue, and in the latter case, the non-smoothness of c' causes an issue. The p'' collects the input variables that avoid these two failure modes and also do not influence the termination of the sequence. As we show in our soundness theorem, doing so is sufficient because it amounts to using a version of chain rule for the target smoothness property.

The abstract semantics of an if command conservatively assumes that any variable in its condition b may affect the value of any output variable (by influencing whether the true or false branch of the command gets executed) and this influence is potentially non-smooth. For every output variable v , the smooth set $p''(v)$ for the if command implements this assumption by excluding free variables in b , and the computed dependency set does the same but this time by including free variables in b .

The abstract semantics of a loop computes the least fixed point of a monotone operator $F^\# : \mathcal{D}^\# \rightarrow \mathcal{D}^\#$ using the standard Kleene iteration. The operator $F^\#$ describes the effect of one iteration of the loop, and it is derived from the standard loop unrolling and our abstract semantics of sequencing and the if command.

The abstract semantics of an observe command $\text{obs}(\text{dist}_N(e_1, e_2), r)$ uses the fact that the command has the same concrete semantics as the assignment $\text{like} := \text{like} \times \text{pdf}_N(r; e_1, e_2)$, where pdf_N is the density function of the normal distribution. The semantics computes (p, d, V) according to that of the assignment, which we explained earlier.

The final case is the abstract semantics of a sample command. The semantics performs a case analysis on the first argument of sam . If it is a constant expression not involving any variables, then the abstract semantics constructs the name μ of the sampled random variable, and updates p, d , and V according to the concrete semantics of the command. Otherwise, the abstract semantics acknowledges that the precise name μ of the random variable cannot be known statically, and performs so called *weak update* by joining two pieces of information before and after the update of the command in the concrete semantics. Note that the abstract semantics does not require the third argument of sam should be the identity function. The ability of dealing with a general function in the third argument is needed since our analysis is intended to be applied to programs after the transformation of the SPGE, which may introduce such an argument.

The abstract semantics is well-defined under the following relatively weak assumption:

ASSUMPTION 1 (EXPRESSION ANALYSIS AND FREE VARIABLES). $\llbracket e \rrbracket^\# \supseteq \text{fv}(e)^c$ for all expressions e .

This assumption is satisfied by the instantiations of the semantics with differentiability and local Lipschitzness, which are used in our implementation. It will be assumed in the rest of the paper.

THEOREM 5.6. *If Assumption 1 holds, then for all commands c , we have $\llbracket c \rrbracket^\# \in \mathcal{D}^\#$, that is, when we let $(p, d, V) \triangleq \llbracket c \rrbracket^\#$, we have $p(v) \supseteq d(v)^c$ and $d(v) \supseteq V$ for all variables $v \in \text{Var}$.*

Example 5.7 (Differentiability). Consider the differentiability property and the example program of Example 5.5. Let (p_1, d_1, V_1) and (p_2, d_2, V_2) be the results of analysing the first assignment command $y := x * x$ and the following if command of the program. Then,

$$(p_1, d_1, V_1) = (\lambda v. \text{Var}, \lambda v. (v \equiv y) ? \{x\} : \{v\}, \emptyset), \quad (p_2, d_2, V_2) = (\lambda v. \{x\}^c, \lambda v. (v \equiv s) ? \{x\} : \{x, v\}, \{x\}).$$

Let (p, d, V) be the analysis result for the entire program. Then,

$$p(v) = (V_1 \cup (p_1)_{\cap} (d_2(v))^c \cup (d_1)_{\cup} (p_2(v)^c))^c = d_1(x)^c = \{x\}^c, \quad V = V_1 \cup (d_1)_{\cup} (V_2) = \{x\}.$$

Also, $d(v) = V_1 \cup (d_1)_{\cup} (d_2(v)) = \{x\}$. As shown in Fig. 3, the variable x that may affect the condition expression of the if command is removed from the smoothness sets, and $p(s) = p(y) = \{x\}^c$. Note that this result is conservative with respect to y . \square

5.2.2 Analysis Soundness and Assumptions. The soundness of our analysis states that for every command c , its abstract semantics $\llbracket c \rrbracket^\# \in \mathcal{D}^\#$ over-approximates the concrete semantics $\llbracket c \rrbracket$ via γ : $\llbracket c \rrbracket \in \gamma(\llbracket c \rrbracket^\#)$. The soundness is conditioned on Assumption 1 and six new assumptions. One of the new assumptions is about the soundness of $\langle e \rangle^\#$. The remaining five assumptions are concerned with the predicate family for the target smoothness property $\phi = (\phi_{K,L} : K, L \subseteq \text{Var})$, and say that certain canonical operators are smooth according to ϕ so that using them in the abstract semantics should not cause an issue. In this subsection, we present the six assumptions one by one, and sketch how those assumptions are related to the soundness.

We start with the assumption that the analysis of each expression $\langle e \rangle^\#$ under-approximates the set of variables in which the evaluation of e is smooth (and thus over-approximates the smoothness property of e).

ASSUMPTION 2 (EXPRESSION ANALYSIS SOUNDNESS). *For all expressions e , variables x , subsets $K, L \subseteq \text{Var}$, and states $\tau \in \text{St}[\text{Var} \setminus K]$ such that $K = \langle e \rangle^\#$ and $L = \{x\}$, if we let $g : \text{St}[K] \rightarrow \text{St}[L]$ be the function defined by $g(\sigma) \triangleq [x \mapsto \llbracket e \rrbracket(\sigma \oplus \tau)]$, the function g satisfies $\phi_{K,L}$ (i.e., $g \in \phi_{K,L}$).*

This assumption is used in our soundness argument whenever the abstract semantics uses $\langle e \rangle^\#$ for computing smoothness information about an expression e .

The next two assumptions assert the smoothness of the standard operators on the product spaces.

ASSUMPTION 3 (PROJECTION). *For all $K, L \subseteq \text{Var}$ with $K \supseteq L$, the projection $\pi_{K,L}$ satisfies $\phi_{K,L}$.*

ASSUMPTION 4 (PAIRING). *For all $K, L, M \subseteq \text{Var}$ with $L \cap M = \emptyset$, if $f \in \phi_{K,L}$ and $g \in \phi_{K,M}$, we have $\langle f, g \rangle \in \phi_{K, L \cup M}$, where $\langle f, g \rangle$ is the pairing of two partial functions: $\langle f, g \rangle(\sigma) \triangleq$ if $(\sigma \in \text{dom}(f) \cap \text{dom}(g))$ then $f(\sigma) \oplus g(\sigma)$ else undefined.*

Note that $\text{St}[L \cup M]$ is isomorphic to $\text{St}[L] \times \text{St}[M]$, the product space that we referred to above. The assumptions say that the projection is smooth, and the pairing of smooth functions is smooth. Our analysis uses Assumption 3 to deal with variables not modified by a command. For instance, when analysing an assignment $x := e$, the analysis uses Assumption 3 and concludes that on every output variable v other than x , the assignment is smooth in all the input variables. Assumption 4 is used to justify the handling of a sequence $c; c'$ by our analysis, in particular, the part that the analysis combines smoothness information over multiple output variables after the first command c .

The projection and pairing assumptions are about how shrinking and expanding output variables affect the target smoothness property. The next restriction assumption is about shrinking the input variables. It validates the weakening of the K part of $\models \Phi(f, K, L)$, and is used in the abstract semantics of $c; c'$ (and other composite commands).

ASSUMPTION 5 (RESTRICTION). *For all $K, K', L \subseteq \text{Var}$ with $K \supseteq K'$, and $\tau \in \text{St}[K \setminus K']$, if $f \in \phi_{K,L}$, then we have $g \in \phi_{K',L}$, where $g(\sigma) \triangleq$ if $(\sigma \oplus \tau \in \text{dom}(f))$ then $f(\sigma \oplus \tau)$ else undefined.*

The following assumption says that the function composition preserves smoothness. It is related to the chain rule for differentiation, and used to justify the abstract semantics of a sequence $c; c'$.

ASSUMPTION 6 (COMPOSITION). *For all $K, L, M \subseteq \text{Var}$, if $f \in \phi_{K,L}$ and $g \in \phi_{L,M}$, we have $g \circ f \in \phi_{K,M}$, where $g \circ f$ is the standard composition of two partial functions: $(g \circ f)(\sigma) \triangleq$ if $(\sigma \in \text{dom}(f) \wedge f(\sigma) \in \text{dom}(g))$ then $g(f(\sigma))$ else undefined.*

The final assumption lets the analysis infer smoothness information about the completely-undefined function. It is used to justify the handling of loops by our analysis.

ASSUMPTION 7 (STRICTNESS). *For all $K, L \subseteq \text{Var}$, we have $(\lambda \sigma \in \text{St}[K]. \text{undefined}) \in \phi_{K,L}$.*

THEOREM 5.8 (SOUNDNESS). *If Assumptions 1–7 hold, the analysis computes the sound abstraction of the concrete semantics of commands in the following sense: for all commands c ,*

$$\llbracket c \rrbracket \in \gamma(\llbracket c \rrbracket^\sharp).$$

Remark 5.9. A standard method for proving a property of a loop or more generally a recursively defined function is so called Scott induction. In this method, we view a property as a set \mathcal{T} of state transformers and a loop as the least fixed point of a continuous function F on state transformers. Then, we prove the three conditions: (i) \mathcal{T} contains the least state transformer, (ii) it is closed under the least upper bound of any increasing sequence of state transformers, and (iii) \mathcal{T} is preserved by F . The first and second conditions are called strictness and admissibility, respectively, and these three conditions imply that the least fixed point of F belongs to \mathcal{T} .

Our soundness proof for the loop case deviates slightly from this standard method. If it followed the method instead, we would need, in addition to the strictness assumption, the following assumption, which corresponds to the second admissibility condition:

ASSUMPTION 8 (ADMISSIBILITY). *Let $K, L \subseteq \text{Var}$, and order partial functions in $[\text{St}[K] \rightarrow \text{St}[L]]$ by the inclusion of the graphs of partial functions. Then, for every increasing sequence $\{f_n : \text{St}[K] \rightarrow \text{St}[L]\}_{n \in \mathbb{N}}$ (i.e., the graph of f_{n+1} includes that of f_n for all $n \in \mathbb{N}$), if every f_n satisfies $\phi_{K,L}$, so does the least upper bound f_∞ of the sequence (defined by its graph being the union of the graphs of all f_n 's).*

The inclusion of this admissibility assumption would, then, limit the applicability of our program analysis, since some well-known smoothness properties, such as (global) Lipschitz continuity and local boundedness, do not satisfy the assumption, although they satisfy our five assumptions (Assumptions 3–7) (see Table 2). On the plus side, the inclusion of the admissibility assumption could enable our analysis to handle loops more accurately, possibly by tracking the impact of the boolean condition of each loop on smoothness more precisely. Our soundness proof avoids the admissibility assumption by exploiting the fact that our analysis handles loop conditions conservatively: our analysis drops all the variables that loop conditions may depend on from the set of smooth variables, so that it avoids finding too precise inductive predicates that can break soundness.¹⁴ \square

5.3 Instantiations

Our program analysis requires that the family of smoothness predicates should satisfy Assumptions 3–7. Although these assumptions are violated by some smoothness properties, such as partial differentiability and partial continuity, they are met by our leading example $\phi^{(d)}$ for differentiability (Example 5.1), and also by the predicate family $\phi^{(l)}$ for local Lipschitzness, which is used in our implementation. Recall the definitions of the predicate families $\phi^{(d)}$ and $\phi^{(l)}$: for all $K, L \subseteq \text{Var}$,

$$\phi_{K,L}^{(d)} \triangleq \{f : \text{St}[K] \rightarrow \text{St}[L] \mid \text{dom}(f) \text{ is open and } f \text{ is (jointly) differentiable in its domain}\},$$

$$\phi_{K,L}^{(l)} \triangleq \{f : \text{St}[K] \rightarrow \text{St}[L] \mid \text{dom}(f) \text{ is open, and for all } \sigma \in \text{dom}(f), \text{ there are } C > 0 \text{ and an open } O \subseteq \text{dom}(f) \text{ s.t. } \sigma \in O \text{ and } \|f(\sigma_0) - f(\sigma_1)\|_2 \leq C\|\sigma_0 - \sigma_1\|_2 \text{ for all } \sigma_0, \sigma_1 \in O\}.$$

THEOREM 5.10. *Both $\phi^{(d)}$ and $\phi^{(l)}$ satisfy Assumptions 3–7.*

Remark 5.11. The requirement of open domain in $\phi^{(d)}$ is sometimes too constraining and hurts the accuracy of the analysis. It can, however, be relaxed, and we can generalise $\phi^{(d)}$ to the following

¹⁴To be precise, our analysis does not require the admissibility assumption, not because our abstract domain is finite, but because given a loop, our analysis finds an inductive predicate that is “sufficiently admissible” in the sense that it is closed under the least upper bound of any chain *that matters for soundness*: the chain should be definable by some loop. More details are in §F.4.

Table 1. Failure cases of the composition assumption. For each given ϕ , we have $f, g \in \phi$ but $g \circ f \notin \phi$, where f and g are interpreted as (total or partial) functions from $\text{St}[K]$ to $\text{St}[L]$. Let $c_1 \equiv (y = x^2; z = g(y))$ and $c_2 \equiv (y = x; z = g(x, y))$. Then, for each i -th ϕ , $\llbracket c_i \rrbracket^\sharp$ incorrectly concludes that z is smooth with respect to x .

ϕ	f	g
$\phi_{K,L}^{(d'')}$	$f(x) = x^2$ defined on \mathbb{R}	$g(x) = 1_{[x>0]}$ defined on $[0, 1]$
$\phi_{K,L}^{(pd)}, \phi_{K,L}^{(pc)}$	$f(x) = (x, x)$ defined on \mathbb{R}	$g(x, y) = \begin{cases} xy/(x^2 + y^2) & \text{if } (x, y) \neq (0, 0) \\ 0 & \text{otherwise} \end{cases}$ defined on \mathbb{R}^2

predicate family $\phi^{(d')}$, which corresponds to the standard definition of differentiability on a manifold with boundary in differential geometry [Lee 2012, Chapter 2]:

$$\phi_{K,L}^{(d')} \triangleq \{f : \text{St}[K] \rightarrow \text{St}[L] \mid \text{for all } \sigma \in \text{dom}(f), \text{ there exist an open } U \subseteq \text{St}[K] \text{ and } g : U \rightarrow \text{St}[L] \text{ such that } \sigma \in U, f = g \text{ on } U \cap \text{dom}(f), \text{ and } g \text{ is (jointly) differentiable}\}.$$

Note the weakening of open-domain requirement in $\phi^{(d')}$: the open domain U in the above definition does not have to be included in $\text{dom}(f)$. The family $\phi^{(d')}$ satisfies Assumptions 3–7, and can lead to a more permissive instantiation of our program analysis than the family $\phi^{(d)}$, especially in handling atomic commands, such as assignment, sample, and observe. We point out that $\phi^{(d')}$ does not satisfy Assumption 8 (i.e., the admissibility assumption), while $\phi^{(d)}$ does satisfy it. As a result, if the handling of loops in our analysis is changed such that loop conditions are analysed more accurately, the analysis may remain sound only for $\phi^{(d)}$, not for $\phi^{(d')}$, as explained in Remark 5.9. \square

Remark 5.12. At this point, the reader might feel that Assumptions 3–7 are satisfied by nearly all smoothness properties. This impression is not accurate. For instance, the composition assumption does not hold for the notions of differentiability of partial functions formalised by the following $\phi^{(d'')}$ and $\phi^{(pd)}$, nor for the partial continuity formalised by $\phi^{(pc)}$:

$$\begin{aligned} \phi_{K,L}^{(d'')} &\triangleq \{f : \text{St}[K] \rightarrow \text{St}[L] \mid f \text{ is (jointly) differentiable in the interior of its domain}\}, \\ \phi_{K,L}^{(pd)} &\triangleq \{f : \text{St}[K] \rightarrow \text{St}[L] \mid \text{dom}(f) \text{ is open, and for all } v \in K, f \text{ is partially differentiable in } v\}, \\ \phi_{K,L}^{(pc)} &\triangleq \{f : \text{St}[K] \rightarrow \text{St}[L] \mid \text{dom}(f) \text{ is open, and for all } v \in K, f \text{ is partially continuous in } v\}. \end{aligned}$$

Table 1 contains counterexamples that show the failure of the composition assumption for these predicate families. In fact, when instantiated with these families, our program analysis is not sound. The same table shows example programs and incorrect conclusions derived by our analysis.

Table 2 shows more target smoothness properties from mathematics, and whether each property satisfies Assumptions 3–7 (and Assumption 8). Recall that our program analysis does not require Assumption 8 for soundness; the table shows the assumption just for reference. The target properties from “cont.” to “real analytic” in the table (and three more) satisfy Assumptions 3–7, so that our analysis can be immediately applied to those target properties while remaining sound. \square

6 ALGORITHM FOR THE SPGE VARIABLE-SELECTION PROBLEM

We now put together the results from §4 and §5 to formally define and soundly (yet approximately) solve the SPGE variable-selection problem. We start with the formal definition of the problem:

Definition 6.1 (SPGE Variable-Selection Problem; Formal). Assume we are given a model c_m , a guide c_g , and a (initial) simple reparameterisation plan π_0 such that c_m , c_g , and \bar{c}_g^π always terminate and have no double-sampling errors for all $\pi \sqsubseteq \pi_0$. Here we write $\pi \sqsubseteq \pi'$ if the graph of π is included in that of π' . Given these c_m , c_g , and π_0 , find a reparameterisation plan $\pi \sqsubseteq \pi_0$ such that (i) π is simple and satisfies (R2) and (R3) in §4.2, and (ii) $|rv(\pi)|$ is maximised. We say that π is a *sound solution* if it satisfies (i), and an *optimal solution* if it satisfies (i) and (ii). \square

Table 2. Various well-known target smoothness properties from mathematics, and whether each of them satisfies Assumptions 3–7 (and Assumption 8). Here “cont.” and “diff.” denote “continuous” and “differentiable”. The properties above the double horizontal line are defined such that they include the open-domain requirement as in $\phi^{(d)}$ and $\phi^{(l)}$, and the properties below the line are defined without the open-domain requirement.

Target smoothness property	A3 (proj.)	A4 (pair.)	A5 (rest.)	A6 (comp.)	A7 (stri.)	A8 (admi.)
cont. (C^0)	○	○	○	○	○	○
locally Lipschitz (= $\phi^{(l)}$)	○	○	○	○	○	○
uniformly cont.	○	○	○	○	○	✗
Lipschitz cont.	○	○	○	○	○	✗
diff. (= $\phi^{(d)}$)	○	○	○	○	○	○
continuously diff. (C^1)	○	○	○	○	○	○
smooth (C^∞)	○	○	○	○	○	○
real analytic (C^ω)	○	○	○	○	○	○
partially cont. (= $\phi^{(pc)}$)	○	○	○	✗	○	○
partially diff. (= $\phi^{(pd)}$)	○	○	○	✗	○	○
almost-everywhere cont.	○	○	✗	✗	○	○
almost-everywhere diff.	○	○	✗	✗	○	○
coordinatewise non-decreasing	○	○	○	○	○	○
locally bounded	○	○	○	○	○	✗
bounded	✗	○	○	○	○	✗
Borel measurable	○	○	○	○	○	○
locally integrable	○	○	✗	✗	○	✗
integrable	✗	○	✗	✗	○	✗

The input π_0 in the problem is a newcomer. It fixes a semantics-preserving transformation for all the sample commands. Typically, π_0 is defined on the entire $\text{NameEx} \times \text{DistEx} \times \text{LamEx}$, and remains fixed across all input model-guide pairs (c_m, c_g) . More importantly, it is valid so that the change of any sample command by π_0 preserves the semantics of the command when we take into account both the second distribution argument and the third lambda argument of the sample command. The validity of π_0 is inherited by any sound solution π of the SPGE variable-selection problem since validity as a property on reparameterisation plans is down-closed with respect to the \sqsubseteq order. In our setup, π_0 is fixed to be the following reparameterisation plan from §4.1:

$$\pi_0(n, \text{dist}_N(e_1, e_2), \lambda y.e_3) \triangleq (\text{dist}_N(0, 1), \lambda y.e_3[(y \times \sqrt{e_2} + e_1)/y]) \quad (11)$$

for all $n \in \text{NameEx}$ and expressions e_1, e_2 , and e_3 .

As an example of the SPGE variable-selection problem, consider the problem for the π_0 in Eq. (11) and the model-guide pair (c_m, c_g) given in Fig. 1, where “ z_1 ” in the figure is interpreted as $\text{name}("z_1", 0)$. Then, as discussed in §2, the problem has the following optimal solution: $\pi \triangleq \pi_0|_{S \times \text{DistEx} \times \text{LamEx}}$ for $S \triangleq \{\text{name}(\alpha, e) \in \text{NameEx} \mid \alpha \neq "z_2"\}$.

We present an algorithm for computing a sound (yet possibly suboptimal) solution to the problem.

- (1) By running our program analysis instantiated with differentiability (described in §5.2 and §5.3), compute $(\mathbb{P}_m, \mathbb{D}_m, \mathbb{V}_m) \triangleq \llbracket c_m \rrbracket^\#$ and $(\mathbb{P}_g, \mathbb{D}_g, \mathbb{V}_g) \triangleq \llbracket c_g \rrbracket^\#$, where we use \mathbb{P} , \mathbb{D} , and \mathbb{V} for the output of the analysis to distinguish them from densities p and distributions d .
- (2) Using \mathbb{P}_m and \mathbb{P}_g , check

$$\theta \subseteq K, \quad \text{where } K \triangleq \mathbb{P}_m(\text{like}) \cap \bigcap_{\mu \in \text{Name}} \mathbb{P}_m(p r_\mu) \cap \bigcap_{\mu \in \text{Name}} \mathbb{P}_g(p r_\mu). \quad (12)$$

If the check fails, return an error message that our analysis cannot discharge (R2) for any π , since the analysis concludes that the density function of c_m or c_g can be non-differentiable in θ (even when $rv(\pi) = \emptyset$). If the check passes, initialise the set of reparameterised random variables by

$$S \triangleq \{(\alpha, i) \in \text{Name} \mid \text{for all } i' \in \mathbb{N}, (\alpha, i') \in \text{Name} \implies (\alpha, i') \in K\}.$$

- (3) Using S and π_0 , construct a reparameterisation plan $\pi \sqsubseteq \pi_0$ by $\pi \triangleq \pi_0[S]$, where $\pi_0[S](n, d, l)$ is $\pi_0(n, d, l)$ if $(n, d, l) \in \text{dom}(\pi_0)$, $n = \text{name}(\alpha, _)$, and $(\alpha, _) \in S$; otherwise, it is undefined.
- (4) By running the differentiability analysis on $\overline{c_g^\pi}$, compute $(\overline{\mathbb{P}}_g, \overline{\mathbb{D}}_g, \overline{\mathbb{V}}_g) \triangleq \llbracket \overline{c_g^\pi} \rrbracket^\#$ and check

$$\theta \subseteq \bigcap_{\mu \in \text{Name}} \overline{\mathbb{P}}_g(\text{pr}_\mu) \cap \bigcap_{\mu \in \text{Name}} \overline{\mathbb{D}}_g(\text{val}_\mu). \quad (13)$$

If the check passes, return π as the output of the algorithm. If not, update S by $S \setminus \{(\alpha, i) \in \text{Name}\}$ after choosing some $(\alpha, _) \in S$, and then repeat the above procedure (from the step (3), the point where we construct π using S) until S becomes empty.

Our algorithm computes a sound solution (in the sense stated in Definition 6.1), partly because of the soundness of our program analysis:

THEOREM 6.2. *Let c_m, c_g , and π_0 be the inputs to the SPGE variable-selection problem. If the above algorithm returns π for (c_m, c_g, π_0) , then π is a sound solution for the problem.*

We point out that the theorem holds for the local Lipschitzness case as well (in addition to the differentiability case). That is, if the above algorithm runs our program analysis instantiated with local Lipschitzness (instead of differentiability), and if it returns an output π , then π is a sound solution to the SPGE variable-selection problem that uses (R2') and (R3') (instead of (R2) and (R3)).

Our algorithm solves the problem only approximately: there is no formal guarantee that it always computes an optimal solution. The suboptimality may arise due to two approximations: the over-approximation of our program analysis when it computes differentiability (or local Lipschitzness) information, and the heuristic choices made by our algorithm when the algorithm computes the random-variable set S . We demonstrate, however, that our algorithm finds optimal solutions for all the benchmarks in §7.

Our algorithm calls our program analysis at most $|\{\alpha \in \text{Str} \mid (\alpha, _) \in S_0\}| + 2$ times, where S_0 is the initial value of S (i.e., the set of random variables whose sample commands are to be transformed) in the algorithm. However, for all the benchmarks in §7, our algorithm terminated with the initial set S_0 and thus called our analysis only 3 times (on the model, the guide, and the reparameterised guide according to S_0). Based on these results and our intuition on the algorithm, we conjecture that our algorithm always terminates with the initial set S_0 under a mild condition on π_0 and our analysis (e.g., $(e)^\#$ is computed inductively on e). Since the conjecture is still open, our algorithm might not succeed in the first iteration, and if so, it continues to search for a sound solution greedily. Note that there are many other ways to continue the search and our algorithm uses just one of them (as it is linear-time).

7 EXPERIMENTAL EVALUATION

In our experiments, we consider two research questions. First, can the analysis proposed in §5 be instantiated and implemented so that it can produce meaningful smoothness results on real-world probabilistic programs? Second, can the algorithm proposed in §6 find near-optimal solutions to the SPGE variable-selection problem on real-world probabilistic programs? To assess the two questions, we have implemented a static smoothness analyser for Pyro programs based on §5, and a variable selector based on §6 which (approximately) solves the variable-selection problem.¹⁵ Our analyser and variable selector are implemented in OCaml, and support a subset of the Pyro PPL and two smoothness properties: differentiability and local Lipschitzness.

Implementation. Although the analysis described in §5 may look simple when considering a basic PPL, real-world PPLs such as Pyro are of a much higher degree of complexity. First, they provide a large panel of continuous/discrete probability distributions for sample and observe commands, and

¹⁵Our implementation is available at <https://github.com/wonyeol/smoothness-analysis> and [Lee et al. 2022a].

library functions for tensors and neural networks. Second, programs in real-world PPLs may fail to be smooth for reasons other than if-else and while commands. In particular, values sampled from discrete distributions, and arguments to operators and distribution constructors that are well-defined only on a strict subset of values, may induce non-smoothness. A straightforward treatment of these will result in an overly conservative analysis, treating far too many variables as potentially non-smooth. Third, Pyro programs typically rely on tensors (of large, statically unknown size) to deal with large datasets, and it is generally infeasible to reason about each (real-valued) element of tensors individually. In the following, we discuss how our static analyser addresses these issues and provides sound, useful information about smoothness of Pyro programs.

Distributions and library functions. Our analyser supports 17 distributions (continuous or discrete). Each distribution is characterized by a pair (b, a) for a boolean b and an array of booleans a , where b (or a_i) denotes whether its probability density is differentiable or locally Lipschitz with respect to the sampled value (or the i -th argument) of the distribution. For example, a normal distribution is described by $(\text{true}, [\text{true}, \text{true}])$ (assuming that the second argument is positive) and a Poisson distribution by $(\text{false}, [\text{true}])$. Similarly, the analyser supports a large number of PyTorch/Pyro library functions for tensors and neural networks, and assumes the correct smoothness information about them. For instance, the ReLU function is considered locally Lipschitz but not differentiable.

Refining smoothness information based on safety pre-analysis. Although the expression x/y is generally non-smooth with respect to y (even if it is well-defined for $y=0$), if more information is available, for instance that y always lies in range $[1, 10]$, we can safely consider it smooth with respect to both x and y . Likewise, the density of a normal distribution is generally non-smooth with respect to the standard deviation argument σ (even if it is well-defined for $\sigma \leq 0$), so more precise smoothness information can be produced when σ is known to be always positive. Thus, establishing precise smoothness information requires to first establish safety properties related to program operations. To achieve this, our tool actually performs two analyses in sequence: (i) a safety pre-analysis infers ranges over all numerical variables and marks each argument to an operator or a distribution constructor as either “safe” or “potentially unsafe”; (ii) the program analysis formalised in §5 utilises information computed in the first phase to produce precise smoothness information. The first phase boils down to a forward abstract interpretation based on basic abstract domains like intervals and signs [Cousot and Cousot 1977]. It logs safety information for each program statement just like static analyses for runtime errors and undefined behaviors [Blanchet et al. 2003]. As formalised in §5.2, the second analysis is compositional. Due to their different nature, the two analyses need to be done in sequence.

Tensors. Pyro programs commonly use nested loops and indexed tensors. As the number of scalar values in each tensor is often statically unknown (or known but huge), treating each scalar as a separate variable is not feasible; so we rely on a conservative summarisation of tensors. Intuitively, we treat all scalars in a tensor as a “block”: e.g., when density might not be smooth with respect to some scalar(s) of a tensor, the analysis conservatively concludes that it might not be smooth with respect to all scalars in the tensor. In our experiments, this abstraction does not result in any precision loss.

Evaluation. We evaluated our analyser and variable selector on 13 representative Pyro examples from the Pyro webpage [Uber AI Labs 2022] that use standard SVI engines and contain explicitly written model-guide pairs (without AutoGuide). They include advanced models with deep neural networks such as attend-infer-repeat [Eslami et al. 2016] and single-cell annotation using variational inference [Xu et al. 2021]. Additionally, we included the example in Fig. 1, for which Pyro offers an unsound reparameterisation plan. Table 3 lists half of these 14 Pyro examples with their code size and conceptual complexity (see §I for the rest). Experiments were performed on a Macbook Pro with 2.3GHz Core i9 and 32GB RAM.

Table 3. Subset of Pyro examples used in experiments and their key features (see §I for the rest). The last five columns show the total number of code lines (excluding comments), loops, sample commands, observe commands, and learnable parameters (declared explicitly by `pyro.param` or implicitly by a neural network module). Each number is the sum of the counts in the model and guide.

Name	Probabilistic model	LoC	while	sam	obs	param
spnor	Splitting normal example in Fig. 1	16	0	2	1	2
sgdef	Deep exponential family	105	0	12	1	12
dmm	Deep Markov model	112	3	2	1	13
mhmm	Hidden Markov models	137	1	5	5	12
scanvi	Single-cell annotation using variational inference	147	0	7	2	21
air	Attend-infer-repeat	174	2	6	1	16
cvae	Conditional variational autoencoder	205	0	2	1	15

Table 4. Results of smoothness analyses. “Manual” and “Ours” denote the number of continuous random variables and learnable parameters in which the density of the program is smooth, computed by hand and by our analyser. “Time” denotes the runtime of our analyser in seconds. “#CRP” denotes the total number of continuous random variables and learnable parameters in the program. -m and -g denote model and guide. We consider $\{(\alpha, i) \in \text{Name}\}$ as one random variable for each $\alpha \in \text{Name}$. See §I for the rest of results.

Name	Differentiable			Locally Lipschitz			#CRP
	Manual	Ours	Time	Manual	Ours	Time	
spnor-m	1	1	0.006	1	1	0.009	2
spnor-g	4	4	0.007	4	4	0.008	4
sgdef-m	6	6	0.003	6	6	0.006	6
sgdef-g	18	18	0.016	18	18	0.015	18
dmm-m	4	4	0.014	10	10	0.016	10
dmm-g	4	4	0.026	5	5	0.020	5
mhmm-m	10	10	0.063	10	10	0.075	10
mhmm-g	6	6	0.007	6	6	0.008	6
scanvi-m	6	6	0.032	12	12	0.032	12
scanvi-g	8	8	0.052	15	15	0.058	15
air-m	1	1	0.108	4	4	0.105	4
air-g	3	3	0.075	15	15	0.072	16
cvae-m	3	3	0.025	8	8	0.027	8
cvae-g	5	5	0.031	9	9	0.023	9

Smoothness analyser. We assess our smoothness analyser on the 14 Pyro examples for differentiability and local Lipschitzness (§5.3), and show a subset of results in Table 4 (see §I for the rest). The results demonstrate that our analysis can cope successfully with real-world Pyro programs. First, our analysis is accurate. For all examples, the analysis identifies the exact ground-truth set of random variables and parameters in which the density of the program is differentiable (or locally Lipschitz). In many of them, information computed by the pre-analysis is required to achieve these exact results; e.g., some examples (e.g., `dpmm` and `air`) require precise information about which distribution arguments can be proved to be always in the proper range of values. Second, the runtime of our analysis is low. Typical probabilistic programming applications are not of a very large size, and conceptual complexity is generally the main issue, thus the analysis performance presents no scalability concern.

We draw two more observations from the results. First, for `spnor-m` and `air-g`, the density of each program is not locally Lipschitz in one continuous random variable. These non-local-Lipschitznesses arise as follows: for the former, the random variable (“`z2`” in Fig. 1) is used in the branch condition of an if-else command that contains observe commands, thereby creating discontinuity; and for the latter, the random variable (“`zwhere`”) is passed into the denominator of a division operator, thereby causing a division-by-zero error for some value.

Second, for all the other examples, the density is locally Lipschitz in all continuous random variables and parameters, but is often non-differentiable in many parameters (and continuous random

Table 5. Results of variable selections. “Ours-Time” denote the runtime of our variable selector in seconds. “Ours-Sound” and “Pyro \ Ours” denote the number of random variables in the example that are in π_{ours} , and that are in π_0 but not in π_{ours} , respectively, where π_{ours} and π_0 denote the reparameterisation plans given by our variable selector and by Pyro. “Pyro \ Ours” is partitioned into “Sound” and “Unsound”: the latter denotes the number of random variables that make (R2’) or (R3’) violated when added to π_{ours} , and the former denotes the number of the rest. “#CR” and “#DR” denote the total number of continuous and discrete random variables in the example. We consider $\{(\alpha, i) \in \text{Name}\}$ as one random variable for each $\alpha \in \text{Name}$. See §I for the rest of results.

Name	Ours		Pyro \ Ours		#CR	#DR
	Time	Sound	Sound	Unsound		
spnor	0.021	1	0	1	2	0
sgdef	0.034	6	0	0	6	0
dmm	0.054	1	0	0	1	0
mhmm	0.083	2	0	0	2	1
scanvi	0.143	3	0	0	3	1
air	0.247	1	0	1	2	1
cvae	0.063	1	0	0	1	0

variables too); see, for instance, `scanvi` and `cvae`. Due to this, the requirement (R2) is not satisfied for these examples even with the empty reparameterisation plan (corresponding to the score estimator); that is, if we use the differentiability requirements (R2) and (R3) to validate the unbiasedness of gradient estimators, even the score estimator cannot be validated for these examples. From manual inspection, we checked that the non-differentiabilities from these examples all arise by the use of locally Lipschitz but non-differentiable operators (e.g., `relu` and `grid_sample`). Since many practical models (and guides) use locally Lipschitz but non-differentiable operators, this observation strongly suggests that a right smoothness requirement for validating gradient estimators is not differentiability (which has been used as a standard requirement), but rather local Lipschitzness (e.g., (R2’) and (R3’)).

Variable selector. To evaluate our variable selector, we consider the SPGE variable-selection problem with local Lipschitzness requirements, i.e., the problem that uses (R2’) and (R3’) in §4.3 instead of (R2) and (R3) in §4.2. We do not consider the original problem (with differentiability requirements), since for many examples the differentiability requirements are not satisfied even by the empty reparameterisation plan (i.e., score estimator) as observed above. For an initial reparameterisation plan π_0 for the problem, we use the plan given by Pyro’s default variable selector: it is defined for all continuous random variables and applies standard reparameterisations (e.g., Eq. (11) for a normal distribution). In this settings, we apply our variable selector to the problem on the 14 Pyro examples. Table 5 displays the results (only for 7 examples; see §I for the rest) and compares them with π_0 .

The results demonstrate that for all examples, our variable selector finds the optimal reparameterisation plan with a small runtime. We also observe that for all cases, it terminates in the first iteration and calls our smoothness analyser only three times, as mentioned in §6. Note that the reparameterisation plan given by Pyro is also optimal for all but two examples. We emphasise, however, that our variable selector not only finds a reparameterisation plan but also verifies the local Lipschitzness requirements (R2’) and (R3’), whereas Pyro’s default variable selector does not do so. Indeed, for two examples, Pyro’s reparameterisation plan is unsound as it violates the local Lipschitzness requirements. Hence, these results should be interpreted as: for all but two examples, our variable selector (and smoothness analyser) successfully verifies that the default gradient estimator used by Pyro satisfies important smoothness-related preconditions for unbiasedness, namely the local Lipschitzness requirements.

The two examples for which Pyro becomes unsound are `spnor` and `air`. Recall that they have two continuous random variables (one for each) in which their densities are not locally Lipschitz. The unsoundness of Pyro on these examples stems precisely from the fact that it reparameterises the two non-locally-Lipschitz random variables without checking any local Lipschitzness requirements.

8 RELATED WORK

The high-level idea of using program transformation for improved posterior inference and model learning in PPLs has been explored previously [Claret et al. 2013; Gorinova et al. 2020; Nori et al. 2014; Ritchie et al. 2016b; Schulman et al. 2015]. In particular, Schulman et al. [2015] proposed a method for implementing the SPGE for stochastic computation graphs via graph transformation, and this method was adopted in the implementation of the same estimator in Pyro and also in our work. However, the method lacks a formal analysis on the implemented estimator especially in the context of probabilistic programs; it does not have a version of Theorem 4.5, which formally identifies requirements for the unbiasedness of the estimator. Also, the method does not check the required smoothness properties of given probabilistic programs. Our work fills in these gaps. Gorinova et al. [2020] proposed an automatic technique to transform models in a PPL using the same or closely-related transformation of sample commands in the SPGE. The work is, however, concerned with transforming models and taming their posterior distributions, while ours focuses on transforming guides. Also, the work does not check smoothness properties of transformed models that are required for running efficient inference algorithms, such as Hamiltonian Monte Carlo, on those models, while our work checks those properties using our program analysis.

Program analyses or type systems for PPLs have been developed to detect common errors [Lee et al. 2020; Lew et al. 2020; Wang et al. 2021], infer important probabilistic properties such as conditional independence [Gorinova et al. 2022], or automatically plan inference algorithms [WebPPL 2019] as in our work. In particular, WebPPL runs a simple program analysis (checking if there are interleaving sample and observe commands) to decide if it is worth applying sequential Monte Carlo.

The smoothness properties computed by our program analysis, such as differentiability and local Lipschitzness, fall in the scope of *hyperliveness* in the hierarchy of hyperproperties [Clarkson and Schneider 2008]. Intuitively, hyperliveness properties are those that cannot be refuted based on any finite counterexample (i.e., made of finitely-many finite execution traces), and counterexamples for differentiability and local Lipschitzness should indeed require infinitely-many execution traces due to the use of limit or all neighbouring inputs in their definitions. Not so many analyses have considered such hyperliveness properties. Among those, the most relevant to our work are the continuity analyses of Chaudhuri et al. [2010, 2012]. It uses a program abstraction that is rather similar to ours, but their analyses suffer from soundness issues, partly due to the incorrect joining of continuity sets [Chaudhuri et al. 2010] and also to an unsound rule for sequential composition [Chaudhuri et al. 2012] (see §A for details). We do not claim that these issues are difficult to fix. Our point is just that developing program analyses for smoothness properties requires special care. Chaudhuri et al.'s work focuses on proving smoothness properties of control software, or revealing the unexpected continuity of discrete algorithms. On the other hand, our program analysis is designed to assist variational inference and model learning for probabilistic programs. Barthe et al. [2020] proposed a refinement type system, which considers a higher-order functional language and ensures that every typable first-order program is continuous in all variables. On the other hand, our program analysis considers a first-order imperative language and can prove that a program is continuous in some (not necessarily all) variables. Other existing program analyses for smoothness properties include [Laurel et al. 2022] which over-approximates the Clarke generalised Jacobian, and [Mangal et al. 2020] which proves probabilistic Lipschitzness.

ACKNOWLEDGMENTS

We thank Hangeyol Yu for helping us prove Theorems 4.2 and 4.5, and anonymous reviewers for giving constructive comments. Lee was supported by Samsung Scholarship. Yang was supported by the Engineering Research Center Program through the National Research Foundation of Korea (NRF) funded by the Korean Government MSIT (NRF-2018R1A5A1059921) and also by the Institute for Basic Science (IBS-R029-C1). Rival was supported by the French ANR VeriAMOS project.

REFERENCES

- Martin Arjovsky, Soumith Chintala, and Léon Bottou. 2017. Wasserstein Generative Adversarial Networks. In *International Conference on Machine Learning (ICML)*. 214–223.
- Gilles Barthe, Raphaëlle Crubillé, Ugo Dal Lago, and Francesco Gavazzo. 2020. On the Versatility of Open Logical Relations - Continuity, Automatic Differentiation, and a Containment Theorem. In *European Symposium on Programming (ESOP)*. 56–83. https://doi.org/10.1007/978-3-030-44914-8_3
- Eli Bingham, Jonathan P. Chen, Martin Jankowiak, Fritz Obermeyer, Neeraj Pradhan, Theofanis Karaletsos, Rohit Singh, Paul A. Szerlip, Paul Horsfall, and Noah D. Goodman. 2019. Pyro: Deep Universal Probabilistic Programming. *Journal of Machine Learning Research* 20, 28 (2019), 1–6.
- B. Blanchet, P. Cousot, R. Cousot, J. Feret, L. Mauborgne, A. Miné, D. Monniaux, and X. Rival. 2003. A Static Analyzer for Large Safety Critical Software. In *Programming Languages, Design and Implementation (PLDI)*. 196–207. <https://doi.org/10.1145/781131.781153>
- Vladimir I. Bogachev. 2007. *Measure Theory* (first ed.). Springer. <https://doi.org/10.1007/978-3-540-34514-5>
- Bob Carpenter, Andrew Gelman, Matthew Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Marcus Brubaker, Jiqiang Guo, Peter Li, and Allen Riddell. 2017. Stan: A Probabilistic Programming Language. *Journal of Statistical Software* 76, 1 (2017), 1–32. <https://doi.org/10.18637/jss.v076.i01>
- Arun Tejasvi Chaganty, Aditya V. Nori, and Sriram K. Rajamani. 2013. Efficiently Sampling Probabilistic Programs via Program Analysis. In *Artificial Intelligence and Statistics (AISTATS)*. 153–160.
- Swarat Chaudhuri, Sumit Gulwani, and Roberto Lubliner. 2010. Continuity analysis of programs. In *Principles of Programming Languages (POPL)*. 57–70. <https://doi.org/10.1145/1706299.1706308>
- Swarat Chaudhuri, Sumit Gulwani, and Roberto Lubliner. 2012. Continuity and robustness of programs. *Commun. ACM* 55, 8 (2012), 107–115. <https://doi.org/10.1145/2240236.2240262>
- Guillaume Claret, Sriram K. Rajamani, Aditya V. Nori, Andrew D. Gordon, and Johannes Borgström. 2013. Bayesian inference using data flow analysis. In *Foundations of Software Engineering (FSE)*. 92–102. <https://doi.org/10.1145/2491411.2491423>
- M. R. Clarkson and F. B. Schneider. 2008. Hyperproperties. In *Computer Security Foundations (CSF)*. 51–65. <https://doi.org/10.1109/CSF.2008.7>
- Patrick Cousot and Radhia Cousot. 1977. Abstract Interpretation: A Unified Lattice Model for Static Analysis of Programs by Construction or Approximation of Fixpoints. In *Principles of Programming Languages (POPL)*. 238–252. <https://doi.org/10.1145/512950.512973>
- Patrick Cousot and Radhia Cousot. 1979. Systematic design of program analysis frameworks. In *Principles of Programming Languages (POPL)*. 269–282. <https://doi.org/10.1145/567752.567778>
- S. M. Ali Eslami, Nicolas Heess, Theophane Weber, Yuval Tassa, David Szepesvari, Koray Kavukcuoglu, and Geoffrey E. Hinton. 2016. Attend, Infer, Repeat: Fast Scene Understanding with Generative Models. In *Neural Information Processing Systems (NIPS)*. 3233–3241.
- Hong Ge, Kai Xu, and Zoubin Ghahramani. 2018. Turing: A Language for Flexible Probabilistic Inference. In *Artificial Intelligence and Statistics (AISTATS)*. 1682–1690.
- Timon Gehr, Sasa Misailovic, and Martin T. Vechev. 2016. PSI: Exact Symbolic Inference for Probabilistic Programs. In *Computer Aided Verification (CAV)*. 62–83. https://doi.org/10.1007/978-3-319-41528-4_4
- Noah Goodman, Vikash Mansinghka, Daniel M Roy, Keith Bonawitz, and Joshua B Tenenbaum. 2008. Church: a language for generative models. In *Uncertainty in Artificial Intelligence (UAI)*. 220–229.
- Andrew D. Gordon, Thore Graepel, Nicolas Rolland, Claudio Russo, Johannes Borgstrom, and John Guiver. 2014. Tabular: A Schema-driven Probabilistic Programming Language. In *Principles of Programming Languages (POPL)*. 321–334. <https://doi.org/10.1145/2578855.2535850>
- Maria I. Gorinova, Andrew D. Gordon, Charles Sutton, and Matthijs Vákár. 2022. Conditional Independence by Typing. *ACM Trans. Program. Lang. Syst.* 44, 1 (2022), 4:1–4:54. <https://doi.org/10.1145/3490421>
- Maria I. Gorinova, Dave Moore, and Matthew D. Hoffman. 2020. Automatic Reparameterisation of Probabilistic Programs. In *International Conference on Machine Learning (ICML)*. 3648–3657.
- Steven Holtzen, Guy Van den Broeck, and Todd D. Millstein. 2020. Scaling exact inference for discrete probabilistic programs. *Proc. ACM Program. Lang.* 4, OOPSLA (2020), 140:1–140:31. <https://doi.org/10.1145/3428208>
- Hyunjik Kim, George Papamakarios, and Andriy Mnih. 2021. The Lipschitz Constant of Self-Attention. In *International Conference on Machine Learning (ICML)*. 5562–5571.
- Diederik P. Kingma and Max Welling. 2014. Auto-Encoding Variational Bayes. In *International Conference on Learning Representations (ICLR)*.
- Alp Kucukelbir, Rajesh Ranganath, Andrew Gelman, and David M. Blei. 2015. Automatic Variational Inference in Stan. In *Neural Information Processing Systems (NIPS)*. 568–576.
- Jacob Laurel, Rem Yang, Gagandeep Singh, and Sasa Misailovic. 2022. A dual number abstraction for static analysis of Clarke Jacobians. *Proc. ACM Program. Lang.* 6, POPL (2022), 1–30. <https://doi.org/10.1145/3498718>

- John M. Lee. 2012. *Introduction to Smooth Manifolds* (second ed.). Springer. <https://doi.org/10.1007/978-1-4419-9982-5>
- Wonyeol Lee, Xavier Rival, and Hongseok Yang. 2022a. Artifact for the Paper “Smoothness Analysis for Probabilistic Programs with Application to Optimised Variational Inference”. <https://doi.org/10.5281/zenodo.7246597>
- Wonyeol Lee, Xavier Rival, and Hongseok Yang. 2022b. Smoothness Analysis for Probabilistic Programs with Application to Optimised Variational Inference. *arXiv:2208.10530* (2022). <https://doi.org/10.48550/ARXIV.2208.10530>
- Wonyeol Lee, Hangeyi Yu, Xavier Rival, and Hongseok Yang. 2020. Towards verified stochastic variational inference for probabilistic programs. *Proc. ACM Program. Lang.* 4, POPL (2020), 16:1–16:33. <https://doi.org/10.1145/3371084>
- Alexander K. Lew, Marco F. Cusumano-Towner, Benjamin Sherman, Michael Carbin, and Vikash K. Mansinghka. 2020. Trace types and denotational semantics for sound programmable inference in probabilistic languages. *Proc. ACM Program. Lang.* 4, POPL (2020), 19:1–19:32. <https://doi.org/10.1145/3371087>
- Ravi Mangal, Kartik Sarangmath, Aditya V. Nori, and Alessandro Orso. 2020. Probabilistic Lipschitz Analysis of Neural Networks. In *Static Analysis Symposium (SAS)*. 274–309. https://doi.org/10.1007/978-3-030-65474-0_13
- Vikash K. Mansinghka, Daniel Selsam, and Yura N. Perov. 2014. Venture: a higher-order probabilistic programming platform with programmable inference. *arXiv:1404.0099* (2014). <https://doi.org/10.48550/ARXIV.1404.0099>
- T. Minka, J.M. Winn, J.P. Guiver, S. Webster, Y. Zaykov, B. Yangel, A. Spengler, and J. Bronskill. 2014. Infer.NET 2.6. <https://dotnet.github.io/infer/>.
- Praveen Narayanan, Jacques Carette, Wren Romano, Chung-chieh Shan, and Robert Zinkov. 2016. Probabilistic inference by program transformation in Hakaru (system description). In *Functional and Logic Programming (FLOPS)*. 62–79. https://doi.org/10.1007/978-3-319-29604-3_5
- Radford M. Neal. 2011. MCMC Using Hamiltonian Dynamics. In *Handbook of Markov Chain Monte Carlo*. 113–162. <https://doi.org/10.1201/b10905>
- Aditya V. Nori, Chung-Kil Hur, Sriram K. Rajamani, and Selva Samuel. 2014. R2: An Efficient MCMC Sampler for Probabilistic Programs. In *AAAI Conference on Artificial Intelligence (AAAI)*. 2476–2482. <https://doi.org/10.1609/aaai.v28i1.9060>
- André Platzer. 2018. *Logical Foundations of Cyber-Physical Systems*. Springer. <https://doi.org/10.1007/978-3-319-63588-0>
- Rajesh Ranganath, Sean Gerrish, and David M. Blei. 2014. Black Box Variational Inference. In *Artificial Intelligence and Statistics (AISTATS)*. 814–822.
- Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. 2014. Stochastic Backpropagation and Approximate Inference in Deep Generative Models. In *International Conference on Machine Learning (ICML)*. 7344–7353.
- Daniel Ritchie, Paul Horsfall, and Noah D. Goodman. 2016a. Deep Amortized Inference for Probabilistic Programs. *arXiv:1610.05735* (2016). <https://doi.org/10.48550/ARXIV.1610.05735>
- Daniel Ritchie, Andreas Stuhlmüller, and Noah D. Goodman. 2016b. C3: Lightweight Incrementalized MCMC for Probabilistic Programs using Continuations and Callsite Caching. In *Artificial Intelligence and Statistics (AISTATS)*. 28–37.
- John Salvatier, Thomas V. Wiecki, and Christopher Fonnesbeck. 2016. Probabilistic programming in Python using PyMC3. *PeerJ Comput. Sci.* 2 (2016), e55. <https://doi.org/10.7717/peerj-cs.55>
- John Schulman, Nicolas Heess, Theophane Weber, and Pieter Abbeel. 2015. Gradient Estimation Using Stochastic Computation Graphs. In *Neural Information Processing Systems (NIPS)*. 3528–3536.
- N. Siddharth, Brooks Paige, Jan-Willem van de Meent, Alban Desmaison, Noah D. Goodman, Pushmeet Kohli, Frank Wood, and Philip Torr. 2017. Learning Disentangled Representations with Semi-Supervised Deep Generative Models. In *Neural Information Processing Systems (NIPS)*. 5927–5937.
- Sam Staton, Hongseok Yang, Frank D. Wood, Chris Heunen, and Ohad Kammar. 2016. Semantics for probabilistic programming: higher-order functions, continuous distributions, and soft constraints. In *Logic in Computer Science (LICS)*. 525–534. <https://doi.org/10.1145/2933575.2935313>
- David Tolpin, Jan-Willem van de Meent, Hongseok Yang, and Frank D. Wood. 2016. Design and Implementation of Probabilistic Programming Language Anglican. In *Implementation and Application of Functional Programming Languages (IFL)*. 6:1–6:12. <https://doi.org/10.1145/3064899.3064910>
- Dustin Tran, Matthew D. Hoffman, Dave Moore, Christopher Suter, Srinivas Vasudevan, and Alexey Radul. 2018. Simple, Distributed, and Accelerated Probabilistic Programming. In *Neural Information Processing Systems (NeurIPS)*. 7609–7620.
- Dustin Tran, Alp Kucukelbir, Adji B. Dieng, Maja R. Rudolph, Dawen Liang, and David M. Blei. 2016. Edward: A library for probabilistic modeling, inference, and criticism. *arXiv:1610.09787* (2016). <https://doi.org/10.48550/ARXIV.1610.09787>
- Uber AI Labs. 2022. Pyro examples. <http://pyro.ai/examples/>. Version used: June 18, 2022.
- Jan-Willem van de Meent, Brooks Paige, Hongseok Yang, and Frank Wood. 2018. An Introduction to Probabilistic Programming. *arXiv:1809.10756* (2018). <https://doi.org/10.48550/ARXIV.1809.10756>
- Di Wang, Jan Hoffmann, and Thomas W. Reps. 2021. Sound probabilistic inference via guide types. In *Programming Language Design and Implementation (PLDI)*. 788–803. <https://doi.org/10.1145/3453483.3454077>
- WebPPL. 2019. <https://github.com/probmods/webppl/blob/v0.9.15/src/header.wppl#L510>.
- Ronald J. Williams. 1992. Simple Statistical Gradient-Following Algorithms for Connectionist Reinforcement Learning. *Machine Learning* 8, 3-4 (1992), 229–256. <https://doi.org/10.1007/BF00992696>

- David Wingate, Noah D. Goodman, Andreas Stuhlmüller, and Jeffrey Mark Siskind. 2011a. Nonstandard Interpretations of Probabilistic Programs for Efficient Inference. In *Neural Information Processing Systems (NIPS)*. 1152–1160.
- David Wingate, Andreas Stuhlmüller, and Noah D. Goodman. 2011b. Lightweight Implementations of Probabilistic Programming Languages Via Transformational Compilation. In *Artificial Intelligence and Statistics (AISTATS)*. 770–778.
- David Wingate and Theophane Weber. 2013. Automated Variational Inference in Probabilistic Programming. *arXiv:1301.1299* (2013). <https://doi.org/10.48550/ARXIV.1301.1299>
- Frank Wood, Jan Willem van de Meent, and Vikash Mansinghka. 2014. A New Approach to Probabilistic Programming Inference. In *Artificial Intelligence and Statistics (AISTATS)*. 1024–1032.
- Chenling Xu, Romain Lopez, Edouard Mehlman, Jeffrey Regier, Michael I Jordan, and Nir Yosef. 2021. Probabilistic harmonization and annotation of single-cell transcriptomics data with deep generative models. *Molecular systems biology* 17, 1 (2021), e9620. <https://doi.org/10.15252/msb.20209620>
- Yuan Zhou, Hongseok Yang, Yee Whye Teh, and Tom Rainforth. 2020. Divide, Conquer, and Combine: a New Inference Strategy for Probabilistic Programs with Stochastic Support. In *International Conference on Machine Learning (ICML)*. 11534–11545.

Received 2022-07-07; accepted 2022-11-07

A DEFERRED RESULTS IN §1

A.1 Unsoundness of Continuity Analyses in [Chaudhuri et al. 2010, 2012]

The continuity analysis in [Chaudhuri et al. 2010] considers joint continuity, whereas the continuity analysis in [Chaudhuri et al. 2012] considers partial continuity. That is, given a command c , an output variable v of c , and some input variables u_1, \dots, u_m to c , the former analyses whether v is continuous in $\{u_1, \dots, u_m\}$ jointly, whereas the latter analyses whether v is continuous in u_i separately for every $1 \leq i \leq m$.

Join and Sequence rules. The former analysis contains a rule called Join [Chaudhuri et al. 2010, Figure 3], and the latter analysis contains a rule called Sequence [Chaudhuri et al. 2012, Figure 1]. The two rules can be rewritten (with some simplifications) as follows, in terms of functions between \mathbb{R}^n : for any $f, g : \mathbb{R}^n \rightarrow \mathbb{R}^n$ and $S, S', T, U \subseteq \{1, \dots, n\}$,

$$\frac{\begin{array}{l} \text{For each } j \in T, f_j \text{ is continuous in } \{x_i \mid i \in S\} \\ \text{For each } j \in T, f_j \text{ is continuous in } \{x_i \mid i \in S'\} \end{array}}{\text{For each } j \in T, f_j \text{ is continuous in } \{x_i \mid i \in S \cup S'\}} \text{ (Join)}$$

$$\frac{\begin{array}{l} \text{For each } j \in T, f_j \text{ is continuous in } x_i \text{ for each } i \in S \\ \text{For each } k \in U, g_k \text{ is continuous in } y_j \text{ for each } j \in T \end{array}}{\text{For each } k \in U, (g \circ f)_k \text{ is continuous in } x_i \text{ for each } i \in S} \text{ (Sequence)}$$

where f and g are functions of variables x_1, \dots, x_n and y_1, \dots, y_n , respectively, and $h_i \triangleq \text{proj}_i \circ h$ for $h : \mathbb{R}^n \rightarrow \mathbb{R}^n$ and $i \in \{1, \dots, n\}$ denotes the i -th component of h . As mentioned above, the Join rule analyses joint continuity, while the Sequence rule analyses partial continuity. Further, the Join rule says that joint continuity is preserved under the union of input variables, while the Sequence rule says that partial continuity is preserved under the composition of functions.

The two rules, however, are unsound with the following counterexamples. Let $h : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be the function

$$h(x_1, x_2) \triangleq \begin{cases} (x_1 x_2 / (x_1^2 + x_2^2), x_2) & \text{if } (x_1, x_2) \neq (0, 0) \\ (0, x_2) & \text{otherwise.} \end{cases}$$

Note that h_1 is continuous in x_1 and in x_2 separately, but not in $\{x_1, x_2\}$ jointly. First, for the Join rule, consider the following $f : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ and $S, S', T \subseteq \{1, 2\}$:

$$f(x_1, x_2) \triangleq h(x_1, x_2), \quad S \triangleq \{1\}, \quad S' \triangleq \{2\}, \quad T \triangleq \{1, 2\}.$$

Then, the premise of the Join rule holds, and so the conclusion of the rule must hold. But this is *not* the case since $f_1 = h_1$ is not continuous in $\{x_1, x_2\}$. Hence, the Join rule is unsound. Next, for the Sequence rule, consider the following $f, g : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ and $S, T, U \subseteq \{1, 2\}$:

$$f(x_1, x_2) \triangleq (x_1, x_1), \quad g(x_1, x_2) \triangleq h(x_1, x_2), \quad S \triangleq T \triangleq U \triangleq \{1, 2\}.$$

Then, the premise of the Sequence rule holds, and so the conclusion of the rule must hold. But this is *not* the case since $(g \circ f)_1$ is not continuous in x_1 (due to $(g \circ f)_1(x_1, x_2) = \mathbf{1}_{\{x_1 \neq 0\}} \cdot \frac{1}{2}$). Hence, the Sequence rule is unsound. These counterexamples show that joint continuity is *not* preserved under the union of input variables, and partial continuity is *not* preserved under the composition of functions.

The two aforementioned counterexamples can be easily translated into programs: the first becomes $c_1 \equiv (z := h_1(x, y))$ and the second becomes $c_2 \equiv (y := x; z := h_1(x, y))$, where x, y , and z are program variables and h_1 is the binary operator defined above. The analysis in [Chaudhuri et al. 2010] deduces that in c_1 , z is continuous in x and y (jointly), and the analysis in [Chaudhuri et al. 2012] deduces that in c_2 , z is continuous in x (separately). Both deductions, however, are incorrect as seen above, and the two analyses are thus unsound.

Sequence rule (again). Both of the two analyses in [Chaudhuri et al. 2010, 2012] contain the Sequence rule (discussed above) which has the following rule as an instance: for all variables v_0, v_1, v_2 that are not necessarily distinct, and for all commands c_1, c_2 ,

$$\frac{\begin{array}{l} \text{In } c_1, v_1 \text{ is continuous in } v_0 \\ \text{In } c_2, v_2 \text{ is continuous in } v_1 \end{array}}{\text{In } (c_1; c_2), v_2 \text{ is continuous in } v_0}$$

The above instance of the Sequence rule is, however, unsound because it incorrectly handles the dependencies between variables. For instance, consider commands $c_1 \equiv (\text{if } (x < 0) \{y := 0\} \text{ else } \{y := 1\})$ and $c_2 \equiv (z := x + y)$, and variables $v_0 \equiv x, v_1 \equiv x$, and $v_2 \equiv z$. Then, in $(c_1; c_2)$, z is *not* continuous in x (because z after $(c_1; c_2)$ is x if $x < 0$, and is $x + 1$ if $x \geq 0$). However, the premise of the above rule holds for this case, and so the rule incorrectly concludes that z is continuous in x under $(c_1; c_2)$, by ignoring the dependency of z on y (which is discontinuous in x).

Loop rule. The analysis in [Chaudhuri et al. 2010] contains a rule called Simple-loop (Figure 5 in the paper) to analyse loops, and the analysis in [Chaudhuri et al. 2012] contains a rule called Loop (Figure 1 in the paper) which is essentially the same as the Simple-loop rule.

The two rules, however, incorrectly assume that (joint or partial) continuity without any restriction on the domain of functions satisfies the admissibility assumption discussed in Remark 5.9; and this in turn makes the rules unsound. To illustrate the unsoundness, consider a command $c \equiv (\text{while } (0 < x < 1) \{c'\})$ with $c' \equiv (y := y + f(x); x := g(x))$, where f and g are the continuous operators defined by:

$$f(x) \triangleq \begin{cases} x & \text{if } 0 < x \leq 1/2 \\ 1 - x & \text{if } 1/2 < x \leq 1 \\ 0 & \text{otherwise,} \end{cases} \quad g(x) \triangleq \begin{cases} 0 & \text{if } x \leq 0 \\ 2x & \text{if } 0 < x \leq 1/2 \\ 1 & \text{otherwise.} \end{cases}$$

Then, the premises of the two rules are satisfied, mainly because x and y after c' are continuous jointly in x and y before c' , and x and y do not change in c' if $x = 0$ or $x = 1$ (i.e., at the “boundary” of the loop condition of c). Hence, the two rules conclude that x and y after c are continuous in x and y before c (jointly or separately). However, this is an unsound conclusion because in c , y is *not* continuous in x at $x = 0$: we have $y' = y$ if $x = 0$, but $y' \rightarrow y + 1$ as $x \rightarrow 0^+$ (more precisely, $y' = y$ if $x \leq 0$ or $x \geq 1$, and $y' = y + (1 - x)$ if $0 < x < 1$), where x' and y' denote the values of x and y after c .

B DEFERRED RESULTS IN §2

B.1 Table Summarising §2

Table 6. Gradient estimators for variational inference, and requirements for each estimator. “Req.” denotes “Requirement” and “diff.” denotes “differentiable”. Recall that $f_\theta(z) \triangleq \log(p_{c_m}(z)/p_{c_g, \theta}(z))$.

	SCE	PGE	SPGE
Setup	$q_\theta(z)$	$p_{c_g, \theta}(z)$	$p_{c_g, \theta}(z)$
	$v_\theta(z)$	z	$v_{c_g, \theta}(z)$
	$g_\theta(z)$	$f_\theta(v_\theta(z)) \cdot \nabla_\theta \log q_\theta(z)$	$\nabla_\theta (f_\theta(v_\theta(z))) + f_\theta(v_\theta(z)) \cdot \nabla_\theta \log q_\theta(z)$
Req.	$q_\theta(z)$	diff. in θ	diff. in θ
	$v_\theta(z)$	–	diff. in θ
	$p_{c_m}(z)$	diff. in θ and z	diff. in θ and “selected” z_i 's
	$p_{c_g, \theta}(z)$	diff. in θ and z	diff. in θ and “selected” z_i 's

Table 6 compares key aspects of the three gradient estimators (SCE, PGE, and SPGE) explained in §2.

C DEFERRED RESULTS IN §4.1

C.1 Proof of Theorem 4.2

We introduce several definitions, state lemmas, and prove Theorem 4.2 using the lemmas. We prove the lemmas in §C.2 and §C.3.

Recall the partition $\text{Var} = \text{PVar} \uplus \text{Name} \uplus \text{AVar}$ of Var . We use the following letters to denote the values of each part: $\sigma_p \in \text{St}[\text{PVar}]$, $\sigma_n \in \text{St}[\text{Name}]$, and $\sigma_a \in \text{St}[\text{AVar}]$. Based on the partition, we define the next functions:

$$\begin{aligned} \text{prs}(c) &: \text{St}[\text{PVar}] \times \text{St}[\text{Name}] \times \text{St}[\text{AVar}] \rightarrow [0, \infty), \\ \text{prs}(c)(\sigma_p, \sigma_n, \sigma_a) &\triangleq \begin{cases} \llbracket c \rrbracket(\sigma_p \oplus \sigma_n \oplus \sigma_a)(\text{like}) & \text{if } \text{noerr}(c, \sigma_p \oplus \sigma_n \oplus \sigma_a) \\ \cdot \prod_{\mu \in \text{Name}} \llbracket c \rrbracket(\sigma_p \oplus \sigma_n \oplus \sigma_a)(pr_\mu) & \\ 0 & \text{otherwise,} \end{cases} \\ \text{vals}(c) &: \text{St}[\text{PVar}] \times \text{St}[\text{Name}] \times \text{St}[\text{AVar}] \rightarrow \text{St}[\text{Name}], \\ \text{vals}(c)(\sigma_p, \sigma_n, \sigma_a) &\triangleq \begin{cases} \lambda \mu \in \text{Name}. \llbracket c \rrbracket(\sigma_p \oplus \sigma_n \oplus \sigma_a)(val_\mu) & \text{if } \text{noerr}(c, \sigma_p \oplus \sigma_n \oplus \sigma_a) \\ \lambda \mu \in \text{Name}. 0 & \text{otherwise,} \end{cases} \end{aligned}$$

where $\text{noerr}(c, \sigma)$ is a predicate for a command c and $\sigma \in \text{St}$, defined by

$$\text{noerr}(c, \sigma) \iff \llbracket c \rrbracket \sigma \in \text{St} \wedge (\forall \mu \in \text{Name}. \llbracket c \rrbracket \sigma(\text{cnt}_\mu) - \sigma(\text{cnt}_\mu) \leq 1).$$

The predicate $\text{noerr}(c, \sigma)$ says that c terminates for σ without a double-sampling error. The functions prs and vals generalise the density function p and the value function v , respectively; in particular, they do not assume a particular initial state σ_0 used in Eq. (3). We consider the generalisation of p and v so as to enable inductive proofs.

Although generalising p and v , the functions prs and vals are not sufficient to enable inductive proofs since their inputs and outputs contain some unnecessary parts, which stops induction from working well (especially in the sequential composition case): namely, the part of $\text{St}[\text{Name}]$ that is not read during execution, and the part of $\text{St}[\text{AVar}]$ that is not updated during execution. To exclude those unnecessary parts, we first define the set of substates of $\text{St}[\text{Name}]$ as follows:

$$\xi_n \in \text{St}_\square[\text{Name}] \triangleq \bigcup_{K \subseteq \text{Name}} \text{St}[K].$$

Based on these substates, we define the next functions:

$$\begin{aligned} \text{prs}_\square(c) &: \text{St}[\text{PVar}] \times \text{St}_\square[\text{Name}] \rightarrow [0, \infty), \\ \text{prs}_\square(c)(\sigma_p, \xi_n) &\triangleq \begin{cases} \llbracket c \rrbracket(\sigma_p \oplus \xi_n \oplus \sigma_r)(\text{like}) & \text{if } \exists \sigma_r. \text{used}(c, \sigma_p \oplus \xi_n \oplus \sigma_r, \xi_n) \\ \cdot \prod_{\mu \in \text{dom}(\xi_n)} \llbracket c \rrbracket(\sigma_p \oplus \xi_n \oplus \sigma_r)(pr_\mu) & \\ 0 & \text{otherwise,} \end{cases} \\ \text{vals}_\square(c) &: \text{St}[\text{PVar}] \times \text{St}_\square[\text{Name}] \rightarrow \text{St}_\square[\text{Name}], \\ \text{vals}_\square(c)(\sigma_p, \xi_n) &\triangleq \begin{cases} \lambda \mu \in \text{dom}(\xi_n). \llbracket c \rrbracket(\sigma_p \oplus \xi_n \oplus \sigma_r)(val_\mu) & \text{if } \exists \sigma_r. \text{used}(c, \sigma_p \oplus \xi_n \oplus \sigma_r, \xi_n) \\ \lambda \mu \in \text{dom}(\xi_n). 0 & \text{otherwise,} \end{cases} \\ \text{pvars}_\square(c) &: \text{St}[\text{PVar}] \times \text{St}_\square[\text{Name}] \rightarrow \text{St}[\text{PVar}], \\ \text{pvars}_\square(c)(\sigma_p, \xi_n) &\triangleq \begin{cases} \lambda x \in \text{PVar}. \llbracket c \rrbracket(\sigma_p \oplus \xi_n \oplus \sigma_r)(x) & \text{if } \exists \sigma_r. \text{used}(c, \sigma_p \oplus \xi_n \oplus \sigma_r, \xi_n) \\ \lambda x \in \text{PVar}. 0 & \text{otherwise,} \end{cases} \end{aligned}$$

where $used(c, \sigma, \xi_n)$ is a predicate for a command c , $\sigma \in \text{St}$, and $\xi_n \in \text{St}_\square[\text{Name}]$, defined by

$$used(c, \sigma, \xi_n) \iff noerr(c, \sigma) \wedge (\sigma(like) = 1) \wedge (\xi_n = \sigma|_{\text{dom}(\xi_n)}) \\ \wedge (\text{dom}(\xi_n) = \{\mu \in \text{Name} \mid \llbracket c \rrbracket \sigma(cnt_\mu) - \sigma(cnt_\mu) = 1\}).$$

The predicate $used(c, \sigma, \xi_n)$ says that c terminates for σ without a double-sampling error, $like$ is initialised to 1 in σ , and ξ_n is the Name part of σ that is sampled during the execution of c from σ . By using $used(-, -, -)$, the three functions do not take the unnecessary part of a state as an input, and do not return the unnecessary part of a state in the output. The three functions are well-defined.

LEMMA C.1. prs_\square , $vals_\square$, and $pvars_\square$ are well-defined, i.e., they do not depend on the choice of σ_r .

We now state two main lemmas for Theorem 4.2. The first lemma describes how prs and $vals$ are connected with prs_\square and $vals_\square$. The second lemma says that a particular integral involving prs_\square , $vals_\square$, and $pvars_\square$ is the same for c and \bar{c}^π if a reparameterisation plan π is valid.

LEMMA C.2. Let c be a command, and $f_i : \mathbb{R} \rightarrow \mathbb{R}$ for $i \in \{1, 2, 3\}$ be measurable functions such that $f_1(r) \geq 0$ for all $r \in \mathbb{R}$. Define $f_* : \text{St}[\text{Name}] \rightarrow \text{St}[\text{AVar}]$ by

$$f_*(\sigma_n)(a) \triangleq \begin{cases} 1 & \text{if } a \equiv like \\ f_1(\sigma_n(\mu)) & \text{if } a \equiv pr_\mu \text{ for } \mu \in \text{Name} \\ f_2(\sigma_n(\mu)) & \text{if } a \equiv val_\mu \text{ for } \mu \in \text{Name} \\ f_3(\sigma_n(\mu)) & \text{if } a \equiv cnt_\mu \text{ for } \mu \in \text{Name}. \end{cases}$$

Then, for all $\sigma_p \in \text{St}[\text{PVar}]$ and all measurable $h : \text{St}[\text{Name}] \rightarrow \mathbb{R}$,

$$\int d\sigma_n \left(prs(c)(\sigma_p, \sigma_n, f_*(\sigma_n)) \cdot h \left(vals(c)(\sigma_p, \sigma_n, f_*(\sigma_n)) \right) \right) \\ = \int d\xi_n \left(prs_\square(c)(\sigma_p, \xi_n) \cdot g \left(vals_\square(c)(\sigma_p, \xi_n) \right) \right)$$

where the integral on the LHS is defined if and only if the one on the RHS is defined, and the function $g : \text{St}_\square[\text{Name}] \rightarrow \mathbb{R}$ is defined by

$$g(\xi'_n) = \int d\xi'_n \left(\mathbf{1}_{[\text{dom}(\xi'_n) \uplus \text{dom}(\xi'_n) = \text{Name}]} \cdot \left(\prod_{\mu \in \text{dom}(\xi'_n)} f_1(\xi'_n(\mu)) \right) \cdot h \left(\xi'_n \oplus \lambda \mu \in \text{dom}(\xi'_n) \cdot f_2(\xi'_n(\mu)) \right) \right). \quad (14)$$

LEMMA C.3. Let c be a command and $g : \text{St}[\text{PVar}] \times \text{St}_\square[\text{Name}] \rightarrow \mathbb{R}$ be a measurable function. Then, for all $\sigma_p \in \text{St}[\text{PVar}]$,

$$\int d\xi_n \left(prs_\square(c)(\sigma_p, \xi_n) \cdot g \left(pvars_\square(c)(\sigma_p, \xi_n), vals_\square(c)(\sigma_p, \xi_n) \right) \right) \\ = \int d\xi_n \left(prs_\square(\bar{c}^\pi)(\sigma_p, \xi_n) \cdot g \left(pvars_\square(\bar{c}^\pi)(\sigma_p, \xi_n), vals_\square(\bar{c}^\pi)(\sigma_p, \xi_n) \right) \right)$$

where the integral on the LHS is defined if and only if the one on the RHS is defined.

We now prove Theorem 4.2 using these two lemmas.

PROOF OF THEOREM 4.2. Let π be a valid reparameterisation plan, c be a command, $\sigma_\theta \in \text{St}[\theta]$, and $h : \text{St}[\text{Name}] \rightarrow \mathbb{R}$ be a measurable function. Suppose that the integral on the LHS of Theorem 4.2 is defined. Recall that for a given $\sigma_n \in \text{St}[\text{Name}]$, the definitions of p and v (in §3 and §4.1) use the

initial state $\sigma \triangleq \sigma_\theta \oplus \sigma_n \oplus \sigma_0 \in \text{St}$, where $\sigma_0 \in \text{St}[(\text{PVar} \setminus \theta) \cup \text{AVar}]$ depends on σ_n and has the following definition:

$$\sigma_0(v) \triangleq \begin{cases} 0 & \text{if } v \in \text{PVar} \setminus \theta \\ 1 & \text{if } v \equiv \text{like} \\ \mathcal{N}(\sigma_n(\mu); 0, 1) & \text{if } v \equiv \text{pr}_\mu \text{ for } \mu \in \text{Name} \\ \sigma_n(\mu) & \text{if } v \equiv \text{val}_\mu \text{ for } \mu \in \text{Name} \\ 0 & \text{if } v \equiv \text{cnt}_\mu \text{ for } \mu \in \text{Name}. \end{cases}$$

The initial state σ can be re-expressed as

$$\sigma = \sigma_p \oplus \sigma_n \oplus f_*(\sigma_n) \quad (15)$$

using the following $\sigma_p \in \text{St}[\text{PVar}]$ and $f_* : \text{St}[\text{Name}] \rightarrow \text{St}[\text{AVar}]$:

$$\sigma_p(x) \triangleq \begin{cases} \sigma_\theta(x) & \text{if } x \in \theta \\ 0 & \text{if } x \in \text{PVar} \setminus \theta, \end{cases} \quad f_*(\sigma_n)(a) \triangleq \begin{cases} 1 & \text{if } a \equiv \text{like} \\ f_1(\sigma_n(\mu)) & \text{if } a \equiv \text{pr}_\mu \text{ for } \mu \in \text{Name} \\ f_2(\sigma_n(\mu)) & \text{if } a \equiv \text{val}_\mu \text{ for } \mu \in \text{Name} \\ f_3(\sigma_n(\mu)) & \text{if } a \equiv \text{cnt}_\mu \text{ for } \mu \in \text{Name}, \end{cases}$$

where $f_1(r) \triangleq \mathcal{N}(r; 0, 1)$, $f_2(r) \triangleq r$, and $f_3(r) \triangleq 0$. Using this, we get the desired equation:

$$\begin{aligned} & \int d\sigma_n \left(p_{c, \sigma_\theta}(\sigma_n) \cdot h \left(v_{c, \sigma_\theta}(\sigma_n) \right) \right) \\ &= \int d\sigma_n \left(\text{prs}(c)(\sigma_p, \sigma_n, f_*(\sigma_n)) \cdot h \left(\text{vals}(c)(\sigma_p, \sigma_n, f_*(\sigma_n)) \right) \right) \\ &= \int d\xi_n \left(\text{prs}_\square(c)(\sigma_p, \xi_n) \cdot g \left(\text{vals}_\square(c)(\sigma_p, \xi_n) \right) \right) \\ &= \int d\xi_n \left(\text{prs}_\square(\bar{c}^\pi)(\sigma_p, \xi_n) \cdot g \left(\text{vals}_\square(\bar{c}^\pi)(\sigma_p, \xi_n) \right) \right) \\ &= \int d\sigma_n \left(\text{prs}(\bar{c}^\pi)(\sigma_p, \sigma_n, f_*(\sigma_n)) \cdot h \left(\text{vals}(\bar{c}^\pi)(\sigma_p, \sigma_n, f_*(\sigma_n)) \right) \right) \\ &= \int d\sigma_n \left(p_{\bar{c}^\pi, \sigma_\theta}(\sigma_n) \cdot h \left(v_{\bar{c}^\pi, \sigma_\theta}(\sigma_n) \right) \right) \end{aligned}$$

where $g : \text{St}_\square[\text{Name}] \rightarrow \mathbb{R}$ is defined as Eq. (14). The first equality holds by Eq. (15) and the definition of p_{c, σ_θ} , v_{c, σ_θ} , prs , and vals . The second equality holds by Lemma C.2 (applied to c). The third equality follows from Lemma C.3. The fourth equality holds by Lemma C.2 (applied to \bar{c}^π). The fifth equality holds by Eq. (15) and the definitions of $p_{\bar{c}^\pi, \sigma_\theta}$, $v_{\bar{c}^\pi, \sigma_\theta}$, prs , and vals . Note that the same equational reasoning with the reverse direction can be used to prove the claimed equation of the theorem when the integral on the RHS of the equation is defined. \square

C.2 Proofs of Lemmas C.1 and C.2

PROOF OF LEMMA C.1. Let c be a command, $\sigma_p \in \text{St}[\text{PVar}]$, and $\xi_n \in \text{St}_\square[\text{Name}]$. Consider $\sigma_r \in \text{St}[\text{Var} \setminus (\text{dom}(\sigma_p) \cup \text{dom}(\xi_n))]$ such that $\text{used}(c, \sigma_p \oplus \xi_n \oplus \sigma_r, \xi_n)$. We want to show that $\text{prs}_\square(c)(\sigma_p, \xi_n)$, $\text{vals}_\square(c)(\sigma_p, \xi_n)$, and $\text{pvars}_\square(c)(\sigma_p, \xi_n)$ do not depend on the choice of σ_r . To do so, consider $\sigma'_r \in \text{St}[\text{Var} \setminus (\text{dom}(\sigma_p) \cup \text{dom}(\xi_n))]$ such that $\text{used}(c, \sigma_p \oplus \xi_n \oplus \sigma'_r, \xi_n)$. Let $\sigma \triangleq \sigma_p \oplus \xi_n \oplus \sigma_r$

and $\sigma' \triangleq \sigma_p \oplus \xi_n \oplus \sigma'_r$. Then, it suffices to show that

$$\begin{aligned} \llbracket c \rrbracket \sigma(\text{like}) \cdot \prod_{\mu \in \text{dom}(\xi_n)} \llbracket c \rrbracket \sigma(\text{pr}_\mu) &= \llbracket c \rrbracket \sigma'(\text{like}) \cdot \prod_{\mu \in \text{dom}(\xi_n)} \llbracket c \rrbracket \sigma'(\text{pr}_\mu), \\ \llbracket c \rrbracket \sigma(\text{val}_\mu) &= \llbracket c \rrbracket \sigma'(\text{val}_\mu) \quad \text{for all } \mu \in \text{dom}(\xi_n), \\ \llbracket c \rrbracket \sigma(x) &= \llbracket c \rrbracket \sigma'(x) \quad \text{for all } x \in \text{PVar}. \end{aligned} \quad (16)$$

Since $\text{used}(c, \sigma, \xi_n)$ and $\text{used}(c, \sigma', \xi_n)$, we have $\sigma(\text{like}) = 1 = \sigma'(\text{like})$ and so $\sigma|_V = \sigma'|_V$ for $V = \text{PVar} \cup \text{dom}(\xi_n) \cup \{\text{like}\}$. Using this and $\text{used}(c, \sigma, \xi_n)$, we can apply Lemma C.6-(3) and (4) to get $\llbracket c \rrbracket \sigma(v) = \llbracket c \rrbracket \sigma'(v)$ for all $v \in \text{PVar} \cup \{\text{like}\} \cup \{\text{pr}_\mu, \text{val}_\mu \mid \mu \in \text{dom}(\xi_n)\}$. Hence, we obtain the desired equations in Eq. (16). \square

PROOF OF LEMMA C.2. Let c be a command, $h : \text{St}[\text{Name}] \rightarrow \mathbb{R}$ be a measurable function, $f_* : \text{St}[\text{Name}] \rightarrow \text{St}[\text{AVar}]$ be the function defined in the statement of this lemma, and $\sigma_p \in \text{St}[\text{PVar}]$.

We first prove that the following equations hold for any measurable $h' : \text{St}[\text{Name}] \rightarrow \mathbb{R}$:

$$\begin{aligned} &\int d\sigma_n \left(\mathbf{1}_{[\text{noerr}(c, \sigma_p \oplus \sigma_n \oplus f_*(\sigma_n))]} \cdot h'(\sigma_n) \right) \\ &= \int d\sigma_n \left(\mathbf{1}_{[\text{noerr}(c, \sigma_p \oplus \sigma_n \oplus f_*(\sigma_n))]} \cdot h'(\sigma_n) \cdot \sum_{K \subseteq \text{Name}} \mathbf{1}_{[\text{used}(c, \sigma_p \oplus \sigma_n \oplus f_*(\sigma_n), \sigma_n|_K)}] \right) \\ &= \sum_{K \subseteq \text{Name}} \int d\sigma_n \left(\mathbf{1}_{[\text{used}(c, \sigma_p \oplus \sigma_n \oplus f_*(\sigma_n), \sigma_n|_K)}] \cdot \mathbf{1}_{[\text{noerr}(c, \sigma_p \oplus \sigma_n \oplus f_*(\sigma_n))]} \cdot h'(\sigma_n) \right) \\ &= \sum_{K \subseteq \text{Name}} \int_{[K \rightarrow \mathbb{R}]} d\xi_n \int_{[\text{Name} \setminus K \rightarrow \mathbb{R}]} d\xi'_n \left(\mathbf{1}_{[\text{used}(c, \sigma_p \oplus (\xi_n \oplus \xi'_n) \oplus f_*(\xi_n \oplus \xi'_n), \xi_n)]} \right. \\ &\quad \left. \cdot \mathbf{1}_{[\text{noerr}(c, \sigma_p \oplus (\xi_n \oplus \xi'_n) \oplus f_*(\xi_n \oplus \xi'_n))]} \cdot h'(\xi_n \oplus \xi'_n) \right) \\ &= \sum_{K \subseteq \text{Name}} \int_{[K \rightarrow \mathbb{R}]} d\xi_n \left(\sum_{L \subseteq \text{Name}} \int_{[L \rightarrow \mathbb{R}]} d\xi'_n \left(\mathbf{1}_{[\text{dom}(\xi_n) \uplus \text{dom}(\xi'_n) = \text{Name}]} \cdot \mathbf{1}_{[\text{used}(c, \sigma_p \oplus (\xi_n \oplus \xi'_n) \oplus f_*(\xi_n \oplus \xi'_n), \xi_n)]} \right. \right. \\ &\quad \left. \left. \cdot \mathbf{1}_{[\text{noerr}(c, \sigma_p \oplus (\xi_n \oplus \xi'_n) \oplus f_*(\xi_n \oplus \xi'_n))]} \cdot h'(\xi_n \oplus \xi'_n) \right) \right) \\ &= \int d\xi_n \int d\xi'_n \left(\mathbf{1}_{[\text{dom}(\xi_n) \uplus \text{dom}(\xi'_n) = \text{Name}]} \cdot \mathbf{1}_{[\text{used}(c, \sigma_p \oplus (\xi_n \oplus \xi'_n) \oplus f_*(\xi_n \oplus \xi'_n), \xi_n)]} \right. \\ &\quad \left. \cdot \mathbf{1}_{[\text{noerr}(c, \sigma_p \oplus (\xi_n \oplus \xi'_n) \oplus f_*(\xi_n \oplus \xi'_n))]} \cdot h'(\xi_n \oplus \xi'_n) \right). \end{aligned}$$

All of these equations mean that one side of the equation is defined if and only if the other side is defined, and when both sides are defined, they are the same. The first equality holds because $\text{noerr}(c, \sigma_p \oplus \sigma_n \oplus f_*(\sigma_n))$ implies that there exists a unique $K \subseteq \text{Name}$ with $\text{used}(c, \sigma_p \oplus \sigma_n \oplus f_*(\sigma_n), \sigma_n|_K) = 1$. The second equality holds since Name is finite. The third equality holds because $\text{St}[\text{Name}]$ is isomorphic to $[K \rightarrow \mathbb{R}] \times [\text{Name} \setminus K \rightarrow \mathbb{R}]$. The fourth equality holds since $\xi'_n \in [L \rightarrow \mathbb{R}]$ with $L \neq \text{Name} \setminus K$ implies $\mathbf{1}_{[\text{dom}(\xi_n) \uplus \text{dom}(\xi'_n) = \text{Name}]} = 0$. The fifth equality holds by the definition of $\text{St}_\square[\text{Name}]$ and its underlying measure.

Using this result, we obtain the desired equation:

$$\begin{aligned} &\int d\sigma_n \left(\text{prs}(c)(\sigma_p, \sigma_n, f_*(\sigma_n)) \cdot h \left(\text{vals}(c)(\sigma_p, \sigma_n, f_*(\sigma_n)) \right) \right) \\ &= \int d\sigma_n \left(\mathbf{1}_{[\text{noerr}(c, \sigma_p \oplus \sigma_n \oplus f_*(\sigma_n))]} \cdot \text{prs}(c)(\sigma_p, \sigma_n, f_*(\sigma_n)) \cdot h \left(\text{vals}(c)(\sigma_p, \sigma_n, f_*(\sigma_n)) \right) \right) \\ &= \int d\xi_n \int d\xi'_n \left(\mathbf{1}_{[\text{dom}(\xi_n) \uplus \text{dom}(\xi'_n) = \text{Name}]} \cdot \mathbf{1}_{[\text{used}(c, \sigma_p \oplus (\xi_n \oplus \xi'_n) \oplus f_*(\xi_n \oplus \xi'_n), \xi_n)]} \right. \end{aligned}$$

$$\begin{aligned}
& \cdot \mathbf{1}_{[\text{noerr}(c, \sigma_p \oplus (\xi_n \oplus \xi'_n) \oplus f_*(\xi_n \oplus \xi'_n))]} \\
& \cdot \text{prs}(c)(\sigma_p, \xi_n \oplus \xi'_n, f_*(\xi_n \oplus \xi'_n)) \cdot h\left(\text{vals}(c)(\sigma_p, \xi_n \oplus \xi'_n, f_*(\xi_n \oplus \xi'_n))\right) \\
= & \int d\xi_n \int d\xi'_n \left(\mathbf{1}_{[\text{dom}(\xi_n) \uplus \text{dom}(\xi'_n) = \text{Name}]} \cdot \mathbf{1}_{[\text{used}(c, \sigma_p \oplus (\xi_n \oplus \xi'_n) \oplus f_*(\xi_n \oplus \xi'_n), \xi_n)]} \right. \\
& \left. \cdot \text{prs}(c)(\sigma_p, \xi_n \oplus \xi'_n, f_*(\xi_n \oplus \xi'_n)) \cdot h\left(\text{vals}(c)(\sigma_p, \xi_n \oplus \xi'_n, f_*(\xi_n \oplus \xi'_n))\right) \right) \\
= & \int d\xi_n \int d\xi'_n \left(\mathbf{1}_{[\text{dom}(\xi_n) \uplus \text{dom}(\xi'_n) = \text{Name}]} \cdot \mathbf{1}_{[\text{used}(c, \sigma_p \oplus (\xi_n \oplus \xi'_n) \oplus f_*(\xi_n \oplus \xi'_n), \xi_n)]} \right. \\
& \left. \cdot \text{prs}_{\square}(c)(\sigma_p, \xi_n) \cdot \left(\prod_{\mu \in \text{dom}(\xi'_n)} f_1(\xi'_n(\mu)) \right) \cdot h\left(\text{vals}_{\square}(c)(\sigma_p, \xi_n) \oplus (\lambda\mu \in \text{dom}(\xi'_n). f_2(\xi'_n(\mu)))\right) \right) \\
= & \int d\xi_n \int d\xi'_n \left(\mathbf{1}_{[\text{dom}(\xi_n) \uplus \text{dom}(\xi'_n) = \text{Name}]} \right. \\
& \left. \cdot \text{prs}_{\square}(c)(\sigma_p, \xi_n) \cdot \left(\prod_{\mu \in \text{dom}(\xi'_n)} f_1(\xi'_n(\mu)) \right) \cdot h\left(\text{vals}_{\square}(c)(\sigma_p, \xi_n) \oplus (\lambda\mu \in \text{dom}(\xi'_n). f_2(\xi'_n(\mu)))\right) \right) \\
= & \int d\xi_n \left(\text{prs}_{\square}(c)(\sigma_p, \xi_n) \cdot \int d\xi'_n \left(\mathbf{1}_{[\text{dom}(\text{vals}_{\square}(c)(\sigma_p, \xi_n) \uplus \text{dom}(\xi'_n) = \text{Name}]} \cdot \left(\prod_{\mu \in \text{dom}(\xi'_n)} f_1(\xi'_n(\mu)) \right) \right. \right. \\
& \left. \left. \cdot h\left(\text{vals}_{\square}(c)(\sigma_p, \xi_n) \oplus (\lambda\mu \in \text{dom}(\xi'_n). f_2(\xi'_n(\mu)))\right) \right) \right) \\
= & \int d\xi_n \left(\text{prs}_{\square}(c)(\sigma_p, \xi_n) \cdot g\left(\text{vals}_{\square}(c)(\sigma_p, \xi_n)\right) \right)
\end{aligned}$$

where $g : \text{St}_{\square}[\text{Name}] \rightarrow \mathbb{R}$ is defined as in the statement of this lemma, and each equation again means that one side of it is defined if and only if the other side is defined, and when both sides are defined, they are the same. The first and third equalities hold because $\text{prs}(c)(\sigma_p, \sigma_n, f_*(\sigma_n)) \neq 0$ implies $\mathbf{1}_{[\text{noerr}(c, \sigma_p \oplus \sigma_n \oplus f_*(\sigma_n))]} = 1$. The second equality uses the equation that we have shown in the previous paragraph. The fourth equality holds because of the following reason: if

$$\mathbf{1}_{[\text{dom}(\xi_n) \uplus \text{dom}(\xi'_n) = \text{Name}]} \cdot \mathbf{1}_{[\text{used}(c, \sigma_p \oplus \xi_n \oplus \sigma_r, \xi_n)]} = 1 \quad \text{for } \sigma_r \triangleq \xi'_n \oplus f_*(\xi_n \oplus \xi'_n),$$

then

$$\begin{aligned}
\text{prs}(c)(\sigma_p, \xi_n \oplus \xi'_n, f_*(\xi_n \oplus \xi'_n)) &= \llbracket c \rrbracket(\sigma_p \oplus \xi_n \oplus \sigma_r)(\text{like}) \cdot \prod_{\mu \in \text{Name}} \llbracket c \rrbracket(\sigma_p \oplus \xi_n \oplus \sigma_r)(pr_{\mu}) \\
&= \llbracket c \rrbracket(\sigma_p \oplus \xi_n \oplus \sigma_r)(\text{like}) \cdot \prod_{\mu \in \text{dom}(\xi_n)} \llbracket c \rrbracket(\sigma_p \oplus \xi_n \oplus \sigma_r)(pr_{\mu}) \\
&\quad \cdot \prod_{\mu \in \text{Name} \setminus \text{dom}(\xi_n)} \llbracket c \rrbracket(\sigma_p \oplus \xi_n \oplus \sigma_r)(pr_{\mu}) \\
&= \text{prs}_{\square}(c)(\sigma_p, \xi_n) \cdot \prod_{\mu \in \text{Name} \setminus \text{dom}(\xi_n)} \llbracket c \rrbracket(\sigma_p \oplus \xi_n \oplus \sigma_r)(pr_{\mu}) \\
&= \text{prs}_{\square}(c)(\sigma_p, \xi_n) \cdot \prod_{\mu \in \text{dom}(\xi'_n)} f_1(\xi'_n(\mu))
\end{aligned}$$

and

$$\begin{aligned}
\text{vals}(c)(\sigma_p, \xi_n \oplus \xi'_n, f_*(\xi_n \oplus \xi'_n)) &= \lambda\mu \in \text{Name}. \llbracket c \rrbracket(\sigma_p \oplus \xi_n \oplus \sigma_r)(\text{val}_{\mu}) \\
&= (\lambda\mu \in \text{dom}(\xi_n). \llbracket c \rrbracket(\sigma_p \oplus \xi_n \oplus \sigma_r)(\text{val}_{\mu})) \\
&\quad \oplus (\lambda\mu \in \text{Name} \setminus \text{dom}(\xi_n). (\sigma_p \oplus \xi_n \oplus \sigma_r)(\text{val}_{\mu})) \\
&= \text{vals}_{\square}(c)(\sigma_p, \xi_n) \oplus (\lambda\mu \in \text{Name} \setminus \text{dom}(\xi'_n). (\sigma_p \oplus \xi_n \oplus \sigma_r)(\text{val}_{\mu})) \\
&= \text{vals}_{\square}(c)(\sigma_p, \xi_n) \oplus (\lambda\mu \in \text{dom}(\xi'_n). f_2(\xi'_n(\mu))).
\end{aligned}$$

These equalities for $\text{prs}(c)$ and $\text{vals}(c)$ themselves hold for the below reasons:

- The first equalities hold by $used(c, \sigma_p \oplus \xi_n \oplus \sigma_r, \xi_n)$ and the definitions of $prs(c)$ and $vals(c)$.
- The second equalities hold by Lemma C.6-(2), which is applicable since $used(c, \sigma_p \oplus \xi_n \oplus \sigma_r, \xi_n)$.
- The third equalities hold by $used(c, \sigma_p \oplus \xi_n \oplus \sigma_r, \xi_n)$ and the definitions of $prs_{\square}(c)$ and $vals_{\square}(c)$.
- The fourth equalities hold by $\text{dom}(\xi_n) \uplus \text{dom}(\xi'_n) = \text{Name}$ and the definition of f_* .

Returning back to the main equations, we point out that the fifth equality comes from the next fact:

$$\left(\mathbf{1}_{[\text{dom}(\xi_n) \uplus \text{dom}(\xi'_n) = \text{Name}]} \cdot prs_{\square}(c)(\sigma_p, \xi_n) \right) \neq 0 \implies \mathbf{1}_{[used(c, \sigma_p \oplus (\xi_n \oplus \xi'_n) \oplus f_*(\xi_n \oplus \xi'_n), \xi_n)]} = 1.$$

The justification for this implication is given below:

- If the premise holds, then there exists σ_r such that $used(c, \sigma_p \oplus \xi_n \oplus \sigma_r, \xi_n)$. Since $\sigma_r(\text{like}) = 1 = f_*(-)(\text{like})$, $\sigma_p \oplus \xi_n \oplus \sigma_r$ and $\sigma_p \oplus (\xi_n \oplus \xi'_n) \oplus f_*(\xi_n \oplus \xi'_n)$ coincide on $\text{PVar} \cup \text{dom}(\xi_n) \cup \{\text{like}\}$. Thus, Lemma C.7 gives the conclusion.

Again back to the main equations, we note that the sixth equality holds since

$$\text{dom}(\xi_n) = \text{dom}(vals_{\square}(c)(\sigma_p, \xi_n)),$$

and the seventh equality follows from the definition of g . This completes the proof. \square

LEMMA C.4. *For all commands c and states $\sigma \in \text{St}$ such that $\llbracket c \rrbracket \sigma \in \text{St}$, we have*

$$\llbracket c \rrbracket \sigma(\text{cnt}_{\mu}) \geq \sigma(\text{cnt}_{\mu}) \quad \text{and} \quad \llbracket c \rrbracket \sigma(\mu) = \sigma(\mu)$$

for all $\mu \in \text{Name}$.

PROOF. We prove the lemma by induction on the structure of c . Let $\sigma \in \text{St}$ such that $\llbracket c \rrbracket \sigma \in \text{St}$.

Cases $c \equiv \text{skip}$, or $c \equiv (x := e)$, or $c \equiv \text{obs}(d, r)$. In these cases, $\llbracket c \rrbracket \sigma(\text{cnt}_{\mu}) = \sigma(\text{cnt}_{\mu})$ and $\llbracket c \rrbracket \sigma(\mu) = \mu$ for all μ . The claim of the lemma, thus, follows.

Case $c \equiv (x := \text{sam}(n, d, \lambda y. e'))$. Let $\mu \triangleq \llbracket n \rrbracket \sigma$, $p \triangleq \llbracket d \rrbracket \sigma$, and $r \triangleq \llbracket e'[\mu/y] \rrbracket \sigma$. Then,

$$\llbracket c \rrbracket \sigma = \sigma[x \mapsto r, \text{val}_{\mu} \mapsto r, \text{pr}_{\mu} \mapsto p(r), \text{cnt}_{\mu} \mapsto \sigma(\text{cnt}_{\mu}) + 1].$$

Thus, the claim of the lemma follows.

Case $c \equiv (c'; c'')$. Pick $\mu \in \text{Name}$. Then,

$$\llbracket c'; c'' \rrbracket \sigma(\mu) = \llbracket c'' \rrbracket (\llbracket c' \rrbracket \sigma)(\mu) = \llbracket c' \rrbracket \sigma(\mu) = \sigma(\mu).$$

Here the second and third equalities use induction hypothesis on c' and c'' , respectively. Also,

$$\begin{aligned} \llbracket c'; c'' \rrbracket \sigma(\text{cnt}_{\mu}) - \sigma(\text{cnt}_{\mu}) &= \left(\llbracket c'' \rrbracket (\llbracket c' \rrbracket \sigma)(\text{cnt}_{\mu}) - \llbracket c' \rrbracket \sigma(\text{cnt}_{\mu}) \right) + \left(\llbracket c' \rrbracket \sigma(\text{cnt}_{\mu}) - \sigma(\text{cnt}_{\mu}) \right) \\ &\geq 0. \end{aligned}$$

The inequality here uses induction hypothesis on c' and c'' .

Case $c \equiv (\text{if } b \{c'\} \text{ else } \{c''\})$. Assume that $\llbracket b \rrbracket \sigma = \text{true}$. We will prove the claims of the lemma under this assumption. The other case of $\llbracket b \rrbracket = \text{false}$ can be proved similarly. Pick $\mu \in \text{Name}$. Then, by induction hypothesis on c' ,

$$\llbracket c \rrbracket \sigma(\mu) = \llbracket c' \rrbracket \sigma(\mu) = \sigma(\mu) \quad \text{and} \quad \llbracket c \rrbracket \sigma(\text{cnt}_{\mu}) = \llbracket c' \rrbracket \sigma(\text{cnt}_{\mu}) \geq \sigma(\text{cnt}_{\mu}).$$

Case $c \equiv (\text{while } b \{c'\})$. Let \mathcal{T} be the following subset of $[\text{St} \rightarrow \text{St}_\perp]$:

$$f \in \mathcal{T} \iff \forall \sigma \in \text{St}. \left(f(\sigma) \neq \perp \implies \forall \mu \in \text{Name}. f(\sigma)(\mu) = \sigma(\mu) \wedge f(\sigma)(\text{cnt}_\mu) \geq \sigma(\text{cnt}_\mu) \right).$$

Let F be the operator on $[\text{St} \rightarrow \text{St}_\perp]$ whose least fixed point becomes the semantics of the loop c . The desired conclusion follows if we show that \mathcal{T} contains $\lambda \sigma. \perp$ and is closed under taking the limit of a chain in \mathcal{T} , and F preserves \mathcal{T} . The least element $\lambda \sigma. \perp$ belongs to \mathcal{T} since there are no states σ with $(\lambda \sigma. \perp)(\sigma) \neq \perp$. Consider an increasing sequence f_0, f_1, \dots in \mathcal{T} , and let $f_\infty \triangleq \bigsqcup_{n \in \mathbb{N}} f_n$. Pick σ such that $f_\infty(\sigma) \neq \perp$. Then, $f_\infty(\sigma) = f_m(\sigma)$ for some $m \in \mathbb{N}$. Since $f_m \in \mathcal{T}$, we have

$$f_m(\sigma)(\mu) = \sigma(\mu) \quad \text{and} \quad f_m(\sigma)(\text{cnt}_\mu) \geq \sigma(\text{cnt}_\mu)$$

for all $\mu \in \text{Name}$. Since $f_m(\sigma) = f_\infty(\sigma)$, we also have, for every $\mu \in \text{Name}$, $f_\infty(\sigma)(\mu) = \sigma(\mu)$ and $f_\infty(\sigma)(\text{cnt}_\mu) \geq \sigma(\text{cnt}_\mu)$, as desired. It remains to show that $F(f) \in \mathcal{T}$ for all $f \in \mathcal{T}$. Pick $f \in \mathcal{T}$ and $\sigma \in \text{St}$ such that $F(f)(\sigma) \in \text{St}$. If $\llbracket b \rrbracket \sigma = \text{false}$, we have $F(f)(\sigma) = \sigma$, and the claims of the lemma follow. Otherwise, $F(f)(\sigma) = f(\llbracket c' \rrbracket \sigma)$. Pick $\mu \in \text{Name}$. Then, by induction hypothesis on c' and the membership $f \in \mathcal{T}$,

$$F(f)(\sigma)(\mu) = f(\llbracket c' \rrbracket \sigma)(\mu) = \llbracket c' \rrbracket \sigma(\mu) = \sigma(\mu),$$

and

$$F(f)(\sigma)(\text{cnt}_\mu) = f(\llbracket c' \rrbracket \sigma)(\text{cnt}_\mu) \geq \llbracket c' \rrbracket \sigma(\text{cnt}_\mu) \geq \sigma(\text{cnt}_\mu).$$

We have just shown that $F(f) \in \mathcal{T}$, as desired. \square

Definition C.5. Define used_- as the predicate used but without the condition that like should be 1. That is, for all commands c , states $\sigma \in \text{St}$, and $\xi_n \in \text{St}_\square[\text{Name}]$,

$$\begin{aligned} \text{used}_-(c, \sigma, \xi_n) \iff & \llbracket c \rrbracket \sigma \in \text{St} \wedge (\llbracket c \rrbracket \sigma(\text{cnt}_\mu) - \sigma(\text{cnt}_\mu) \leq 1 \text{ for all } \mu \in \text{Name}) \\ & \wedge \xi_n = \sigma|_{\text{dom}(\xi_n)} \\ & \wedge \text{dom}(\xi_n) = \{\mu \in \text{Name} \mid \llbracket c \rrbracket \sigma(\text{cnt}_\mu) - \sigma(\text{cnt}_\mu) = 1\}. \end{aligned}$$

LEMMA C.6. *Let c be a command, $\sigma_0, \sigma_1 \in \text{St}$, and $\xi_n \in \text{St}_\square[\text{Name}]$. Suppose that $\text{used}_-(c, \sigma_0, \xi_n)$ and $\sigma_1|_V = \sigma_0|_V$ for $V \triangleq \text{PVar} \cup \text{dom}(\xi_n)$. Then, the following properties hold:*

- (1) $\llbracket c \rrbracket \sigma_1 \in \text{St}$.
- (2) $\llbracket c \rrbracket \sigma_1(a) = \sigma_1(a)$ for all $a \in \{\text{pr}_\mu, \text{val}_\mu, \text{cnt}_\mu \mid \mu \in \text{Name} \setminus \text{dom}(\xi_n)\}$.
- (3) $\llbracket c \rrbracket \sigma_1(v) = \llbracket c \rrbracket \sigma_0(v)$ for all $v \in \text{PVar} \cup \{\text{pr}_\mu, \text{val}_\mu \mid \mu \in \text{dom}(\xi_n)\}$.
- (4) $\llbracket c \rrbracket \sigma_1(\text{like}) = \llbracket c \rrbracket \sigma_0(\text{like})$, if $\sigma_0(\text{like}) = \sigma_1(\text{like})$.
- (5) $\llbracket c \rrbracket \sigma_1(a) - \sigma_1(a) = \llbracket c \rrbracket \sigma_0(a) - \sigma_0(a)$ for all $a \in \{\text{cnt}_\mu \mid \mu \in \text{Name}\}$.

PROOF. For $\sigma'_0, \sigma'_1 \in \text{St}$ and $\xi'_n \in \text{St}_\square[\text{Name}]$, write

$$\sigma'_0 \sim_{\xi'_n} \sigma'_1$$

to mean that $\sigma'_0|_V = \sigma'_1|_V$ for $V \triangleq \text{PVar} \cup \text{dom}(\xi'_n)$. Note that using this notation, we can write the conditions of the lemma as follows:

$$\text{used}_-(c, \sigma_0, \xi_n) \wedge \sigma_0 \sim_{\xi_n} \sigma_1.$$

We will prove, by induction on the structure of c , that these conditions imply the five properties claimed by the lemma. Our proof will sometimes use a simple observation that the five properties claimed by the lemma and the relationship $\sigma_0 \sim_{\xi_n} \sigma_1$ imply $\text{used}_-(c, \sigma_1, \xi_n)$. One consequence of the observation is that if our lemma holds, its five properties also hold with σ_0 and σ_1 swapped. We will often use this consequence.

Case $c \equiv \text{skip}$. In this case, $\llbracket c \rrbracket \sigma_0 = \sigma_0$ and $\llbracket c \rrbracket \sigma_1 = \sigma_1$. From these equalities, the claimed properties (1), (2), (4) and (5) follow. For the remaining property (3), we note that $\text{dom}(\xi_n) = \emptyset$ and the property, thus, follows from $\sigma_0 \sim_{\xi_n} \sigma_1$.

Case $c \equiv (x := e)$. In this case, $\llbracket c \rrbracket \sigma_0 = \sigma_0[x \mapsto \llbracket e \rrbracket \sigma_0]$ and $\llbracket c \rrbracket \sigma_1 = \sigma_1[x \mapsto \llbracket e \rrbracket \sigma_1]$. The results are not \perp , and they are identical to the pre-states σ_0 and σ_1 as far as auxiliary variables in AVar are concerned. Also, expressions in commands do not depend on variables other than program variables, so that $\sigma_0 \sim_{\xi_n} \sigma_1$ gives $\llbracket e \rrbracket \sigma_0 = \llbracket e \rrbracket \sigma_1$ and $\llbracket c \rrbracket \sigma_0(x) = \llbracket c \rrbracket \sigma_1(x)$ for all $x \in \text{PVar}$. From all of these observations, the claimed properties (1)–(5) follow.

Case $c \equiv (x := \text{sam}(n, d, \lambda y. e'))$. Since $\sigma_0(x) = \sigma_1(x)$ for all $x \in \text{PVar}$, we have $\llbracket n \rrbracket \sigma_0 = \llbracket n \rrbracket \sigma_1$ and $\llbracket d \rrbracket \sigma_0 = \llbracket d \rrbracket \sigma_1$. Let $\mu \triangleq \llbracket n \rrbracket \sigma_0$, $p \triangleq \llbracket d \rrbracket \sigma_0$, and $r \triangleq \llbracket e'[\mu/y] \rrbracket \sigma_0$. By the semantics of the sample commands, we have

$$\llbracket c \rrbracket \sigma_0 = \sigma_0[x \mapsto r, \text{val}_\mu \mapsto r, \text{pr}_\mu \mapsto p(\sigma_0(\mu)), \text{cnt}_\mu \mapsto \sigma_0(\text{cnt}_\mu) + 1].$$

Since $\text{used}_-(c, \sigma_0, \xi_n)$ holds, we have $\xi_n = \sigma_0|_{\{\mu\}}$, which in turn implies $\sigma_0(\mu) = \sigma_1(\mu)$ because $\sigma_0 \sim_{\xi_n} \sigma_1$. Thus, $\llbracket e'[\mu/y] \rrbracket \sigma_1 = \llbracket e'[\mu/y] \rrbracket \sigma_0 = r$, and

$$\begin{aligned} \llbracket c \rrbracket \sigma_1 &= \sigma_1[x \mapsto \llbracket e'[\mu/y] \rrbracket \sigma_1, \text{val}_\mu \mapsto \llbracket e'[\mu/y] \rrbracket \sigma_1, \text{pr}_\mu \mapsto p(\sigma_1(\mu)), \text{cnt}_\mu \mapsto \sigma_1(\text{cnt}_\mu) + 1] \\ &= \sigma_1[x \mapsto r, \text{val}_\mu \mapsto r, \text{pr}_\mu \mapsto p(\sigma_0(\mu)), \text{cnt}_\mu \mapsto \sigma_1(\text{cnt}_\mu) + 1]. \end{aligned}$$

The RHS of the last equality implies that the five properties claimed by the lemma hold.

Case $c \equiv \text{obs}(d, r)$. We have $\llbracket d \rrbracket \sigma_0 = \llbracket d \rrbracket \sigma_1$ since $\sigma_0(x) = \sigma_1(x)$ for all $x \in \text{PVar}$. Let $p \triangleq \llbracket d \rrbracket \sigma_0$. Then,

$$\llbracket c \rrbracket \sigma_0 = \sigma_0[\text{like} \mapsto \sigma_0(\text{like}) \cdot p(r)] \quad \text{and} \quad \llbracket c \rrbracket \sigma_1 = \sigma_1[\text{like} \mapsto \sigma_1(\text{like}) \cdot p(r)].$$

Also, $\text{dom}(\xi_n) = \emptyset$ since $\text{used}_-(c, \sigma_0, \xi_n)$ holds. From what we have proved and also the agreement of σ_0 and σ_1 on program variables, the five properties claimed by the lemma follow.

Case $c \equiv (c'; c'')$. Since $\llbracket c \rrbracket \sigma_0 = \llbracket c' \rrbracket^\dagger(\llbracket c' \rrbracket \sigma_0) \in \text{St}$, we have $\llbracket c' \rrbracket \sigma_0 \in \text{St}$. Let

$$\begin{aligned} \sigma'_0 &\triangleq \llbracket c' \rrbracket \sigma_0, \\ N_0 &\triangleq \{\mu \in \text{Name} \mid \llbracket c'' \rrbracket \sigma'_0(\text{cnt}_\mu) - \sigma_0(\text{cnt}_\mu) = 1\}, \\ N'_0 &\triangleq \{\mu \in \text{Name} \mid \sigma'_0(\text{cnt}_\mu) - \sigma_0(\text{cnt}_\mu) = 1\}. \end{aligned}$$

Then, $N_0 = \text{dom}(\xi_n)$ because $\text{used}_-(c'; c'', \sigma_0, \xi_n)$ holds. We will prove the following facts:

- (1) $N'_0 \subseteq N_0$.
- (2) Let $\xi'_n \triangleq \xi_n|_{N'_0}$, and $\xi''_n \triangleq \xi_n|_{(N_0 \setminus N'_0)}$. Then, $\text{used}_-(c', \sigma_0, \xi'_n)$ and $\text{used}_-(c'', \sigma'_0, \xi''_n)$ hold.
- (3) $\llbracket c' \rrbracket \sigma_1 \in \text{St}$.
- (4) Let $\sigma'_1 \triangleq \llbracket c' \rrbracket \sigma_1$. Then, $\sigma'_0 \sim_{\xi''_n} \sigma'_1$.

These four facts imply the five properties claimed by the lemma. Here is the reason. Note that $\sigma_0 \sim_{\xi'_n} \sigma_1$ since $\text{dom}(\xi'_n) = N'_0 \subseteq N_0 = \text{dom}(\xi_n)$. This relationship between σ_0 and σ_1 and the second fact let us use induction hypothesis on $(c', \xi'_n, \sigma_0, \sigma_1)$. Also, the second and fourth facts allow us to use induction hypothesis on $(c'', \xi''_n, \sigma'_0, \sigma'_1)$. We can derive the five properties from what we get from these two applications of induction hypothesis:

- (1) By induction hypothesis on $(c'', \xi''_n, \sigma'_0, \sigma'_1)$, we have $\llbracket c'; c'' \rrbracket \sigma_1 = \llbracket c'' \rrbracket \sigma'_1 \in \text{St}$.
- (2) For all $a \in \{\text{pr}_\mu, \text{val}_\mu, \text{cnt}_\mu \mid \mu \in \text{Name} \setminus \text{dom}(\xi_n)\}$,

$$\llbracket c'; c'' \rrbracket \sigma_1(a) = \llbracket c'' \rrbracket \sigma'_1(a) = \sigma'_1(a) = \llbracket c' \rrbracket \sigma_1(a) = \sigma_1(a).$$

The second equality comes from induction hypothesis on $(c'', \xi_n'', \sigma'_0, \sigma'_1)$ and $\text{dom}(\xi_n'') \subseteq \text{dom}(\xi_n)$, and the fourth equality from induction hypothesis on $(c', \xi_n', \sigma_0, \sigma_1)$ and $\text{dom}(\xi_n') \subseteq \text{dom}(\xi_n)$.

- (3) For all $v \in \text{PVar} \cup \{pr_\mu, val_\mu \mid \mu \in \text{dom}(\xi_n'')\}$, by induction hypothesis on $(c'', \xi_n'', \sigma'_0, \sigma'_1)$,

$$\llbracket c'; c'' \rrbracket \sigma_1(v) = \llbracket c'' \rrbracket \sigma'_1(v) = \llbracket c'' \rrbracket \sigma'_0(v) = \llbracket c'; c'' \rrbracket \sigma_0(v).$$

Also, for all $a \in \{pr_\mu, val_\mu \mid \mu \in \text{dom}(\xi_n')\}$, we have $a \in \{pr_\mu, val_\mu \mid \mu \in \text{Name} \setminus \text{dom}(\xi_n'')\}$, and we can calculate:

$$\begin{aligned} \llbracket c'; c'' \rrbracket \sigma_1(a) &= \llbracket c'' \rrbracket \sigma'_1(a) = \sigma'_1(a) \\ &= \llbracket c' \rrbracket \sigma_1(a) = \llbracket c' \rrbracket \sigma_0(a) \\ &= \sigma'_0(a) = \llbracket c'' \rrbracket \sigma'_0(a) = \llbracket c'; c'' \rrbracket \sigma_0(a). \end{aligned}$$

The second equality uses induction hypothesis on $(c'', \xi_n'', \sigma'_0, \sigma'_1)$, and the fourth comes from the induction hypothesis on $(c', \xi_n', \sigma_0, \sigma_1)$. The sixth equality follows from induction hypothesis applied to $(c'', \xi_n'', \sigma'_0, \sigma'_1)$ and again to the same tuple but with σ'_0 and σ'_1 swapped.

- (4) If $\sigma_0(\text{like}) = \sigma_1(\text{like})$, by induction hypothesis on $(c', \xi_n', \sigma_0, \sigma_1)$,

$$\sigma'_0(\text{like}) = \llbracket c' \rrbracket \sigma_0(\text{like}) = \llbracket c' \rrbracket \sigma_1(\text{like}) = \sigma'_1(\text{like}),$$

which in turn implies, by induction hypothesis on $(c'', \xi_n'', \sigma'_0, \sigma'_1)$,

$$\llbracket c'; c'' \rrbracket \sigma_0(\text{like}) = \llbracket c'' \rrbracket \sigma'_0(\text{like}) = \llbracket c'' \rrbracket \sigma'_1(\text{like}) = \llbracket c'; c'' \rrbracket \sigma_1(\text{like}).$$

- (5) For all $a \in \{cnt_\mu \mid \mu \in \text{Name}\}$,

$$\begin{aligned} \llbracket c'; c'' \rrbracket \sigma_1(a) - \sigma_1(a) &= \llbracket c'; c'' \rrbracket \sigma_1(a) - \llbracket c' \rrbracket \sigma_1(a) + \llbracket c' \rrbracket \sigma_1(a) - \sigma_1(a) \\ &= \llbracket c'' \rrbracket \sigma'_1(a) - \sigma'_1(a) + \llbracket c' \rrbracket \sigma_1(a) - \sigma_1(a) \\ &= \llbracket c'' \rrbracket \sigma'_0(a) - \sigma'_0(a) + \llbracket c' \rrbracket \sigma_0(a) - \sigma_0(a) \\ &= \llbracket c'; c'' \rrbracket \sigma_0(a) - \sigma_0(a). \end{aligned}$$

The only non-trivial inequality is the third one, and it follows from induction hypothesis on $(c', \xi_n', \sigma_0, \sigma_1)$ and $(c'', \xi_n'', \sigma'_0, \sigma'_1)$.

We prove the four facts as follows:

- (1) Let $\mu \in N'_0$. Since $\text{used}_-(c'; c'', \sigma_0, \xi_n)$, we have

$$\llbracket c'' \rrbracket \sigma'_0(\mu) - \sigma_0(\mu) = \llbracket c'; c'' \rrbracket \sigma_0 - \sigma_0(\mu) \leq 1.$$

Also, by Lemma C.4 and the definition of N'_0 ,

$$\llbracket c'' \rrbracket \sigma'_0(\mu) - \sigma_0(\mu) \geq \sigma'_0(\mu) - \sigma_0(\mu) = 1.$$

Thus, $\llbracket c'' \rrbracket \sigma'_0(\mu) - \sigma_0(\mu) = 1$, which implies that $\mu \in N_0$, as desired.

- (2) We should show that $\text{used}_-(c', \sigma_0, \xi_n)$ and $\text{used}_-(c'', \sigma'_0, \xi_n')$ hold. The conjuncts in the definition of $\text{used}_-(c', \sigma_0, \xi_n)$ except the second follow immediately from $\text{used}_-(c'; c'', \sigma_0, \xi_n)$ and the definition of ξ_n . For the remaining second conjunct, we use Lemma C.4 and $\text{used}_-(c'; c'', \sigma_0, \xi_n)$, and prove the conjunct as shown below: for all $\mu \in \text{Name}$,

$$\llbracket c' \rrbracket \sigma_0(cnt_\mu) - \sigma_0(cnt_\mu) \leq \llbracket c'' \rrbracket (\llbracket c' \rrbracket \sigma_0)(cnt_\mu) - \sigma_0(cnt_\mu) \leq 1.$$

For $\text{used}_-(c'', \sigma'_0, \xi_n')$, we first note that the first and third conjuncts in its definition are direct consequences of $\text{used}_-(c'; c'', \sigma_0, \xi_n)$ and the definition of ξ_n' . We prove the second conjunct in the definition as follows: for all $\mu \in \text{Name}$,

$$\llbracket c'' \rrbracket \sigma'_0(cnt_\mu) - \sigma'_0(cnt_\mu) = \llbracket c'; c'' \rrbracket \sigma_0(cnt_\mu) - \llbracket c' \rrbracket cnt_0(cnt_\mu)$$

$$\begin{aligned} &\leq \llbracket c'; c'' \rrbracket \sigma_0(cnt_\mu) - cnt_0(cnt_\mu) \\ &\leq 1. \end{aligned}$$

The first inequality uses Lemma C.4, and the second comes from $used_-(c'; c'', \sigma_0, \xi_n)$. It remains to show the fourth conjunct in the definition of $used_-(c'', \sigma'_0, \xi''_n)$, which we do below: for all $\mu \in \text{Name}$,

$$\begin{aligned} \llbracket c'' \rrbracket \sigma'_0(cnt_\mu) - \sigma'_0(cnt_\mu) &= 1 \\ \iff \llbracket c'' \rrbracket \sigma'_0(cnt_\mu) - \sigma'_0(cnt_\mu) &= 1 \wedge \sigma'_0(cnt_\mu) - \sigma_0(cnt_\mu) = 0 \\ \iff \mu \in N_0 \wedge \mu \notin N'_0 & \\ \iff \mu \in \text{dom}(\xi''_n). & \end{aligned}$$

The first equivalence comes from Lemma C.4 and $\llbracket c'' \rrbracket \sigma'_0(cnt_\mu) - \sigma_0(cnt_\mu) \leq 1$, which holds because of $used_-(c'; c'', \sigma_0, \xi_n)$. The second equivalence follows from the definitions of N_0 and N'_0 .

- (3) Since $\sigma_0 \sim_{\xi_n} \sigma_1$ implies $\sigma_0 \sim_{\xi'_n} \sigma_1$ and we have $used_-(c', \sigma_0, \xi'_n)$, we can apply induction hypothesis to $(c', \xi'_n, \sigma_0, \sigma_1)$, and get $\llbracket c' \rrbracket \sigma_1 \in \text{St}$.
- (4) We continue our reasoning in the previous item, and derive from induction hypothesis on $(c', \xi'_n, \sigma_0, \sigma_1)$ the fact that for all $x \in \text{PVar}$,

$$\sigma'_0(x) = \llbracket c' \rrbracket \sigma_0(x) = \llbracket c' \rrbracket \sigma_1(x) = \sigma'_1(x).$$

Also, for all $\mu \in \text{dom}(\xi''_n)$,

$$\sigma'_0(\mu) = \sigma_0(\mu) = \sigma_1(\mu) = \sigma'_1(\mu),$$

where the first and third equalities come from Lemma C.4, and the second equality follows from the assumption that $\sigma_0 \sim_{\xi_n} \sigma_1$.

Case $c \equiv (\text{if } b \{c'\} \text{ else } c'')$. Assume that $\llbracket b \rrbracket \sigma_0 = \text{true}$. Then, $\llbracket c \rrbracket \sigma_0 = \llbracket c' \rrbracket \sigma_0$. We prove the five properties claimed by the lemma under this assumption. The proof for the other possibility, namely, $\llbracket b \rrbracket \sigma_0 = \text{false}$ is similar. Since $\sigma_0 \sim_{\xi_n} \sigma_1$, the states σ_0 and σ_1 coincide for the values of program variables. Thus, $\llbracket b \rrbracket \sigma_1 = \text{true}$, and $\llbracket c \rrbracket \sigma_1 = \llbracket c' \rrbracket \sigma_1$. Since $\llbracket c \rrbracket \sigma_0 = \llbracket c' \rrbracket \sigma_0$ as well, it suffices to show the five properties claimed by the lemma for $(c', \sigma_0, \sigma_1, \xi_n)$. This sufficient condition follows from induction hypothesis on $(c', \sigma_0, \sigma_1, \xi_n)$, since $used_-(c, \sigma_0, \xi_n)$ and $\llbracket b \rrbracket \sigma_0 = \text{true}$ imply $used_-(c', \sigma_0, \xi_n)$.

Case $c \equiv (\text{while } b \{c'\})$. Consider the version of $used_-$ where the first parameter can be a state transformer $f : \text{St} \rightarrow \text{St}_\perp$, instead of a command. Similarly, consider the version of the five properties claimed by the lemma where we use a state transformer $f : \text{St} \rightarrow \text{St}_\perp$, again instead of a command. We denote both versions by $used_-(f, \sigma''_0, \xi''_n)$ and $\varphi(f, \xi''_n, \sigma''_0, \sigma''_1)$. Let \mathcal{T} be the subset of $\text{St} \rightarrow \text{St}_\perp$ defined by

$$f \in \mathcal{T} \iff \forall \sigma''_0, \sigma''_1 \in \text{St}. \forall \xi''_n \in \text{St}_\square. \left(\left(used_-(f, \sigma''_0, \xi''_n) \wedge \sigma''_0 \sim_{\xi''_n} \sigma''_1 \right) \implies \varphi(f, \xi''_n, \sigma''_0, \sigma''_1) \right),$$

and $F : [\text{St} \rightarrow \text{St}_\perp] \rightarrow [\text{St} \rightarrow \text{St}_\perp]$ be the following operator used in the semantics of the loop $\llbracket c \rrbracket$:

$$F(f)(\sigma) \triangleq \text{if } (\llbracket b \rrbracket \sigma = \text{true}) \text{ then } f^\dagger(\llbracket c' \rrbracket \sigma) \text{ else } \sigma.$$

We will show that \mathcal{T} contains $\lambda\sigma. \perp$ and is closed under taking the least upper bound of an increasing chain in $[\text{St} \rightarrow \text{St}_\perp]$, and the operator F preserves \mathcal{T} . These three conditions imply that $\llbracket c \rrbracket$ is in \mathcal{T} , which in turn gives the five properties claimed by the lemma.

The first condition holds simply because $used_-(\lambda\sigma. \perp, \sigma''_0, \xi''_n)$ is false for all σ''_0 and ξ''_n . To prove the closure under the least upper bound of a chain, consider an increasing sequence f_0, f_1, \dots in \mathcal{T} . Let $f_\infty \triangleq \bigsqcup_{n \in \mathbb{N}} f_n$. Consider $\sigma''_0, \sigma''_1 \in \text{St}$ and $\xi''_n \in \text{St}_\square$ such that $\sigma''_0 \sim_{\xi''_n} \sigma''_1$ and $used_-(f_\infty, \sigma''_0, \xi''_n)$. We should show that $\varphi(f_\infty, \xi''_n, \sigma''_0, \sigma''_1)$ holds. By the definition of f_∞ , there exists $m \in \mathbb{N}$ such that

$f_\infty(\sigma_0) = f_m(\sigma_0)$. Then, the assumption $used_-(f_\infty, \sigma_0'', \xi_n'')$ implies $used_-(f_m, \sigma_0'', \xi_n'')$. This in turn gives $\varphi(f_m, \xi_n'', \sigma_0'', \sigma_1'')$ because $f_m \in \mathcal{T}$. By what we have proved and the definition of f_∞ , we have

$$f_m(\sigma_0'') = f_\infty(\sigma_0'') \in \text{St} \quad \text{and} \quad f_m(\sigma_1'') = f_\infty(\sigma_1'') \in \text{St}.$$

Thus, $\varphi(f_m, \xi_n'', \sigma_0'', \sigma_1'')$ entails $\varphi(f_\infty, \xi_n'', \sigma_0'', \sigma_1'')$, as desired. It remains to show that $F(f) \in \mathcal{T}$ for all $f \in \mathcal{T}$. Pick $f \in \mathcal{T}$. We first replay our proof for the sequential-composition case after viewing $f^\dagger \circ \llbracket c' \rrbracket$ as the sequential composition of c' and f . This replay, then, gives the membership $f^\dagger \circ \llbracket c' \rrbracket \in \mathcal{T}$. Next, we replay our proof for the if case on $F(f)$ after viewing $f^\dagger \circ \llbracket c' \rrbracket$ as the true branch and $\lambda\sigma. \sigma = \llbracket \text{skip} \rrbracket$ as the false branch. This replay implies the required $F(f) \in \mathcal{T}$. \square

LEMMA C.7. *Let c be a command, $\sigma, \sigma' \in \text{St}$, and $\xi_n \in \text{St}_\square[\text{Name}]$.*

- *If $\sigma|_V = \sigma'|_V$ for $V \triangleq \text{PVar} \cup \text{dom}(\xi_n) \cup \{\text{like}\}$, then $used(c, \sigma, \xi_n)$ implies $used(c, \sigma', \xi_n)$.*
- *If $\sigma|_U = \sigma'|_U$ for $U \triangleq \text{PVar} \cup \text{dom}(\xi_n)$, then $used_-(c, \sigma, \xi_n)$ implies $used_-(c, \sigma', \xi_n)$.*

PROOF. Assume the settings in the statement of this lemma. For the first claim, assume $used(c, \sigma, \xi_n)$. Then, by the definition of $used$ and $noerr$,

$$\begin{aligned} \llbracket c \rrbracket \sigma \in \text{St} \wedge (\forall \mu \in \text{Name}. \llbracket c \rrbracket \sigma(\text{cnt}_\mu) - \sigma(\text{cnt}_\mu) \leq 1) \wedge (\sigma(\text{like}) = 1) \\ \wedge (\xi_n = \sigma|_{\text{dom}(\xi_n)}) \wedge (\text{dom}(\xi_n) = \{\mu \in \text{Name} \mid \llbracket c \rrbracket \sigma(\text{cnt}_\mu) - \sigma(\text{cnt}_\mu) = 1\}). \end{aligned}$$

From this and Lemma C.6 (which is applicable since $used(c, \sigma, \xi_n)$ and $\sigma|_V = \sigma'|_V$), we obtain

$$\begin{aligned} \llbracket c \rrbracket \sigma' \in \text{St} \wedge (\forall \mu \in \text{Name}. \llbracket c \rrbracket \sigma'(\text{cnt}_\mu) - \sigma'(\text{cnt}_\mu) \leq 1) \wedge (\sigma'(\text{like}) = 1) \\ \wedge (\xi_n = \sigma'|_{\text{dom}(\xi_n)}) \wedge (\text{dom}(\xi_n) = \{\mu \in \text{Name} \mid \llbracket c \rrbracket \sigma'(\text{cnt}_\mu) - \sigma'(\text{cnt}_\mu) = 1\}). \end{aligned}$$

Note that we have the first clause by Lemma C.6-(1), the second and fifth clauses by Lemma C.6-(5), and the third and fourth clauses by $\sigma|_V = \sigma'|_V$. Hence, $used(c, \sigma', \xi_n)$ holds. The proof of the second claim is exactly the same except that we apply Lemma C.6 to $\sigma|_U = \sigma'|_U$ to prove only the four clauses of $used$ that exclude $\sigma'(\text{like}) = 1$. \square

C.3 Proof of Lemma C.3

PROOF OF LEMMA C.3. We prove this lemma by induction on the structure of c . Let $g : \text{St}[\text{PVar}] \times \text{St}_\square[\text{Name}] \rightarrow \mathbb{R}$ be a measurable function and $\sigma_p \in \text{St}[\text{PVar}]$. In this proof, each equation involving integrals means (otherwise noted) that one side of the equation is defined if and only if the other side is defined, and when both sides are defined, they are the same.

Case $c \equiv \text{skip}$, $c \equiv (x := e)$, or $c \equiv \text{obs}(d, r)$. In this case, $\bar{c}^\pi \equiv c$ so the desired equation holds.

Case $c \equiv (x := \text{sam}(n, d, \lambda y. e))$. If $(n, d, \lambda y. e) \notin \text{dom}(\pi)$, then $\bar{c}^\pi \equiv c$ and thus the desired equation holds. So assume that $\pi(n, d, \lambda y. e) = (d', \lambda y'. e')$ for some d' and $\lambda y'. e'$. Then, $\bar{c}^\pi \equiv (x := \text{sam}(n, d', \lambda y'. e'))$.

First, by the validity of π , for all states $\sigma \in \text{St}$ and measurable subsets $A \subseteq \mathbb{R}$,

$$\int \mathbf{1}_{\llbracket [e[r/y]] \rrbracket \sigma \in A} \cdot \llbracket [d] \rrbracket \sigma(r) \, dr = \int \mathbf{1}_{\llbracket [e'[r/y']] \rrbracket \sigma \in A} \cdot \llbracket [d'] \rrbracket \sigma(r) \, dr,$$

where both sides are always defined. Using this and the monotone convergence theorem, we can show that for all measurable $f : \mathbb{R} \rightarrow \mathbb{R}$,

$$\int f(\llbracket [e[r/y]] \rrbracket \sigma) \cdot \llbracket [d] \rrbracket \sigma(r) \, dr = \int f(\llbracket [e'[r/y']] \rrbracket \sigma) \cdot \llbracket [d'] \rrbracket \sigma(r) \, dr. \quad (17)$$

Next, choose any $\sigma_{r_0} \in \text{St}[\text{Var} \setminus \text{PVar}]$. Since $\text{fv}(n) \subseteq \text{PVar}$, there exists $\mu \in \text{Name}$ such that

$$\llbracket [n] \rrbracket (\sigma_p \oplus \sigma_r) = \mu \quad \text{for all } \sigma_r \in \text{St}[\text{Var} \setminus \text{PVar}].$$

Using this and $fv(e), fv(d) \subseteq \text{PVar}$, we obtain the following: for any $\sigma_r \in \text{St}[\text{Var} \setminus \text{PVar}]$,

$$\begin{aligned} \llbracket c \rrbracket(\sigma_p \oplus \sigma_r)(like) &= 1, \\ \llbracket c \rrbracket(\sigma_p \oplus \sigma_r)(pr_\mu) &= \llbracket d \rrbracket(\sigma_p \oplus \sigma_r)(\sigma_r(\mu)) = \llbracket d \rrbracket(\sigma_p \oplus \sigma_{r_0})(\sigma_r(\mu)), \\ \llbracket c \rrbracket(\sigma_p \oplus \sigma_r)(val_\mu) &= \llbracket e[\sigma_r(\mu)/y] \rrbracket(\sigma_p \oplus \sigma_r) = \llbracket e[\sigma_r(\mu)/y] \rrbracket(\sigma_p \oplus \sigma_{r_0}), \\ \llbracket c \rrbracket(\sigma_p \oplus \sigma_r)(cnt_\mu) &= 1, \quad \llbracket c \rrbracket(\sigma_p \oplus \sigma_r)(cnt_{\mu'}) = 0 \quad \text{for } \mu' \neq \mu, \\ \llbracket c \rrbracket(\sigma_p \oplus \sigma_r)|_{\text{PVar}} &= \sigma_p[x \mapsto \llbracket e[\sigma_r(\mu)/y] \rrbracket(\sigma_p \oplus \sigma_{r_0})]. \end{aligned}$$

This implies that for any $\xi_n \in \text{St}_\square[\text{Name}]$, if $pr_{\square}(c)(\sigma_p, \xi_n) \neq 0$, then

$$\begin{aligned} \text{dom}(\xi_n) &= \{\mu\}, \\ pr_{\square}(c)(\sigma_p, \xi_n) &= 1 \cdot \llbracket d \rrbracket(\sigma_p \oplus \sigma_{r_0})(\xi_n(\mu)), \\ pvars_{\square}(c)(\sigma_p, \xi_n) &= \sigma_p[x \mapsto \llbracket e[\xi_n(\mu)/y] \rrbracket(\sigma_p \oplus \sigma_{r_0})], \\ vals_{\square}(c)(\sigma_p, \xi_n) &= [\mu \mapsto \llbracket e[\xi_n(\mu)/y] \rrbracket(\sigma_p \oplus \sigma_{r_0})]. \end{aligned}$$

Note that the same equations hold for \bar{c}^π , except that we replace d , e , and y in the RHS of the above equations by d' , e' , and y' . Using these, we obtain:

$$\begin{aligned} & \int d\xi_n \left(pr_{\square}(c)(\sigma_p, \xi_n) \cdot g \left(pvars_{\square}(c)(\sigma_p, \xi_n), vals_{\square}(c)(\sigma_p, \xi_n) \right) \right) \\ &= \int_{[\{\mu\} \rightarrow \mathbb{R}]} d\xi_n \left(\llbracket d \rrbracket(\sigma_p \oplus \sigma_{r_0})(\xi_n(\mu)) \cdot \widehat{g} \left(\llbracket e[\xi_n(\mu)/y] \rrbracket(\sigma_p \oplus \sigma_{r_0}) \right) \right) \\ &= \int_{\mathbb{R}} dr \left(\llbracket d \rrbracket(\sigma_p \oplus \sigma_{r_0})(r) \cdot \widehat{g} \left(\llbracket e[r/y] \rrbracket(\sigma_p \oplus \sigma_{r_0}) \right) \right) \\ &= \int_{\mathbb{R}} dr \left(\llbracket d' \rrbracket(\sigma_p \oplus \sigma_{r_0})(r) \cdot \widehat{g} \left(\llbracket e'[r/y'] \rrbracket(\sigma_p \oplus \sigma_{r_0}) \right) \right) \\ &= \int_{[\{\mu\} \rightarrow \mathbb{R}]} d\xi_n \left(\llbracket d' \rrbracket(\sigma_p \oplus \sigma_{r_0})(\xi_n(\mu)) \cdot \widehat{g} \left(\llbracket e'[\xi_n(\mu)/y'] \rrbracket(\sigma_p \oplus \sigma_{r_0}) \right) \right) \\ &= \int d\xi_n \left(pr_{\square}(\bar{c}^\pi)(\sigma_p, \xi_n) \cdot g \left(pvars_{\square}(\bar{c}^\pi)(\sigma_p, \xi_n), vals_{\square}(\bar{c}^\pi)(\sigma_p, \xi_n) \right) \right) \end{aligned}$$

where $\widehat{g}: \mathbb{R} \rightarrow \mathbb{R}$ is defined as $\widehat{g}(r) = g(\sigma_p[x \mapsto r], [\mu \mapsto r])$. Here the first and fifth equalities use the equations proven above, the second and fourth equalities use that $[\{\mu\} \rightarrow \mathbb{R}]$ is isomorphic to \mathbb{R} , and the third equality uses Eq. (17). This proves the desired equation.

Case $c \equiv (\text{if } b \{c'\} \text{ else } c'')$. In this case, since $fv(b) \subseteq \text{PVar}$ and $\bar{c}^\pi \equiv \text{if } b \{\bar{c}^{\pi'}\} \text{ else } \bar{c}^{\pi''}$, we have only two subcases:

- For all $\sigma_r \in \text{St}[\text{Var} \setminus \text{PVar}]$, $\llbracket c \rrbracket(\sigma_p \oplus \sigma_r) = \llbracket c' \rrbracket(\sigma_p \oplus \sigma_r)$ and $\llbracket \bar{c}^\pi \rrbracket(\sigma_p \oplus \sigma_r) = \llbracket \bar{c}^{\pi'} \rrbracket(\sigma_p \oplus \sigma_r)$.
- For all $\sigma_r \in \text{St}[\text{Var} \setminus \text{PVar}]$, $\llbracket c \rrbracket(\sigma_p \oplus \sigma_r) = \llbracket c'' \rrbracket(\sigma_p \oplus \sigma_r)$ and $\llbracket \bar{c}^\pi \rrbracket(\sigma_p \oplus \sigma_r) = \llbracket \bar{c}^{\pi''} \rrbracket(\sigma_p \oplus \sigma_r)$.

If the first subcase holds, we have

$$\begin{aligned} & \int d\xi_n \left(pr_{\square}(c)(\sigma_p, \xi_n) \cdot g \left(pvars_{\square}(c)(\sigma_p, \xi_n), vals_{\square}(c)(\sigma_p, \xi_n) \right) \right) \\ &= \int d\xi_n \left(pr_{\square}(c')(\sigma_p, \xi_n) \cdot g \left(pvars_{\square}(c')(\sigma_p, \xi_n), vals_{\square}(c')(\sigma_p, \xi_n) \right) \right) \\ &= \int d\xi_n \left(pr_{\square}(\bar{c}^{\pi'}) (\sigma_p, \xi_n) \cdot g \left(pvars_{\square}(\bar{c}^{\pi'}) (\sigma_p, \xi_n), vals_{\square}(\bar{c}^{\pi'}) (\sigma_p, \xi_n) \right) \right) \end{aligned}$$

$$= \int d\xi_n \left(\text{prs}_\square(\bar{c}^\pi)(\sigma_p, \xi_n) \cdot g \left(\text{pvars}_\square(\bar{c}^\pi)(\sigma_p, \xi_n), \text{vals}_\square(\bar{c}^\pi)(\sigma_p, \xi_n) \right) \right)$$

where the second equality is by IH on c' . If the second subcase holds, we obtain a similar equation by IH on c'' . Hence, the desired equation holds in all subcases.

Case $c \equiv (c'; c'')$. In this case, we obtain the following equation:

$$\begin{aligned} & \int d\xi_n \left(\text{prs}_\square(c'; c'')(\sigma_p, \xi_n) \cdot g \left(\text{pvars}_\square(c'; c'')(\sigma_p, \xi_n), \text{vals}_\square(c'; c'')(\sigma_p, \xi_n) \right) \right) \\ &= \int d\xi'_n \left(\text{prs}_\square(c')(\sigma_p, \xi'_n) \cdot \int d\xi''_n \left(\text{prs}_\square(c'')(\text{pvars}_\square(c')(\sigma_p, \xi'_n), \xi''_n) \cdot \mathbf{1}_{[\text{dom}(\xi'_n) \cap \text{dom}(\xi''_n) = \emptyset]} \right. \right. \\ & \quad \left. \left. \cdot g \left(\text{pvars}_\square(c'')(\text{pvars}_\square(c')(\sigma_p, \xi'_n), \xi''_n), \text{vals}_\square(c')(\sigma_p, \xi'_n) \oplus \text{vals}_\square(c'')(\text{pvars}_\square(c')(\sigma_p, \xi'_n), \xi''_n) \right) \right) \right) \\ &= \int d\xi'_n \left(\text{prs}_\square(c')(\sigma_p, \xi'_n) \cdot g' \left(\text{pvars}_\square(c')(\sigma_p, \xi'_n), \text{vals}_\square(c')(\sigma_p, \xi'_n) \right) \right) \end{aligned}$$

$$\text{where } g'(\widehat{\sigma}'_p, \widehat{\xi}'_n) \triangleq \int d\xi''_n \left(\text{prs}_\square(c'')(\widehat{\sigma}'_p, \xi''_n) \cdot \mathbf{1}_{[\text{dom}(\widehat{\xi}'_n) \cap \text{dom}(\xi''_n) = \emptyset]} \right.$$

$$\left. \cdot g \left(\text{pvars}_\square(c'')(\widehat{\sigma}'_p, \xi''_n), \widehat{\xi}'_n \oplus \text{vals}_\square(c'')(\widehat{\sigma}'_p, \xi''_n) \right) \right)$$

$$= \int d\xi'_n \left(\text{prs}_\square(\bar{c}^\pi)(\sigma_p, \xi'_n) \cdot g' \left(\text{pvars}_\square(\bar{c}^\pi)(\sigma_p, \xi'_n), \text{vals}_\square(\bar{c}^\pi)(\sigma_p, \xi'_n) \right) \right) \quad \dots \quad (*)$$

where the first equality is from Lemma C.9, the second equality uses $\text{dom}(\xi'_n) = \text{dom}(\text{vals}_\square(c')(\sigma_p, \xi'_n))$, and the third equality is by IH on c' . We now analyse $g'(\widehat{\sigma}'_p, \widehat{\xi}'_n)$ as follows:

$$g'(\widehat{\sigma}'_p, \widehat{\xi}'_n)$$

$$= \int d\xi''_n \left(\text{prs}_\square(c'')(\widehat{\sigma}'_p, \xi''_n) \cdot \mathbf{1}_{[\text{dom}(\widehat{\xi}'_n) \cap \text{dom}(\xi''_n) = \emptyset]} \cdot g \left(\text{pvars}_\square(c'')(\widehat{\sigma}'_p, \xi''_n), \widehat{\xi}'_n \oplus \text{vals}_\square(c'')(\widehat{\sigma}'_p, \xi''_n) \right) \right)$$

$$= \int d\xi''_n \left(\text{prs}_\square(c'')(\widehat{\sigma}'_p, \xi''_n) \cdot g'' \left(\text{pvars}_\square(c'')(\widehat{\sigma}'_p, \xi''_n), \text{vals}_\square(c'')(\widehat{\sigma}'_p, \xi''_n) \right) \right)$$

$$\text{where } g''(\widehat{\sigma}'_p, \widehat{\xi}'_n) \triangleq \mathbf{1}_{[\text{dom}(\widehat{\xi}'_n) \cap \text{dom}(\xi''_n) = \emptyset]} \cdot g \left(\widehat{\sigma}'_p, \widehat{\xi}'_n \oplus \widehat{\xi}'_n \right)$$

$$= \int d\xi''_n \left(\text{prs}_\square(\bar{c}''^\pi)(\widehat{\sigma}'_p, \xi''_n) \cdot g'' \left(\text{pvars}_\square(\bar{c}''^\pi)(\widehat{\sigma}'_p, \xi''_n), \text{vals}_\square(\bar{c}''^\pi)(\widehat{\sigma}'_p, \xi''_n) \right) \right)$$

$$= \int d\xi''_n \left(\text{prs}_\square(\bar{c}''^\pi)(\widehat{\sigma}'_p, \xi''_n) \cdot \mathbf{1}_{[\text{dom}(\widehat{\xi}'_n) \cap \text{dom}(\xi''_n) = \emptyset]} \cdot g \left(\text{pvars}_\square(\bar{c}''^\pi)(\widehat{\sigma}'_p, \xi''_n), \widehat{\xi}'_n \oplus \text{vals}_\square(\bar{c}''^\pi)(\widehat{\sigma}'_p, \xi''_n) \right) \right)$$

where the second and fourth equalities use $\text{dom}(\xi''_n) = \text{dom}(\text{vals}_\square(c'')(\widehat{\sigma}'_p, \xi''_n))$, and the third equality is by IH on c'' . Using this, we obtain the following equation for the main quantity (*):

$$\begin{aligned} (*) &= \int d\xi'_n \left(\text{prs}_\square(\bar{c}^\pi)(\sigma_p, \xi'_n) \cdot \int d\xi''_n \left(\text{prs}_\square(\bar{c}''^\pi)(\text{pvars}_\square(\bar{c}^\pi)(\sigma_p, \xi'_n), \xi''_n) \cdot \mathbf{1}_{[\text{dom}(\xi'_n) \cap \text{dom}(\xi''_n) = \emptyset]} \right. \right. \\ & \quad \left. \left. \cdot g \left(\text{pvars}_\square(\bar{c}''^\pi)(\text{pvars}_\square(\bar{c}^\pi)(\sigma_p, \xi'_n), \xi''_n), \right. \right. \\ & \quad \left. \left. \text{vals}_\square(\bar{c}^\pi)(\sigma_p, \xi'_n) \oplus \text{vals}_\square(\bar{c}''^\pi)(\text{pvars}_\square(\bar{c}^\pi)(\sigma_p, \xi'_n), \xi''_n) \right) \right) \right) \\ &= \int d\xi_n \left(\text{prs}_\square(\bar{c}^\pi; \bar{c}''^\pi)(\sigma_p, \xi_n) \cdot g \left(\text{pvars}_\square(\bar{c}^\pi; \bar{c}''^\pi)(\sigma_p, \xi_n), \text{vals}_\square(\bar{c}^\pi; \bar{c}''^\pi)(\sigma_p, \xi_n) \right) \right) \end{aligned}$$

where the first equality uses $\text{dom}(\xi'_n) = \text{dom}(\text{vals}_\square(c')(\sigma_p, \xi'_n))$, and the second equality is by Lemma C.9, as we did above. By $c'; c''^\pi \equiv \bar{c}^\pi; \bar{c}''^\pi$, we get the desired equation.

Case $c \equiv (\text{while } b \{c'\})$. In this case, $\bar{c}^\pi \equiv (\text{while } b \{\bar{c}'^\pi\})$. Without loss of generality, assume that g is a nonnegative function; we can prove the general case of g directly from the nonnegative case of g , by considering the nonnegative part and the negative part of g separately.

Consider the version of $\text{prs}_\square(-)$, $\text{pvars}_\square(-)$, and $\text{vals}_\square(-)$, where the parameter can be a state transformer $f : \text{St} \rightarrow \text{St}_\perp$, instead of a command. We denote the versions by $\text{prs}_\square(f)$, $\text{pvars}_\square(f)$, and $\text{vals}_\square(f)$. Define $\mathcal{T} \subseteq [\text{St} \rightarrow \text{St}_\perp]^2$ and $T : [\text{St} \rightarrow \text{St}_\perp]^2 \rightarrow [\text{St} \rightarrow \text{St}_\perp]^2$ by

$$(f, \bar{f}) \in \mathcal{T} \iff \int d\xi_n G_{g', \sigma'_p}(f)(\xi_n) = \int d\xi_n G_{g', \sigma'_p}(\bar{f})(\xi_n)$$

for all measurable $g' : \text{St}[\text{PVar}] \times \text{St}_\square[\text{Name}] \rightarrow \mathbb{R}_{\geq 0}$ and $\sigma'_p \in \text{St}[\text{PVar}]$,

$$T(f, \bar{f}) \triangleq (F(f), \bar{F}(\bar{f})),$$

where $G_{g', \sigma'_p}(f) \in \text{St}_\square[\text{Name}] \rightarrow \mathbb{R}_{\geq 0}$ and $F, \bar{F} : [\text{St} \rightarrow \text{St}_\perp] \rightarrow [\text{St} \rightarrow \text{St}_\perp]$ are defined by

$$G_{g', \sigma'_p}(f)(\xi_n) \triangleq \text{prs}_\square(f)(\sigma'_p, \xi_n) \cdot g'(\text{pvars}_\square(f)(\sigma'_p, \xi_n), \text{vals}_\square(f)(\sigma'_p, \xi_n)),$$

$$F(f)(\sigma) \triangleq \text{if } (\llbracket b \rrbracket \sigma = \text{true}) \text{ then } (f^\dagger \circ \llbracket c' \rrbracket)(\sigma) \text{ else } \sigma,$$

$$\bar{F}(\bar{f})(\sigma) \triangleq \text{if } (\llbracket b \rrbracket \sigma = \text{true}) \text{ then } (\bar{f}^\dagger \circ \llbracket \bar{c}'^\pi \rrbracket)(\sigma) \text{ else } \sigma.$$

Note that F and \bar{F} are the operators used in the semantics of the loops $\llbracket c \rrbracket$ and $\llbracket \bar{c}^\pi \rrbracket$, respectively. We will show that \mathcal{T} contains $(\lambda\sigma. \perp, \lambda\sigma. \perp)$, the operator T preserves \mathcal{T} , and \mathcal{T} is closed under taking the least upper bound of an increasing chain in $[\text{St} \rightarrow \text{St}_\perp]^2$, where the order on $[\text{St} \rightarrow \text{St}_\perp]^2$ is defined as: $(f_0, \bar{f}_0) \sqsubseteq (f_1, \bar{f}_1) \iff f_0 \sqsubseteq f_1 \wedge \bar{f}_0 \sqsubseteq \bar{f}_1$. These three conditions imply $(\llbracket c \rrbracket, \llbracket \bar{c}^\pi \rrbracket) \in \mathcal{T}$, which in turn proves the desired equation:

$$\begin{aligned} & \int d\xi_n \left(\text{prs}_\square(c)(\sigma_p, \xi_n) \cdot g(\text{pvars}_\square(c)(\sigma_p, \xi_n), \text{vals}_\square(c)(\sigma_p, \xi_n)) \right) \\ &= \int d\xi_n G_{g, \sigma_p}(\llbracket c \rrbracket)(\xi_n) \\ &= \int d\xi_n G_{g, \sigma_p}(\llbracket \bar{c}^\pi \rrbracket)(\xi_n) \\ &= \int d\xi_n \left(\text{prs}_\square(\bar{c}^\pi)(\sigma_p, \xi_n) \cdot g(\text{pvars}_\square(\bar{c}^\pi)(\sigma_p, \xi_n), \text{vals}_\square(\bar{c}^\pi)(\sigma_p, \xi_n)) \right), \end{aligned}$$

where the second equality follows from $(\llbracket c \rrbracket, \llbracket \bar{c}^\pi \rrbracket) \in \mathcal{T}$.

The first condition holds simply because $G_{g', \sigma'_p}(\lambda\sigma. \perp)(\xi_n) = 0$ for all g', σ'_p , and ξ_n . To show the second condition, pick $(f, \bar{f}) \in \mathcal{T}$. Our goal is to show $T(f, \bar{f}) \in \mathcal{T}$. We first replay our proof for the sequential-composition case on $(f^\dagger \circ \llbracket c' \rrbracket, \bar{f}^\dagger \circ \llbracket \bar{c}'^\pi \rrbracket)$, after viewing $f^\dagger \circ \llbracket c' \rrbracket$ and $\bar{f}^\dagger \circ \llbracket \bar{c}'^\pi \rrbracket$ as the sequential composition of c' and f , and of \bar{c}'^π and \bar{f} , respectively. This replay, then, gives the membership $(f^\dagger \circ \llbracket c' \rrbracket, \bar{f}^\dagger \circ \llbracket \bar{c}'^\pi \rrbracket) \in \mathcal{T}$. Next, we replay our proof for the if case on $(F(f), \bar{F}(\bar{f}))$, after viewing $f^\dagger \circ \llbracket c' \rrbracket$ and $\bar{f}^\dagger \circ \llbracket \bar{c}'^\pi \rrbracket$ as the true branches, and $\lambda\sigma. \sigma = \llbracket \text{skip} \rrbracket$ as the false branch. This replay implies the required $T(f, \bar{f}) = (F(f), \bar{F}(\bar{f})) \in \mathcal{T}$.

To show the third condition, consider an increasing sequence $\{(f_k, \bar{f}_k)\}_{k \in \mathbb{N}}$ in \mathcal{T} . Let $f_\infty \triangleq \bigsqcup_{k \in \mathbb{N}} f_k$ and $\bar{f}_\infty \triangleq \bigsqcup_{k \in \mathbb{N}} \bar{f}_k$. Consider a measurable $g' : \text{St}[\text{PVar}] \times \text{St}_\square[\text{Name}] \rightarrow \mathbb{R}_{\geq 0}$ and $\sigma'_p \in \text{St}[\text{PVar}]$. We should show that $\int d\xi_n G_{g', \sigma'_p}(f_\infty)(\xi_n) = \int d\xi_n G_{g', \sigma'_p}(\bar{f}_\infty)(\xi_n)$. Since $\{f_k\}_{k \in \mathbb{N}}$ is increasing, for any $\sigma \in \text{St}$, $f_k(\sigma) \in \text{St}$ implies that $f_{k'}(\sigma) = f_k(\sigma) \in \text{St}$ for all $k' \geq k$. Hence, $\{G_{g', \sigma'_p}(f_k)\}_{k \in \mathbb{N}}$ is a pointwise increasing sequence: for all $\xi_n \in \text{St}_\square[\text{Name}]$,

$$0 \leq G_{g', \sigma'_p}(f_k)(\xi_n) \leq G_{g', \sigma'_p}(f_{k+1})(\xi_n) \quad \text{for all } k \in \mathbb{N}.$$

Also, by the definition of f_∞ , for any $\sigma \in \text{St}$, there exists $K \in \mathbb{N}$ such that $f_\infty(\sigma) = f_K(\sigma)$; thus, $G_{g',\sigma'_p}(f_\infty)$ is the pointwise limit of $\{G_{g',\sigma'_p}(f_k)\}_{k \in \mathbb{N}}$: for all $\xi_n \in \text{St}_\square[\text{Name}]$,

$$G_{g',\sigma'_p}(f_\infty)(\xi_n) = \lim_{k \rightarrow \infty} G_{g',\sigma'_p}(f_k)(\xi_n).$$

Note that the corresponding results hold for $\overline{f_\infty}$ and $\overline{f_k}$. Using these results, we finally obtain the following as desired:

$$\begin{aligned} \int d\xi_n G_{g',\sigma'_p}(f_\infty)(\xi_n) &= \lim_{k \rightarrow \infty} \int d\xi_n G_{g',\sigma'_p}(f_k)(\xi_n) \\ &= \lim_{k \rightarrow \infty} \int d\xi_n G_{g',\sigma'_p}(\overline{f_k})(\xi_n) = \int d\xi_n G_{g',\sigma'_p}(\overline{f_\infty})(\xi_n). \end{aligned}$$

The first and third equalities follow from the monotone convergence theorem, applied to the above results. The second equality holds since $(f_k, \overline{f_k}) \in \mathcal{T}$. This completes the proof of the while case. \square

LEMMA C.8. *Let c be a command, $\sigma_0, \sigma_1 \in \text{St}$, and $r_0 \in \mathbb{R}$. Suppose that $\sigma_1(\text{like}) = \sigma_0(\text{like}) \cdot r_0$ and $\sigma_1|_V = \sigma_0|_V$ for $V \triangleq \text{Var} \setminus \{\text{like}\}$. If $\llbracket c \rrbracket \sigma_0 \in \text{St}$, then*

$$\llbracket c \rrbracket \sigma_1 \in \text{St}, \quad \llbracket c \rrbracket \sigma_1(\text{like}) = \llbracket c \rrbracket \sigma_0(\text{like}) \cdot r_0, \quad (\llbracket c \rrbracket \sigma_1)|_V = (\llbracket c \rrbracket \sigma_0)|_V.$$

PROOF. Let $V \triangleq \text{Var} \setminus \{\text{like}\}$. Pick an arbitrary command c . We prove the lemma by induction on the structure of c . Let $\sigma_0, \sigma_1 \in \text{St}$ and $r_0 \in \mathbb{R}$ such that $\llbracket c \rrbracket \sigma_0 \in \text{St}$, $\sigma_1(\text{like}) = \sigma_0(\text{like}) \cdot r_0$, and $\sigma_1|_V = \sigma_0|_V$. We should show that $\llbracket c \rrbracket \sigma_1 \in \text{St}$, $\llbracket c \rrbracket \sigma_1(\text{like}) = \llbracket c \rrbracket \sigma_0(\text{like}) \cdot r_0$, and $\llbracket c \rrbracket \sigma_1|_V = \llbracket c \rrbracket \sigma_0|_V$.

Case $c \equiv \text{skip}$. In this case, what we need to prove is identical to the assumption on $(\sigma_0, \sigma_1, r_0)$.

Case $c \equiv (x := e)$. By the semantics of the assignments, we have $\llbracket c \rrbracket \sigma_1 \in \text{St}$, and

$$\llbracket c \rrbracket \sigma_1(\text{like}) = \sigma_1(\text{like}) = \sigma_0(\text{like}) \cdot r_0 = \llbracket c \rrbracket \sigma_0(\text{like}) \cdot r_0.$$

The last requirement also holds since $\llbracket e \rrbracket \sigma_0 = \llbracket e \rrbracket \sigma_1$ and $\sigma_0|_V = \sigma_1|_V$.

Case $c \equiv (x := \text{sam}(n, d, \lambda y. e'))$. By the semantics of the sample commands, we have $\llbracket c \rrbracket \sigma_1 \in \text{St}$. Also, the assignments do not change the value of *like*, so that $\llbracket c \rrbracket \sigma_1(\text{like}) = \sigma_1(\text{like}) = \sigma_0(\text{like}) \cdot r_0 = \llbracket c \rrbracket \sigma_0(\text{like}) \cdot r_0$. It remains to show that $\llbracket c \rrbracket \sigma_0|_V = \llbracket c \rrbracket \sigma_1|_V$. Since $\sigma_0|_V = \sigma_1|_V$, we have $\llbracket n \rrbracket \sigma_0 = \llbracket n \rrbracket \sigma_1$ and $\llbracket d \rrbracket \sigma_0 = \llbracket d \rrbracket \sigma_1$. Let $\mu \triangleq \llbracket n \rrbracket \sigma_0$. Then, by the same reason, $\llbracket e'[\mu/y] \rrbracket \sigma_0 = \llbracket e'[\mu/y] \rrbracket \sigma_1$. Let $f \triangleq \llbracket d \rrbracket \sigma_0$ and $r \triangleq \llbracket e'[\mu/y] \rrbracket \sigma_0$. We prove the required equality as follows:

$$\begin{aligned} \llbracket c \rrbracket \sigma_0|_V &= \sigma_0[x \mapsto r, \text{val}_\mu \mapsto r, \text{pr}_\mu \mapsto f(\sigma_0(\mu)), \text{cnt}_\mu \mapsto \sigma_0(\text{cnt}_\mu) + 1]|_V \\ &= \sigma_1[x \mapsto r, \text{val}_\mu \mapsto r, \text{pr}_\mu \mapsto f(\sigma_0(\mu)), \text{cnt}_\mu \mapsto \sigma_0(\text{cnt}_\mu) + 1]|_V \\ &= \sigma_1[x \mapsto r, \text{val}_\mu \mapsto r, \text{pr}_\mu \mapsto f(\sigma_1(\mu)), \text{cnt}_\mu \mapsto \sigma_1(\text{cnt}_\mu) + 1]|_V \\ &= \llbracket c \rrbracket \sigma_1|_V. \end{aligned}$$

Case $c \equiv (\text{obs}(d, r))$. By the semantics of the observe commands, we have $\llbracket c \rrbracket \sigma_1 \in \text{St}$. Also, the observe commands do not change any variable except *like*. So, $\llbracket c \rrbracket \sigma_0|_V = \sigma_0|_V = \sigma_1|_V = \llbracket c \rrbracket \sigma_1|_V$. The remaining requirement for *like* can be proved as follows:

$$\llbracket c \rrbracket \sigma_1(\text{like}) = \sigma_1(\text{like}) \cdot \llbracket d \rrbracket \sigma_1(r) = \sigma_0(\text{like}) \cdot r_0 \cdot \llbracket d \rrbracket \sigma_1(r) = \sigma_0(\text{like}) \cdot r_0 \cdot \llbracket d \rrbracket \sigma_0(r) = \llbracket c \rrbracket \sigma_0(\text{like}) \cdot r_0.$$

Case $c \equiv (c'; c'')$. We have $\llbracket c' \rrbracket \sigma_0 \in \text{St}$ and $\llbracket c'' \rrbracket (\llbracket c' \rrbracket \sigma_0) \in \text{St}$. We apply induction hypothesis first to $(c', \sigma_0, \sigma_1, r_0)$, and then to $(c'', \llbracket c' \rrbracket \sigma_0, \llbracket c' \rrbracket \sigma_1, r_0)$. Then, we get the requirements of the lemma.

Case $c \equiv (\text{if } b \{c'\} \text{ else } \{c''\})$. We deal with the case that $\llbracket b \rrbracket_{\sigma_0} = \text{true}$. The other case of $\llbracket b \rrbracket_{\sigma_0} = \text{false}$ can be proved similarly. Since $\llbracket b \rrbracket_{\sigma_0} = \text{true}$, we have $\llbracket c' \rrbracket_{\sigma_0} = \llbracket c \rrbracket_{\sigma_0} \in \text{St}$. Thus, we can apply induction hypothesis to c' . If we do so, we get

$$\llbracket c' \rrbracket_{\sigma_1} \in \text{St}, \quad \llbracket c' \rrbracket_{\sigma_1}(\text{like}) = \llbracket c' \rrbracket_{\sigma_0}(\text{like}) \cdot r_0, \quad \text{and} \quad \llbracket c' \rrbracket_{\sigma_1}|_V = \llbracket c' \rrbracket_{\sigma_0}|_V.$$

This gives the desired conclusion because $\llbracket b \rrbracket_{\sigma_1} = \llbracket b \rrbracket_{\sigma_0} = \text{true}$ and so $\llbracket c \rrbracket_{\sigma_1} = \llbracket c' \rrbracket_{\sigma_1}$, and $\llbracket c \rrbracket_{\sigma_0} = \llbracket c' \rrbracket_{\sigma_0}$.

Case $c \equiv (\text{while } b \{c'\})$. Let F be the operator on $[\text{St} \rightarrow \text{St}_\perp]$ such that $\llbracket c \rrbracket$ is the least fixed point of F . Define a subset \mathcal{T} of $[\text{St} \rightarrow \text{St}_\perp]$ as follows: a function $f \in [\text{St} \rightarrow \text{St}_\perp]$ is in \mathcal{T} if and only if for all $\sigma'_0, \sigma'_1 \in \text{St}$ such that $\sigma'_1|_V = \sigma'_0|_V$ and $\sigma'_1(\text{like}) = \sigma'_0(\text{like}) \cdot r_0$, we have

$$f(\sigma'_0) \neq \perp \implies \left(f(\sigma'_1) \neq \perp \wedge f(\sigma'_1)|_V = f(\sigma'_0)|_V \wedge f(\sigma'_1)(\text{like}) = f(\sigma'_0)(\text{like}) \cdot r_0 \right).$$

The set \mathcal{T} contains the least function $\lambda \sigma. \perp$, and is closed under the least upper bound of any chain in $[\text{St} \rightarrow \text{St}_\perp]$. It is also closed under F . This F -closure follows essentially from our arguments for sequential composition, if command, and skip, and induction hypothesis on c' . What we have shown for \mathcal{T} implies that \mathcal{T} contains the least fixed point of F , which gives the desired property for c . \square

LEMMA C.9. *Let c', c'' be commands and $g : \text{St}[\text{PVar}] \times \text{St}_\square[\text{Name}] \rightarrow \mathbb{R}$ be a measurable function. Then, for any $\sigma_p \in \text{St}[\text{PVar}]$,*

$$\begin{aligned} & \int d\xi_n \left(\text{pr}_{\text{St}_\square}(c'; c'')(\sigma_p, \xi_n) \cdot g \left(\text{pvars}_{\text{St}_\square}(c'; c'')(\sigma_p, \xi_n), \text{vals}_{\text{St}_\square}(c'; c'')(\sigma_p, \xi_n) \right) \right) \\ &= \int d\xi'_n \left(\text{pr}_{\text{St}_\square}(c')(\sigma_p, \xi'_n) \cdot \int d\xi''_n \left(\text{pr}_{\text{St}_\square}(c'')(\text{pvars}_{\text{St}_\square}(c')(\sigma_p, \xi'_n), \xi''_n) \cdot \mathbf{1}_{\text{dom}(\xi'_n) \cap \text{dom}(\xi''_n) = \emptyset} \right. \right. \\ & \quad \left. \left. \cdot g \left(\text{pvars}_{\text{St}_\square}(c'')(\text{pvars}_{\text{St}_\square}(c')(\sigma_p, \xi'_n), \xi''_n), \text{vals}_{\text{St}_\square}(c')(\sigma_p, \xi'_n) \oplus \text{vals}_{\text{St}_\square}(c'')(\text{pvars}_{\text{St}_\square}(c')(\sigma_p, \xi'_n), \xi''_n) \right) \right) \right). \end{aligned}$$

PROOF. Let c', c'' be commands, $g : \text{St}[\text{PVar}] \times \text{St}_\square[\text{Name}] \rightarrow \mathbb{R}$ a measurable function, and $\sigma_p \in \text{St}[\text{PVar}]$. In this proof, each equation involving integrals means (otherwise noted) that one side of the equation is defined if and only if the other side is defined, and when both sides are defined, they are the same.

First, to convert a single-integral on ξ_n to a double-integral on ξ'_n and ξ''_n as in the desired equation, we show the following claim: for any measurable $f : \text{St}_\square[\text{Name}] \rightarrow \text{St}$ and $h : \text{St}_\square[\text{Name}] \rightarrow \mathbb{R}$ such that $f(\xi_n)|_{\text{dom}(\xi_n)} = \xi_n$ for all $\xi_n \in \text{St}_\square[\text{Name}]$, we have

$$\begin{aligned} & \int d\xi_n \left(\mathbf{1}_{[\text{used}(c'; c''), f(\xi_n), \xi_n]} \cdot h(\xi_n) \right) \\ &= \sum_{K \subseteq \text{Name}} \int_{[K \rightarrow \mathbb{R}]} d\xi_n \left(\mathbf{1}_{[\text{used}(c'; c''), f(\xi_n), \xi_n]} \cdot h(\xi_n) \right) \\ &= \sum_{K \subseteq \text{Name}} \int_{[K \rightarrow \mathbb{R}]} d\xi_n \left(\mathbf{1}_{[\text{used}(c'; c''), f(\xi_n), \xi_n]} \cdot h(\xi_n) \cdot \sum_{L \subseteq K} \mathbf{1}_{[\text{used}(c', f(\xi_n), \xi_n|_L]} \right) \\ &= \sum_{K \subseteq \text{Name}} \sum_{L \subseteq K} \int_{[K \rightarrow \mathbb{R}]} d\xi_n \left(\mathbf{1}_{[\text{used}(c', f(\xi_n), \xi_n|_L]} \cdot \mathbf{1}_{[\text{used}(c'; c''), f(\xi_n), \xi_n]} \cdot h(\xi_n) \right) \\ &= \sum_{K \subseteq \text{Name}} \sum_{L \subseteq K} \int_{[L \rightarrow \mathbb{R}]} d\xi'_n \int_{[K \setminus L \rightarrow \mathbb{R}]} d\xi''_n \left(\mathbf{1}_{[\text{used}(c', f(\xi'_n \oplus \xi''_n), \xi'_n)]} \cdot \mathbf{1}_{[\text{used}(c'; c''), f(\xi'_n \oplus \xi''_n), \xi'_n \oplus \xi''_n]} \cdot h(\xi'_n \oplus \xi''_n) \right) \\ &= \sum_{L' \subseteq \text{Name}} \sum_{\substack{M' \subseteq \text{Name} \\ L' \cap M' = \emptyset}} \int_{[L' \rightarrow \mathbb{R}]} d\xi'_n \int_{[M' \rightarrow \mathbb{R}]} d\xi''_n \left(\mathbf{1}_{[\text{used}(c', f(\xi'_n \oplus \xi''_n), \xi'_n)]} \cdot \mathbf{1}_{[\text{used}(c'; c''), f(\xi'_n \oplus \xi''_n), \xi'_n \oplus \xi''_n]} \cdot h(\xi'_n \oplus \xi''_n) \right) \end{aligned}$$

$$\begin{aligned}
&= \sum_{L' \subseteq \text{Name}} \int_{[L' \rightarrow \mathbb{R}]} d\xi'_n \left(\sum_{M' \subseteq \text{Name}} \int_{[M' \rightarrow \mathbb{R}]} d\xi''_n \left(\mathbf{1}_{[\text{dom}(\xi'_n) \cap \text{dom}(\xi''_n) = \emptyset]} \cdot \mathbf{1}_{[\text{used}(c', f(\xi'_n \oplus \xi''_n), \xi'_n)]} \right. \right. \\
&\quad \left. \left. \cdot \mathbf{1}_{[\text{used}(c', c'', f(\xi'_n \oplus \xi''_n), \xi'_n \oplus \xi''_n)]} \cdot h(\xi'_n \oplus \xi''_n) \right) \right) \\
&= \int d\xi'_n \int d\xi''_n \left(\mathbf{1}_{[\text{dom}(\xi'_n) \cap \text{dom}(\xi''_n) = \emptyset]} \cdot \mathbf{1}_{[\text{used}(c', f(\xi'_n \oplus \xi''_n), \xi'_n)]} \cdot \mathbf{1}_{[\text{used}(c', c'', f(\xi'_n \oplus \xi''_n), \xi'_n \oplus \xi''_n)]} \cdot h(\xi'_n \oplus \xi''_n) \right).
\end{aligned}$$

The first equality uses the definition of $\text{St}_\square[\text{Name}]$ and its measure. The second equality uses that $\text{used}(c'; c'', \sigma, \xi_n)$ implies a unique existence of $L \subseteq \text{dom}(\xi_n)$ such that $\text{used}(c', \sigma, \xi_n|L)$; we already showed the existence of such L in the proof of Lemma C.6 (for the sequential composition case), and the uniqueness follows from the definition of used . The third equality uses that K is finite, and the fourth equality uses that $[K \rightarrow \mathbb{R}]$ is isomorphic to $[L \rightarrow \mathbb{R}] \times [K \setminus L \rightarrow \mathbb{R}]$ for any $L \subseteq K$. The fifth equality holds uses that $\{(L, K \setminus L) \mid K \subseteq \text{Name}, L \subseteq K\} = \{(L', M') \mid L', M' \subseteq \text{Name}, L' \cap M' = \emptyset\}$. The sixth equality uses that Name is finite, $\text{dom}(\xi'_n) = L'$, and $\text{dom}(\xi''_n) = M'$. The seventh equality uses the definition of $\text{St}_\square[\text{Name}]$ and its measure.

Second, to decompose $\text{prs}_\square(c'; c'')$, $\text{pvars}_\square(c'; c'')$, and $\text{vals}_\square(c'; c'')$ as in the desired equation, we show the following claim. Suppose that $\sigma \in \text{St}$ and $\xi'_n, \xi''_n \in \text{St}_\square[\text{Name}]$ with $\text{dom}(\xi'_n) \cap \text{dom}(\xi''_n) = \emptyset$ satisfy $\text{used}(c'; c'', \sigma, \xi'_n \oplus \xi''_n)$ and $\text{used}(c', \sigma, \xi'_n)$. Then, we first get

$$\begin{aligned}
\text{prs}_\square(c'; c'')(\sigma|_{\text{PVar}}, \xi'_n \oplus \xi''_n) &= \llbracket c'; c'' \rrbracket \sigma(\text{like}) \cdot \prod_{\mu \in \text{dom}(\xi_n \oplus \xi''_n)} \llbracket c'; c'' \rrbracket \sigma(\text{pr}_\mu) \\
&= \llbracket c'' \rrbracket (\llbracket c' \rrbracket \sigma)(\text{like}) \\
&\quad \cdot \prod_{\mu \in \text{dom}(\xi'_n)} \llbracket c'' \rrbracket (\llbracket c' \rrbracket \sigma)(\text{pr}_\mu) \cdot \prod_{\mu \in \text{dom}(\xi''_n)} \llbracket c'' \rrbracket (\llbracket c' \rrbracket \sigma)(\text{pr}_\mu) \\
&= \llbracket c' \rrbracket \sigma(\text{like}) \cdot \llbracket c'' \rrbracket ((\llbracket c' \rrbracket \sigma)[\text{like} \mapsto 1])(\text{like}) \\
&\quad \cdot \prod_{\mu \in \text{dom}(\xi'_n)} \llbracket c' \rrbracket \sigma(\text{pr}_\mu) \cdot \prod_{\mu \in \text{dom}(\xi''_n)} \llbracket c'' \rrbracket ((\llbracket c' \rrbracket \sigma)[\text{like} \mapsto 1])(\text{pr}_\mu) \\
&= \text{prs}_\square(c')(\sigma|_{\text{PVar}}, \xi'_n) \cdot \text{prs}_\square(c'')((\llbracket c' \rrbracket \sigma)[\text{like} \mapsto 1]|_{\text{PVar}}, \xi''_n) \\
&= \text{prs}_\square(c')(\sigma|_{\text{PVar}}, \xi'_n) \cdot \text{prs}_\square(c'')(p\text{vars}_\square(c')(\sigma|_{\text{PVar}}, \xi'_n), \xi''_n).
\end{aligned}$$

Here is the proof of each equality.

- The first equality uses $\text{used}(c'; c'', \sigma, \xi'_n \oplus \xi''_n)$.
- The second equality uses $\text{noerr}(c'; c'', \sigma)$, which comes from $\text{used}(c'; c'', \sigma, -)$.
- The third equality comes from Lemma C.8, Lemma C.6-(2), and Lemma C.6-(3). The two applications of Lemma C.6 are valid since $\text{used}_-(c'', \llbracket c' \rrbracket \sigma, \xi''_n)$ and $\text{dom}(\xi'_n) \cap \text{dom}(\xi''_n) = \emptyset$, where the first predicate follows from $\text{used}(c'; c'', \sigma, \xi'_n \oplus \xi''_n)$ and $\text{used}(c', \sigma, \xi'_n)$ by the claim in the proof of Lemma C.6 (for the sequential composition case).
- The fourth equality uses that $\text{used}(c', \sigma, \xi'_n)$ and $\text{used}(c'', (\llbracket c' \rrbracket \sigma)[\text{like} \mapsto 1], \xi''_n)$, where the second predicate follows from $\text{used}_-(c'', \llbracket c' \rrbracket \sigma, \xi''_n)$ and Lemma C.7.
- The fifth equality uses $(\llbracket c' \rrbracket \sigma)[\text{like} \mapsto 1]|_{\text{PVar}} = (\llbracket c' \rrbracket \sigma)|_{\text{PVar}} = p\text{vars}_\square(c')(\sigma|_{\text{PVar}}, \xi'_n)$, where the second part of the equation comes from $\text{used}(c', \sigma, \xi'_n)$.

By the same argument so far (except that pr_μ and \times are replaced by val_μ and \oplus), we next get

$$\text{vals}_\square(c'; c'')(\sigma|_{\text{PVar}}, \xi'_n \oplus \xi''_n) = \text{vals}_\square(c')(\sigma|_{\text{PVar}}, \xi'_n) \oplus \text{vals}_\square(c'')(p\text{vars}_\square(c')(\sigma|_{\text{PVar}}, \xi'_n), \xi''_n).$$

By a similar argument, we lastly get

$$\begin{aligned}
p\text{vars}_\square(c'; c'')(\sigma|_{\text{PVar}}, \xi'_n \oplus \xi''_n) &= \llbracket c'; c'' \rrbracket \sigma|_{\text{PVar}} \\
&= \llbracket c'' \rrbracket (\llbracket c' \rrbracket \sigma)|_{\text{PVar}} \\
&= \llbracket c'' \rrbracket ((\llbracket c' \rrbracket \sigma)[\text{like} \mapsto 1])|_{\text{PVar}}
\end{aligned}$$

$$\begin{aligned}
&= pvars_{\square}(c'')((\llbracket c' \rrbracket \sigma)[like \mapsto 1])|_{PVar, \xi''_n} \\
&= pvars_{\square}(c'')(pvars_{\square}(c')(\sigma)|_{PVar, \xi'_n, \xi''_n}).
\end{aligned}$$

Here is the proof of each equality.

- The first equality uses $used(c'; c'', \sigma, \xi'_n \oplus \xi''_n)$.
- The second equality uses $noerr(c'; c'', \sigma)$ (shown above).
- The third equality uses $used_{-}(c'', \llbracket c' \rrbracket \sigma, \xi''_n)$ (shown above) and Lemma C.6-(3).
- The fourth equality uses $used(c'', (\llbracket c' \rrbracket \sigma)[like \mapsto 1], \xi''_n)$ (shown above).
- The fifth equality uses $(\llbracket c' \rrbracket \sigma)[like \mapsto 1]|_{PVar} = pvars_{\square}(c')(\sigma)|_{PVar, \xi'_n}$ (shown above).

Third, to remove some indicator terms that will appear in our derivation, we show the next claim: for any $\sigma \in St$ and $\xi'_n, \xi''_n \in St_{\square}[Name]$ with $\text{dom}(\xi'_n) \cap \text{dom}(\xi''_n) = \emptyset$, $used_{-}(c', \sigma, \xi'_n)$ and $used_{-}(c'', \llbracket c' \rrbracket \sigma, \xi''_n)$ imply $used_{-}(c'; c'', \sigma, \xi'_n \oplus \xi''_n)$. Assume the premise. Then, we have

$$\begin{aligned}
&\llbracket c' \rrbracket \sigma \in St \wedge \xi'_n = \sigma|_{\text{dom}(\xi'_n)} \\
&\quad \wedge (\forall \mu \in Name. \llbracket c' \rrbracket \sigma(cnt_{\mu}) - \sigma(cnt_{\mu}) \leq 1) \\
&\quad \wedge \text{dom}(\xi'_n) = \{\mu \in Name \mid \llbracket c' \rrbracket \sigma'(cnt_{\mu}) - \sigma(cnt_{\mu}) = 1\}, \\
&\llbracket c'' \rrbracket (\llbracket c' \rrbracket \sigma) \in St \wedge \xi''_n = (\llbracket c' \rrbracket \sigma)|_{\text{dom}(\xi''_n)} \\
&\quad \wedge (\forall \mu \in Name. \llbracket c'' \rrbracket (\llbracket c' \rrbracket \sigma)(cnt_{\mu}) - \llbracket c' \rrbracket \sigma(cnt_{\mu}) \leq 1) \\
&\quad \wedge \text{dom}(\xi''_n) = \{\mu \in Name \mid \llbracket c'' \rrbracket (\llbracket c' \rrbracket \sigma)(cnt_{\mu}) - \llbracket c' \rrbracket \sigma(cnt_{\mu}) = 1\}.
\end{aligned}$$

We should show

$$\begin{aligned}
&\llbracket c'; c'' \rrbracket \sigma \in St \wedge \xi'_n \oplus \xi''_n = \sigma|_{\text{dom}(\xi'_n \oplus \xi''_n)} \\
&\quad \wedge (\forall \mu \in Name. \llbracket c'; c'' \rrbracket \sigma(cnt_{\mu}) - \sigma(cnt_{\mu}) \leq 1) \\
&\quad \wedge \text{dom}(\xi'_n \oplus \xi''_n) = \{\mu \in Name \mid \llbracket c'; c'' \rrbracket \sigma'(cnt_{\mu}) - \sigma(cnt_{\mu}) = 1\}.
\end{aligned}$$

We obtain the four clauses as follows. The first clause follows from $\llbracket c'' \rrbracket (\llbracket c' \rrbracket \sigma) \in St$. The second clause comes from $\xi'_n = \sigma|_{\text{dom}(\xi'_n)}$ and $\xi''_n = (\llbracket c' \rrbracket \sigma)|_{\text{dom}(\xi''_n)} = \sigma|_{\text{dom}(\xi''_n)}$, where the last equality comes from Lemma C.4. The third and fourth clauses hold by the following:

$$\begin{aligned}
&\llbracket c'; c'' \rrbracket \sigma(cnt_{\mu}) - \sigma(cnt_{\mu}) \\
&= \llbracket c'' \rrbracket (\llbracket c' \rrbracket \sigma)(cnt_{\mu}) - \sigma(cnt_{\mu}) \\
&= \begin{cases} \llbracket c' \rrbracket \sigma(cnt_{\mu}) - \sigma(cnt_{\mu}) = 1 & \text{if } \mu \in \text{dom}(\xi'_n) \\ \llbracket c'' \rrbracket (\llbracket c' \rrbracket \sigma)(cnt_{\mu}) - \llbracket c' \rrbracket \sigma(cnt_{\mu}) = 1 & \text{if } \mu \in \text{dom}(\xi''_n) \\ \sigma(cnt_{\mu}) - \sigma(cnt_{\mu}) = 0 & \text{if } \mu \in Name \setminus (\text{dom}(\xi'_n) \cup \text{dom}(\xi''_n)). \end{cases}
\end{aligned}$$

- The first case uses Lemma C.6-(2) (applied to $used_{-}(c'', \llbracket c' \rrbracket \sigma, \xi''_n)$ and $\text{dom}(\xi'_n) \cap \text{dom}(\xi''_n) = \emptyset$) and the fourth clause of $used_{-}(c', \sigma, \xi'_n)$.
- The second case uses Lemma C.6-(2) (applied to $used_{-}(c', \sigma, \xi'_n)$ and $\text{dom}(\xi'_n) \cap \text{dom}(\xi''_n) = \emptyset$) and the fourth clause of $used_{-}(c'', \llbracket c' \rrbracket \sigma, \xi''_n)$.
- The third case uses Lemma C.6-(2) (applied to $used_{-}(c'', \llbracket c' \rrbracket \sigma, \xi''_n)$ and $used_{-}(c', \sigma, \xi'_n)$).

Finally, we put the three results together. Define $f : St_{\square}[Name] \rightarrow St$ as $f(\xi_n) \triangleq \sigma_p \oplus \xi_n \oplus (\lambda v \in Var \setminus (PVar \cup \text{dom}(\xi_n)). 1)$. Then, we obtain the desired equation as follows:

$$\begin{aligned}
&\int d\xi_n \left(prs_{\square}(c'; c'')(\sigma_p, \xi_n) \cdot g \left(pvars_{\square}(c'; c'')(\sigma_p, \xi_n), vals_{\square}(c'; c'')(\sigma_p, \xi_n) \right) \right) \\
&= \int d\xi_n \left(\mathbf{1}_{[used(c'; c'', f(\xi_n), \xi_n)]} \cdot prs_{\square}(c'; c'')(\sigma_p, \xi_n) \cdot g \left(pvars_{\square}(c'; c'')(\sigma_p, \xi_n), vals_{\square}(c'; c'')(\sigma_p, \xi_n) \right) \right)
\end{aligned}$$

$$\begin{aligned}
&= \int d\xi'_n \int d\xi''_n \left(\mathbf{1}_{[\text{dom}(\xi'_n) \cap \text{dom}(\xi''_n) = \emptyset]} \cdot \mathbf{1}_{[\text{used}(c', f(\xi'_n \oplus \xi''_n), \xi'_n)]} \cdot \mathbf{1}_{[\text{used}(c', c'', f(\xi'_n \oplus \xi''_n), \xi'_n \oplus \xi''_n)]} \right. \\
&\quad \left. \cdot \text{prs}_{\square}(c'; c'')(\sigma_p, \xi'_n \oplus \xi''_n) \cdot g\left(\text{pvars}_{\square}(c'; c'')(\sigma_p, \xi'_n \oplus \xi''_n), \text{vals}_{\square}(c'; c'')(\sigma_p, \xi'_n \oplus \xi''_n)\right) \right) \\
&= \int d\xi'_n \int d\xi''_n \left(\mathbf{1}_{[\text{dom}(\xi'_n) \cap \text{dom}(\xi''_n) = \emptyset]} \cdot \mathbf{1}_{[\text{used}(c', f(\xi'_n \oplus \xi''_n), \xi'_n)]} \cdot \mathbf{1}_{[\text{used}(c', c'', f(\xi'_n \oplus \xi''_n), \xi'_n \oplus \xi''_n)]} \right. \\
&\quad \left. \cdot \text{prs}_{\square}(c')(\sigma_p, \xi'_n) \cdot \text{prs}_{\square}(c'')(\text{pvars}_{\square}(c')(\sigma_p, \xi'_n), \xi''_n) \right. \\
&\quad \left. \cdot g\left(\text{pvars}_{\square}(c'')(\text{pvars}_{\square}(c')(\sigma_p, \xi'_n), \xi''_n), \text{vals}_{\square}(c')(\sigma_p, \xi'_n) \oplus \text{vals}_{\square}(c'')(\text{pvars}_{\square}(c')(\sigma_p, \xi'_n), \xi''_n)\right) \right) \\
&= \int d\xi'_n \int d\xi''_n \left(\mathbf{1}_{[\text{dom}(\xi'_n) \cap \text{dom}(\xi''_n) = \emptyset]} \cdot \text{prs}_{\square}(c')(\sigma_p, \xi'_n) \cdot \text{prs}_{\square}(c'')(\text{pvars}_{\square}(c')(\sigma_p, \xi'_n), \xi''_n) \right. \\
&\quad \left. \cdot g\left(\text{pvars}_{\square}(c'')(\text{pvars}_{\square}(c')(\sigma_p, \xi'_n), \xi''_n), \text{vals}_{\square}(c')(\sigma_p, \xi'_n) \oplus \text{vals}_{\square}(c'')(\text{pvars}_{\square}(c')(\sigma_p, \xi'_n), \xi''_n)\right) \right).
\end{aligned}$$

The first equality uses that $\text{prs}_{\square}(c'; c'')(\sigma_p, \xi_n) \neq 0$ implies $\mathbf{1}_{[\text{used}(c', c'', f(\xi_n), \xi_n)]} = 1$:

- Since $\text{prs}_{\square}(c'; c'')(\sigma_p, \xi_n) \neq 0$, there is $\sigma_r \in \text{St}[\text{Var} \setminus (\text{PVar} \cup \text{dom}(\xi_n))]$ such that $\text{used}(c'; c'', \sigma_p \oplus \xi_n \oplus \sigma_r, \xi_n)$. Note $(\sigma_p \oplus \xi_n \oplus \sigma_r)|_V = f(\xi_n)|_V$ for $V = \text{PVar} \cup \text{dom}(\xi_n) \cup \{\text{like}\}$. From these and Lemma C.7, we get $\text{used}(c'; c'', f(\xi_n), \xi_n)$.

The second and third equalities use the first and second results we proved above, respectively. The fourth equality uses the next claim: $\text{prs}_{\square}(c')(\sigma_p, \xi'_n) \cdot \text{prs}_{\square}(c'')(\text{pvars}_{\square}(c')(\sigma_p, \xi'_n), \xi''_n) \neq 0$ and $\text{dom}(\xi'_n) \cap \text{dom}(\xi''_n) = \emptyset$ imply $\mathbf{1}_{[\text{used}(c', f(\xi'_n \oplus \xi''_n), \xi'_n)]} \cdot \mathbf{1}_{[\text{used}(c', c'', f(\xi'_n \oplus \xi''_n), \xi'_n \oplus \xi''_n)]} = 1$. We prove the claim using the third result we proved above:

- Assume the premise. From $\text{prs}_{\square}(c')(\sigma_p, \xi'_n) \neq 0$, there is $\sigma'_r \in \text{St}[\text{Var} \setminus (\text{PVar} \cup \text{dom}(\xi'_n))]$ such that $\text{used}(c', \sigma_p \oplus \xi'_n \oplus \sigma'_r, \xi'_n)$. Note $(\sigma_p \oplus \xi'_n \oplus \sigma'_r)|_{V'} = f(\xi'_n \oplus \xi''_n)|_{V'}$ for $V' = \text{PVar} \cup \text{dom}(\xi'_n) \cup \{\text{like}\}$. From these and Lemma C.7, we get $\text{used}(c', f(\xi'_n \oplus \xi''_n), \xi'_n)$ as desired.
- From $\text{prs}_{\square}(c'')(\text{pvars}_{\square}(c')(\sigma_p, \xi'_n), \xi''_n) \neq 0$, there is $\sigma''_r \in \text{St}[\text{Var} \setminus (\text{PVar} \cup \text{dom}(\xi''_n))]$ such that

$$\text{used}(c'', \text{pvars}_{\square}(c')(\sigma_p, \xi'_n) \oplus \xi''_n \oplus \sigma''_r, \xi''_n).$$

Since $\text{used}(c', f(\xi'_n \oplus \xi''_n), \xi'_n)$, we have $\llbracket c' \rrbracket(f(\xi'_n \oplus \xi''_n)) \in \text{St}$ and $\text{pvars}_{\square}(c')(\sigma_p, \xi'_n) = \text{pvars}_{\square}(c')(f(\xi'_n \oplus \xi''_n)|_{\text{PVar}}, \xi'_n) = \llbracket c' \rrbracket(f(\xi'_n \oplus \xi''_n))|_{\text{PVar}}$; also, by Lemma C.4, $\xi''_n = \llbracket c' \rrbracket(f(\xi'_n \oplus \xi''_n))|_{\text{dom}(\xi''_n)}$. Thus,

$$(\text{pvars}_{\square}(c')(\sigma_p, \xi'_n) \oplus \xi''_n \oplus \sigma''_r)|_{V''} = \llbracket c \rrbracket(f(\xi'_n \oplus \xi''_n))|_{V''}$$

for $V'' = \text{PVar} \cup \text{dom}(\xi''_n)$. From these and Lemma C.7, we get $\text{used}_-(c'', \llbracket c' \rrbracket(f(\xi'_n \oplus \xi''_n)), \xi''_n)$.

- From $\text{dom}(\xi'_n) \cap \text{dom}(\xi''_n) = \emptyset$, $\text{used}_-(c', f(\xi'_n \oplus \xi''_n), \xi'_n)$, and $\text{used}_-(c'', \llbracket c' \rrbracket(f(\xi'_n \oplus \xi''_n)), \xi''_n)$, we can apply the third result proved above, and get $\text{used}_-(c'; c'', f(\xi'_n \oplus \xi''_n), \xi'_n \oplus \xi''_n)$. Since $f(\xi'_n \oplus \xi''_n)(\text{like}) = 1$, we get $\text{used}(c'; c'', f(\xi'_n \oplus \xi''_n), \xi'_n \oplus \xi''_n)$ as desired.

This completes the proof. \square

D DEFERRED RESULTS IN §4.2

D.1 Deferred Statements and Their Proofs

LEMMA D.1. *Let c be a command and π be a simple reparameterisation plan. Suppose that for all $n \in \text{NameEx}$, $d, d' \in \text{DistEx}$, and $(\lambda y.e) \in \text{LamEx}$ such that $\pi(n, d, \lambda y.e) = (d', _)$, we have*

$$d' = \text{dist}_{\mathbb{N}}(r'_1, r'_2) \quad \text{for some } r'_1, r'_2 \in \mathbb{R}. \quad (18)$$

Further, assume that for all $\sigma_n \in \text{St}[\text{Name}]$, the function

$$\sigma_{\theta} \in \text{St}[\theta] \mapsto p_{\bar{c}^{\pi}, \sigma_{\theta}}^{(rv(\pi))}(\sigma_n) \quad (19)$$

is continuous. Then, for all $\sigma_\theta \in \text{St}[\theta]$ and $\sigma_n \in \text{St}[\text{Name}]$,

$$\nabla_\theta p_{\bar{c}^\pi, \sigma_\theta}^{(rv(\pi))}(\sigma_n) = 0. \quad (20)$$

PROOF. Consider c and π that satisfies the given conditions. Fix $\sigma_n \in \text{St}[\text{Name}]$. Let $f : \text{St}[\theta] \rightarrow \mathbb{R}$ be the function in Eq. (19). Suppose that Eq. (20) does not hold. Then, f is not a constant function.

On the one hand, since f is continuous (by assumption) and not constant, the image of f over its domain (i.e., $f(\text{St}[\theta]) \subseteq \mathbb{R}$) is an uncountable set. This can be shown as follows: since the image of a connected set over a continuous function is connected, $f(\text{St}[\theta])$ is a connected set in \mathbb{R} ; since f is not constant, $f(\text{St}[\theta])$ contains at least two points; since $f(\text{St}[\theta])$ is connected, it should contain a non-empty interval, so it should be an uncountable set.

On the other hand, since π is simple and satisfies Eq. (18), and since c has only finitely many sample commands, $f(\text{St}[\theta])$ is a finite set. So this contradicts to that $f(\text{St}[\theta])$ is an uncountable set. Hence, f should satisfy Eq. (20). \square

THEOREM D.2. Let $f : \mathbb{R} \times \mathbb{R}^n \rightarrow \mathbb{R}$ be a measurable function that satisfies the next three conditions:

- For all $x \in \mathbb{R}^n$, $f(-, x) : \mathbb{R} \rightarrow \mathbb{R}$ is differentiable.
- For all $\theta \in \mathbb{R}$, $\int_{\mathbb{R}^n} f(\theta, x) dx$ is finite.
- For all $\theta \in \mathbb{R}$, there is an open $U \subseteq \mathbb{R}$ such that $\theta \in U$ and $\int_{\mathbb{R}^n} \text{Lip}(f(-, x)|_U) dx$ is finite.

Here $\text{Lip}(g)$ for a function $g : V \rightarrow \mathbb{R}$ with $V \subseteq \mathbb{R}$ denotes the smallest Lipschitz constant of g :

$$\text{Lip}(g) \triangleq \sup_{r, r' \in V, r \neq r'} \frac{|g(r') - g(r)|}{|r' - r|} \in \mathbb{R} \cup \{\infty\}.$$

Then, for all $\theta \in \mathbb{R}$, both sides of the following are well-defined and equal:

$$\nabla_\theta \int_{\mathbb{R}^n} f(\theta, x) dx = \int_{\mathbb{R}^n} \nabla_\theta f(\theta, x) dx$$

where ∇_θ denotes the partial differentiation operator with respect to θ .

PROOF. This theorem follows from Theorem E.2 due to the following: the first condition of this theorem implies the first and second conditions of Theorem E.2 (as differentiability implies continuity); the second and third conditions of this theorem are identical to the third and fourth conditions of Theorem E.2; and the conclusion of this theorem is identical to that of Theorem E.2. \square

D.2 Proof of Theorem 4.5

The proof of Theorem 4.5 relies on the following two lemmas, which are proven in §D.3. The first lemma states that if a command contains no observe commands, then its (full) density function can be decomposed into its partial density functions over S and $\text{Name} \setminus S$ for any $S \subseteq \text{Name}$. The second lemma states that if π is simple and c uses only $\lambda y.y$ as the third argument of its sample commands, then the partial density function of \bar{c}^π over non-transformed random variables (i.e., variables in $\text{Name} \setminus rv(\pi)$) is connected to that of c via the value function of \bar{c}^π .

LEMMA D.3. Let c be a command. If c does not contain observe commands, then, for all $S \subseteq \text{Name}$, $\sigma_\theta \in \text{St}[\theta]$, and $\sigma_n \in \text{St}[\text{Name}]$,

$$p_{c, \sigma_\theta}(\sigma_n) = p_{c, \sigma_\theta}^{(S)}(\sigma_n) \cdot p_{c, \sigma_\theta}^{(\text{Name} \setminus S)}(\sigma_n).$$

LEMMA D.4. Let c be a command and π be a reparameterisation plan. Suppose that every sample command in c has $\lambda y.y$ as its third argument. Then, for all $\sigma_\theta \in \text{St}[\theta]$ and $\sigma_n \in \text{St}[\text{Name}]$, if $p_{\bar{c}^\pi, \sigma_\theta}(\sigma_n) > 0$, then

$$p_{\bar{c}^\pi, \sigma_\theta}^{(\text{Name} \setminus rv(\pi))}(\sigma_n) = p_{c, \sigma_\theta}^{(\text{Name} \setminus rv(\pi))}(v_{\bar{c}^\pi, \sigma_\theta}(\sigma_n)).$$

We now prove Theorem 4.5 using the two lemmas.

PROOF OF THEOREM 4.5. Let $S = rv(\pi)$. Before starting the main derivation of the selective gradient estimator, we show the differentiability of several functions which are to be used in the derivation. From (R2) and (R3), the next functions over $\text{St}[\theta]$ are differentiable for all σ_n by the preservation of differentiability under function composition:

$$\begin{aligned} \sigma_\theta \mapsto p_{c_m, \sigma_\theta}(v_{\bar{c}_g^\pi, \sigma_\theta}(\sigma_n)), & \quad \sigma_\theta \mapsto p_{c_g, \sigma_\theta}^{(S)}(v_{\bar{c}_g^\pi, \sigma_\theta}(\sigma_n)), & \quad \sigma_\theta \mapsto p_{c_g, \sigma_\theta}^{(\text{Name} \setminus S)}(v_{\bar{c}_g^\pi, \sigma_\theta}(\sigma_n)), \\ \sigma_\theta \mapsto p_{\bar{c}_g^\pi, \sigma_\theta}^{(S)}(\sigma_n), & \quad \sigma_\theta \mapsto p_{\bar{c}_g^\pi, \sigma_\theta}^{(\text{Name} \setminus S)}(\sigma_n). \end{aligned}$$

From this, the next functions over $\text{St}[\theta]$ are also differentiable for all σ_n by Lemma D.3 with c_g and \bar{c}_g^π and by the fact that the multiplication of differentiable functions is differentiable:

$$\sigma_\theta \mapsto p_{c_g, \sigma_\theta}(v_{\bar{c}_g^\pi, \sigma_\theta}(\sigma_n)), \quad \sigma_\theta \mapsto p_{\bar{c}_g^\pi, \sigma_\theta}(\sigma_n).$$

These differentiability results are required in the below proof to apply several gradients rules (e.g., $\nabla_\theta(f(\theta) + g(\theta)) = \nabla_\theta f(\theta) + \nabla_\theta g(\theta)$) which may fail for non-differentiable functions.

Fix $\sigma_\theta \in \text{St}[\theta]$. Using the above differentiability results, we derive the selective gradient estimator as follows, where we write σ'_n for $v_{\bar{c}_g^\pi, \sigma_\theta}(\sigma_n)$:

$$\begin{aligned} & \nabla_\theta \mathcal{L}_\theta \\ &= \nabla_\theta \int d\sigma_n \left(p_{c_g, \sigma_\theta}(\sigma_n) \cdot \log \frac{p_{c_m, \sigma_\theta}(\sigma_n)}{p_{c_g, \sigma_\theta}(\sigma_n)} \right) \\ &= \nabla_\theta \int d\sigma_n \left(p_{\bar{c}_g^\pi, \sigma_\theta}(\sigma_n) \cdot \log \frac{p_{c_m, \sigma_\theta}(\sigma'_n)}{p_{c_g, \sigma_\theta}(\sigma'_n)} \right) \\ &= \int d\sigma_n \nabla_\theta \left(p_{\bar{c}_g^\pi, \sigma_\theta}(\sigma_n) \cdot \log \frac{p_{c_m, \sigma_\theta}(\sigma'_n)}{p_{c_g, \sigma_\theta}(\sigma'_n)} \right) \\ &= \int d\sigma_n \left(\nabla_\theta p_{\bar{c}_g^\pi, \sigma_\theta}(\sigma_n) \cdot \log \frac{p_{c_m, \sigma_\theta}(\sigma'_n)}{p_{c_g, \sigma_\theta}(\sigma'_n)} + p_{\bar{c}_g^\pi, \sigma_\theta}(\sigma_n) \cdot \nabla_\theta \log \frac{p_{c_m, \sigma_\theta}(\sigma'_n)}{p_{c_g, \sigma_\theta}(\sigma'_n)} \right) \\ &= \int d\sigma_n p_{\bar{c}_g^\pi, \sigma_\theta}(\sigma_n) \left(\nabla_\theta \log p_{\bar{c}_g^\pi, \sigma_\theta}(\sigma_n) \cdot \log \frac{p_{c_m, \sigma_\theta}(\sigma'_n)}{p_{c_g, \sigma_\theta}(\sigma'_n)} - \nabla_\theta \log p_{c_g, \sigma_\theta}(\sigma'_n) + \nabla_\theta \log p_{c_m, \sigma_\theta}(\sigma'_n) \right) \\ &= \int d\sigma_n p_{\bar{c}_g^\pi, \sigma_\theta}(\sigma_n) \left[\left(\nabla_\theta \log p_{\bar{c}_g^\pi, \sigma_\theta}^{(S)}(\sigma_n) + \nabla_\theta \log p_{\bar{c}_g^\pi, \sigma_\theta}^{(\text{Name} \setminus S)}(\sigma_n) \right) \cdot \log \frac{p_{c_m, \sigma_\theta}(\sigma'_n)}{p_{c_g, \sigma_\theta}(\sigma'_n)} \right. \\ & \quad \left. - \left(\nabla_\theta \log p_{c_g, \sigma_\theta}^{(S)}(\sigma'_n) + \nabla_\theta \log p_{c_g, \sigma_\theta}^{(\text{Name} \setminus S)}(\sigma'_n) \right) + \nabla_\theta \log p_{c_m, \sigma_\theta}(\sigma'_n) \right] \\ &= \int d\sigma_n p_{\bar{c}_g^\pi, \sigma_\theta}(\sigma_n) \left[\left(\mathbf{0} + \nabla_\theta \log p_{\bar{c}_g^\pi, \sigma_\theta}^{(\text{Name} \setminus S)}(\sigma_n) \right) \cdot \log \frac{p_{c_m, \sigma_\theta}(\sigma'_n)}{p_{c_g, \sigma_\theta}(\sigma'_n)} \right. \\ & \quad \left. - \left(\nabla_\theta \log p_{c_g, \sigma_\theta}^{(S)}(\sigma'_n) + \mathbf{0} \right) + \nabla_\theta \log p_{c_m, \sigma_\theta}(\sigma'_n) \right] \\ &= \int d\sigma_n p_{\bar{c}_g^\pi, \sigma_\theta}(\sigma_n) \left(\nabla_\theta \log p_{\bar{c}_g^\pi, \sigma_\theta}^{(\text{Name} \setminus S)}(\sigma_n) \cdot \log \frac{p_{c_m, \sigma_\theta}(\sigma'_n)}{p_{c_g, \sigma_\theta}(\sigma'_n)} \right. \\ & \quad \left. - \nabla_\theta \log p_{c_g, \sigma_\theta}^{(S)}(\sigma'_n) + \nabla_\theta \log p_{c_m, \sigma_\theta}(\sigma'_n) \right). \end{aligned}$$

We justify key steps of the above derivation below.

- The second equality comes from Theorem 4.2 and the fact that v_{c_g, σ_θ} is the identity function (since the third argument of every sample command in c_g is the identity function $\lambda y. y$).

- The third equality holds because differentiation there commutes with integration by (R5).
- The fourth comes from the product rule for differentiation: $\nabla_\theta(f(\theta) \cdot g(\theta)) = \nabla_\theta f(\theta) \cdot g(\theta) + f(\theta) \cdot \nabla_\theta g(\theta)$ for all differentiable f and g . Here f and g in the original equation are differentiable because differentiability is preserved under division and log for positive-valued functions.
- The fifth equality holds because $\nabla_\theta f(\theta) = f(\theta) \cdot \nabla_\theta \log f(\theta)$ for all differentiable and positive-valued f .
- The sixth equality follows from Lemma D.3 applied to c_g and \overline{c}_g^π (both of which do not contain observe commands), and from the linearity of differentiation: $\nabla_\theta(f(\theta) + g(\theta)) = \nabla_\theta f(\theta) + \nabla_\theta g(\theta)$ for all differentiable f and g . Here f and g in the original equation are differentiable because differentiability is preserved under log for positive-valued functions.
- The seventh equality follows from (R4) and

$$\mathbb{E}_{p_{\overline{c}_g^\pi, \sigma_\theta}(\sigma_n)} \left[\nabla_\theta \log p_{c_g, \sigma_\theta}^{\langle \text{Name} \setminus S \rangle}(\sigma'_n) \right] = 0. \quad (21)$$

The proof of Eq. (21) will be given after we complete this justification of the derivation.

- The last equality comes from Lemma D.4 applied to c_g .

The only remaining part is to prove Eq. (21). We derive the equation as follows:

$$\begin{aligned} & \int d\sigma_n \left(p_{\overline{c}_g^\pi, \sigma_\theta}(\sigma_n) \cdot \nabla_\theta \log p_{c_g, \sigma_\theta}^{\langle \text{Name} \setminus S \rangle}(\sigma'_n) \right) \\ &= \int d\sigma_n \left(p_{\overline{c}_g^\pi, \sigma_\theta}(\sigma_n) \cdot \nabla_\theta \log p_{\overline{c}_g^\pi, \sigma_\theta}^{\langle \text{Name} \setminus S \rangle}(\sigma_n) \right) \\ &= \int d\sigma_n \left(p_{\overline{c}_g^\pi, \sigma_\theta}(\sigma_n) \cdot \left(\nabla_\theta \log p_{\overline{c}_g^\pi, \sigma_\theta}^{\langle \text{Name} \setminus S \rangle}(\sigma_n) + \nabla_\theta \log p_{\overline{c}_g^\pi, \sigma_\theta}^{(S)}(\sigma_n) \right) \right) \\ &= \int d\sigma_n \left(p_{\overline{c}_g^\pi, \sigma_\theta}(\sigma_n) \cdot \nabla_\theta \log p_{\overline{c}_g^\pi, \sigma_\theta}(\sigma_n) \right) \\ &= \int d\sigma_n \nabla_\theta p_{\overline{c}_g^\pi, \sigma_\theta}(\sigma_n) \\ &= \nabla_\theta \int d\sigma_n p_{\overline{c}_g^\pi, \sigma_\theta}(\sigma_n) \\ &= \nabla_\theta \mathbf{1} = 0. \end{aligned}$$

Here is the justification of the above derivation:

- The first equality comes from Lemma D.4 applied to c_g .
- The second equality follows from (R4).
- The third equality holds because of Lemma D.3 applied to \overline{c}_g^π (which does not contain observe commands), and the linearity of differentiation: $\nabla_\theta(f(\theta) + g(\theta)) = \nabla_\theta f(\theta) + \nabla_\theta g(\theta)$ for all differentiable f and g . Here f and g in the original equation are differentiable because differentiability is preserved under log for positive-valued functions.
- The fourth equality holds because $\nabla_\theta f(\theta) = f(\theta) \cdot \nabla_\theta \log f(\theta)$ for all differentiable and positive-valued f .
- The fifth equality uses (R5), which states the commutativity between differentiation and integration in the equality.
- The six equality comes from that $p_{\overline{c}_g^\pi, \sigma_\theta}$ is a probability density by Remark 4.3.

This completes the proof. □

D.3 Proofs of Lemmas D.3 and D.4

We define the partial density version of $pr_{\square}^{(S)}(c)$ for $S \subseteq \text{Name}$:

$$pr_{\square}^{(S)}(c) : \text{St}[\text{PVar}] \times \text{St}_{\square}[\text{Name}] \rightarrow [0, \infty),$$

$$pr_{\square}^{(S)}(c)(\sigma_p, \xi_n) \triangleq \begin{cases} \prod_{\mu \in \text{dom}(\xi_n) \cap S} \llbracket c \rrbracket(\sigma_p \oplus \xi_n \oplus \sigma_r)(pr_{\mu}) & \text{if } \exists \sigma_r. \text{used}(c, \sigma_p \oplus \xi_n \oplus \sigma_r, \xi_n) \\ 0 & \text{otherwise.} \end{cases}$$

$pr_{\square}^{(S)}(c)$ enjoys many of the properties that $pr_{\square}(c)$ has. For instance, $pr_{\square}^{(S)}(c)$ is a well-defined function (i.e., its value does not depend on the choice of σ_r), as $pr_{\square}(c)$ does. Since the proof of those properties of $pr_{\square}^{(S)}(c)$ would be almost identical to that of $pr_{\square}(c)$, we will use them in the following proofs without explicitly (re)proving them.

PROOF OF LEMMA D.3. Let c be a command that has no observe commands. Let $S \subseteq \text{Name}$, $\sigma_{\theta} \in \text{St}[\theta]$, and $\sigma_n \in \text{St}[\text{Name}]$. We set σ_0 as in the definition of $p_{c, \sigma_{\theta}}$ in Eq. (3) (as a function of σ_n). Let $\sigma = \sigma_{\theta} \oplus \sigma_n \oplus \sigma_0$. If $noerr(c, \sigma)$ does not hold, then the LHS and RHS of the desired equation become zero, so the equation holds. If $noerr(c, \sigma)$ holds, we get the desired equation as follows:

$$\begin{aligned} p_{c, \sigma_{\theta}}^{(S)}(\sigma_n) \cdot p_{c, \sigma_{\theta}}^{(\text{Name} \setminus S)}(\sigma_n) &= (\llbracket c \rrbracket \sigma(\text{like}) \cdot \prod_{\mu \in S} \llbracket c \rrbracket \sigma(pr_{\mu})) \cdot (\llbracket c \rrbracket \sigma(\text{like}) \cdot \prod_{\mu \in \text{Name} \setminus S} \llbracket c \rrbracket \sigma(pr_{\mu})) \\ &= (\llbracket c \rrbracket \sigma(\text{like}))^2 \cdot \prod_{\mu \in \text{Name}} \llbracket c \rrbracket \sigma(pr_{\mu}) \\ &= \llbracket c \rrbracket \sigma(\text{like}) \cdot p_{c, \sigma_{\theta}}(\sigma_n) \\ &= \sigma(\text{like}) \cdot p_{c, \sigma_{\theta}}(\sigma_n) \\ &= p_{c, \sigma_{\theta}}(\sigma_n). \end{aligned}$$

The second last equality uses Lemma D.5, and the last equality uses $\sigma(\text{like}) = 1$ (which holds by the definition of σ_0). This completes the proof. \square

PROOF OF LEMMA D.4. Consider a command c , a reparameterisation plan π , $\sigma_{\theta} \in \text{St}[\theta]$, and $\sigma_n \in \text{St}[\text{Name}]$. Assume that all the sample commands of c have $\lambda y. y$ as their third arguments, and $p_{\bar{c}^{\pi}, \sigma_{\theta}}(\sigma_n) > 0$.

We first define several objects and make observations on them. Let $S \triangleq \text{Name} \setminus rv(\pi)$. Define $f_* : \text{St}[\text{Name}] \rightarrow \text{St}[\text{AVar}]$ to be the function for constructing an initial state:

$$f_*(\sigma_n)(v) \triangleq \begin{cases} f_{pr}(\sigma_n(\mu)) & \text{if } v \equiv pr_{\mu} \text{ for some } \mu \\ f_{vai}(\sigma_n(\mu)) & \text{if } v \equiv val_{\mu} \text{ for some } \mu \\ f_{cnt}(\sigma_n(\mu)) & \text{if } v \equiv cnt_{\mu} \text{ for some } \mu \\ 1 & \text{if } v \equiv \text{like}, \end{cases}$$

where $f_{vai}(r) \triangleq r$, $f_{pr}(r) \triangleq \mathcal{N}(r; 0, 1)$, and $f_{cnt}(r) \triangleq 0$. Define initial states $\bar{\sigma}, \sigma \in \text{St}$ for $p_{\bar{c}^{\pi}, \sigma_{\theta}}(\sigma_n)$ and $p_{c, \sigma_{\theta}}(v_{\bar{c}^{\pi}, \sigma_{\theta}}(\sigma_n))$, respectively, as

$$\bar{\sigma} \triangleq \sigma_p \oplus \sigma_n \oplus f_*(\sigma_n), \quad \sigma \triangleq \sigma_p \oplus v_{\bar{c}^{\pi}, \sigma_{\theta}}(\sigma_n) \oplus f_*(v_{\bar{c}^{\pi}, \sigma_{\theta}}(\sigma_n)),$$

where $\sigma_p \triangleq \sigma_{\theta} \oplus (\lambda v \in \text{PVar} \setminus \theta. 0) \in \text{St}[\text{PVar}]$. Then, the assumption $p_{\bar{c}^{\pi}, \sigma_{\theta}}(\sigma_n) > 0$ implies $noerr(\bar{c}^{\pi}, \bar{\sigma})$ by the definition of p . From this, $\bar{\sigma}(\text{like}) = 1$, and the definition of $used$, there exists $\bar{\xi}_n \in \text{St}_{\square}[\text{Name}]$ such that $used(\bar{c}^{\pi}, \bar{\sigma}, \bar{\xi}_n)$. From this, we have

$$\bar{\sigma} = \sigma_p \oplus \bar{\xi}_n \oplus \bar{\sigma}_r, \quad used(\bar{c}^{\pi}, \sigma_p \oplus \bar{\xi}_n \oplus \bar{\sigma}_r, \bar{\xi}_n),$$

for some $\bar{\sigma}_r$. Next, let

$$\xi_n \triangleq \text{vals}_{\square}(\bar{c}^{\pi})(\sigma_p, \bar{\xi}_n).$$

We can apply Lemma D.6 to $used(\overline{c^\pi}, \overline{\sigma}, \overline{\xi_n})$, since all the sample commands of c have $\lambda y.y$ in their third arguments (by assumption). The application of the lemma gives:

$$\forall \sigma'_r \in \text{St}[\text{Var} \setminus (\text{PVar} \cup \text{dom}(\overline{\xi_n}))]. \quad \sigma'_r(\text{like}) = 1 \implies used(c, \sigma_p \oplus \xi_n \oplus \sigma'_r, \xi_n),$$

$$prs_{\square}^{(S)}(\overline{c^\pi})(\sigma_p, \overline{\xi_n}) = prs_{\square}^{(S)}(c)(\sigma_p, \xi_n),$$

where the for-all part comes from Lemma C.7.

We now show two claims. The first claim is: there exists σ_r such that

$$\sigma = \sigma_p \oplus \xi_n \oplus \sigma_r, \quad used(c, \sigma_p \oplus \xi_n \oplus \sigma_r, \xi_n).$$

By the definition of σ , it suffices to show that $\xi_n = (v_{\overline{c^\pi}, \sigma_\theta}(\sigma_n))|_{\text{dom}(\xi_n)}$. This indeed holds as follows: for any $\mu \in \text{dom}(\xi_n)$, $\xi_n(\mu) = \text{vals}_{\square}(\overline{c^\pi})(\sigma_p, \overline{\xi_n})(\mu) = \llbracket \overline{c^\pi} \rrbracket \overline{\sigma}(\text{val}_\mu) = v_{\overline{c^\pi}, \sigma_\theta}(\sigma_n)(\mu)$, where the second equality uses $used(\overline{c^\pi}, \sigma_p \oplus \overline{\xi_n} \oplus \overline{\sigma_r}, \overline{\xi_n})$. The second claim is: for all $\mu \in S \setminus \text{dom}(\xi_n)$,

$$\sigma(pr_\mu) = \overline{\sigma}(pr_\mu).$$

Here is the proof of the claim: $\sigma(pr_\mu) = f_{pr}(\sigma(\mu)) = f_{pr}(v_{\overline{c^\pi}, \sigma_\theta}(\sigma_n)(\mu)) = f_{pr}(\llbracket \overline{c^\pi} \rrbracket \overline{\sigma}(\text{val}_\mu)) = f_{pr}(\overline{\sigma}(\text{val}_\mu)) = f_{pr}(\overline{\sigma}(\mu))$; and $\overline{\sigma}(pr_\mu) = f_{pr}(\overline{\sigma}(\mu))$; here the second last equality in the first equation uses Lemma C.6-(2) with $\mu \notin \text{dom}(\overline{\xi_n})$ and $used(\overline{c^\pi}, \sigma_p \oplus \overline{\xi_n} \oplus \overline{\sigma_r}, \overline{\xi_n})$, and the last equality in the first equation uses $f_{\text{val}}(r) = r$.

Based on the observations made so far, we show the desired equation as follows:

$$\begin{aligned} p_{\overline{c^\pi}, \sigma_\theta}^{(S)}(\sigma_n) &= \prod_{\mu \in S \cap \text{dom}(\overline{\xi_n})} \llbracket \overline{c^\pi} \rrbracket \overline{\sigma}(pr_\mu) \cdot \prod_{\mu \in S \setminus \text{dom}(\overline{\xi_n})} \llbracket \overline{c^\pi} \rrbracket \overline{\sigma}(pr_\mu) \\ &= prs_{\square}^{(S)}(\overline{c^\pi})(\sigma_p, \overline{\xi_n}) \cdot \prod_{\mu \in S \setminus \text{dom}(\overline{\xi_n})} \overline{\sigma}(pr_\mu) \\ &= prs_{\square}^{(S)}(c)(\sigma_p, \xi_n) \cdot \prod_{\mu \in S \setminus \text{dom}(\xi_n)} \sigma(pr_\mu) \\ &= \prod_{\mu \in S \cap \text{dom}(\xi_n)} \llbracket c \rrbracket \sigma(pr_\mu) \cdot \prod_{\mu \in S \setminus \text{dom}(\xi_n)} \llbracket c \rrbracket \sigma(pr_\mu) \\ &= p_{c, \sigma_\theta}^{(S)}(v_{\overline{c^\pi}, \sigma_\theta}(\sigma_n)). \end{aligned}$$

The first and last equalities are by the definition of p . The second equality uses $used(\overline{c^\pi}, \sigma_p \oplus \overline{\xi_n} \oplus \overline{\sigma_r}, \overline{\xi_n})$ and Lemma C.6-(2) with $\mu \notin \text{dom}(\overline{\xi_n})$. The third equality uses $\text{dom}(\overline{\xi_n}) = \text{dom}(\xi_n)$, the observation made in the first paragraph, and the second claim in the above. The fourth equality uses the first claim in the above, and Lemma C.6-(2) with $\mu \notin \text{dom}(\xi_n)$. \square

LEMMA D.5. *Let c be a command and $\sigma \in \text{St}$. If c has no observe commands and $\llbracket c \rrbracket \sigma \in \text{St}$, then*

$$\llbracket c \rrbracket \sigma(\text{like}) = \sigma(\text{like}).$$

PROOF. Let c be a command that does not contain an observe command. We show the claim of the lemma by induction on the structure of c . Pick $\sigma \in \text{St}$ such that $\llbracket c \rrbracket \sigma \in \text{St}$. We will show that $\llbracket c \rrbracket \sigma(\text{like}) = \sigma(\text{like})$.

Case $c \equiv \text{skip}$. In this case, $\llbracket c \rrbracket \sigma(\text{like}) = \sigma(\text{like})$ by the definition of the semantics.

Case $c \equiv (x := e)$. Again, $\llbracket c \rrbracket \sigma(\text{like}) = \sigma(\text{like})$ by the definition of the semantics.

Case $c \equiv (x := \text{sam}(n, d, \lambda y.e'))$. Once more, $\llbracket c \rrbracket \sigma(\text{like}) = \sigma(\text{like})$ by the definition of the semantics.

Case $c \equiv (c'; c'')$. We have $\llbracket c' \rrbracket \sigma \in \text{St}$ and $\llbracket c'' \rrbracket (\llbracket c' \rrbracket \sigma) \in \text{St}$. We apply induction hypothesis first to (c', σ) , and again to $(c', \llbracket c' \rrbracket \sigma)$. The first application gives $\llbracket c' \rrbracket \sigma(\text{like}) = \sigma(\text{like})$, and the second $\llbracket c'; c'' \rrbracket \sigma(\text{like}) = \llbracket c' \rrbracket \sigma(\text{like})$. The desired conclusion follows from these two equalities.

Case $c \equiv (\text{if } b \{c'\} \text{ else } \{c''\})$. We deal with the case that $\llbracket b \rrbracket \sigma = \text{true}$. The other case of $\llbracket b \rrbracket \sigma = \text{false}$ can be proved similarly. Since $\llbracket b \rrbracket \sigma = \text{true}$, we have $\llbracket c' \rrbracket \sigma = \llbracket c \rrbracket \sigma \in \text{St}$. Thus, we can apply induction hypothesis to c' . If we do so, we get $\llbracket c' \rrbracket \sigma(\text{like}) = \sigma(\text{like})$. This gives the desired conclusion because $\llbracket c \rrbracket \sigma = \llbracket c' \rrbracket \sigma$.

Case $c \equiv (\text{while } b \{c'\})$. Let F be the operator on $[\text{St} \rightarrow \text{St}_\perp]$ such that $\llbracket c \rrbracket$ is the least fixed point of F . Define a subset \mathcal{T} of $[\text{St} \rightarrow \text{St}_\perp]$ as follows:

$$f \in \mathcal{T} \iff \left(\forall \sigma' \in \text{St}. f(\sigma') \in \text{St} \implies f(\sigma')(\text{like}) = \sigma'(\text{like}) \right)$$

The set \mathcal{T} contains the least function $\lambda \sigma. \perp$, and is closed under the least upper bound of any chain in $[\text{St} \rightarrow \text{St}_\perp]$. It is also closed under F . This F -closure follows essentially from our arguments for sequential composition, if command, and skip, and induction hypothesis on c' . What we have shown for \mathcal{T} implies that \mathcal{T} contains the least fixed point of F , which gives the desired property for c . \square

LEMMA D.6. *Let c be a command and π be a reparameterisation plan. Suppose that every sample command in c has $\lambda y. y$ as its third argument. Then, for all $\sigma_p \in \text{St}[\text{PVar}]$ and $\bar{\xi}_n \in \text{St}_\square[\text{Name}]$, if used $(\bar{c}^\pi, \sigma_p \oplus \bar{\xi}_n \oplus \bar{\sigma}_r, \bar{\xi}_n)$ for some $\bar{\sigma}_r$, then*

$$\begin{aligned} & \exists \sigma_r. \text{used}(c, \sigma_p \oplus \bar{\xi}_n \oplus \sigma_r, \bar{\xi}_n), \\ & \text{pvars}_\square(\bar{c}^\pi)(\sigma_p, \bar{\xi}_n) = \text{pvars}_\square(c)(\sigma_p, \bar{\xi}_n), \\ & \text{prs}_\square^{\langle \text{Name} \setminus \text{rv}(\pi) \rangle}(\bar{c}^\pi)(\sigma_p, \bar{\xi}_n) = \text{prs}_\square^{\langle \text{Name} \setminus \text{rv}(\pi) \rangle}(c)(\sigma_p, \bar{\xi}_n), \end{aligned}$$

where

$$\bar{\xi}_n \triangleq \text{vals}_\square(\bar{c}^\pi)(\sigma_p, \bar{\xi}_n).$$

PROOF. Fix a reparameterisation plan π . The proof proceeds by induction on the structure of c . Let $\sigma_p \in \text{St}[\text{PVar}]$, and $\bar{\xi}_n \in \text{St}_\square[\text{Name}]$. Assume that c uses only $\lambda y. y$ in the third argument of its sample commands, and used $(\bar{c}^\pi, \sigma_p \oplus \bar{\xi}_n \oplus \bar{\sigma}_r, \bar{\xi}_n)$ for some $\bar{\sigma}_r$. Let $\bar{\sigma} \triangleq \sigma_p \oplus \bar{\xi}_n \oplus \bar{\sigma}_r$ and $S \triangleq \text{Name} \setminus \text{rv}(\pi)$. Then, we simply have used $(\bar{c}^\pi, \bar{\sigma}, \bar{\xi}_n)$.

Cases $c \equiv \text{skip}$, $c \equiv (x := e)$, or $c \equiv \text{obs}(d, r)$. In this case, $\llbracket c \rrbracket \sigma(\text{cnt}_\mu) = \sigma(\text{cnt}_\mu)$ for all $\sigma \in \text{St}$ and $\mu \in \text{Name}$. So $\text{dom}(\bar{\xi}_n) = \text{dom}(\bar{\xi}_n) = \emptyset$ and thus $\bar{\xi}_n = \bar{\xi}_n$. We also know $\bar{c}^\pi \equiv c$. From these, all of the three conclusions follow immediately.

Case $c \equiv (x := \text{sam}(n, d, \lambda y. e))$. Since $\text{fv}(n) \subseteq \text{PVar}$, there exists $\mu \in \text{Name}$ such that $\llbracket n \rrbracket (\sigma_p \oplus \sigma_r) = \mu$ for all $\sigma_r \in \text{St}[\text{Var} \setminus \text{PVar}]$. So, for all $\sigma_r \in \text{St}[\text{Var} \setminus \text{PVar}]$ and $\mu' \in \text{Name} \setminus \{\mu\}$,

$$\llbracket c \rrbracket (\sigma_p \oplus \sigma_r)(\text{cnt}_{\mu'}) = \begin{cases} (\sigma_p \oplus \sigma_r)(\text{cnt}_{\mu'}) + 1 & \text{if } \mu' = \mu \\ (\sigma_p \oplus \sigma_r)(\text{cnt}_{\mu'}) & \text{otherwise.} \end{cases} \quad (22)$$

From this, we get $\text{dom}(\bar{\xi}_n) = \text{dom}(\bar{\xi}_n) = \{\mu\}$. Further, by assumption, we get $e \equiv y$. We now prove the three conclusions based on these observations and case analysis on $(n, d, \lambda y. e)$.

First, assume $(n, d, \lambda y. e) \notin \text{dom}(\pi)$. Then, $\bar{c}^\pi \equiv c$ and

$$\bar{\xi}_n = \text{vals}_\square(\bar{c}^\pi)(\sigma_p, \bar{\xi}_n) = [\mu \mapsto \llbracket \bar{c}^\pi \rrbracket \bar{\sigma}(\text{val}_\mu)] = [\mu \mapsto \llbracket e[\bar{\sigma}(\mu)/y] \rrbracket \bar{\sigma}] = [\mu \mapsto \bar{\xi}_n(\mu)] = \bar{\xi}_n,$$

where the second last equality uses $e \equiv y$. Hence, the three conclusions clearly hold.

Next, assume $(n, d, \lambda y.e) \in \text{dom}(\pi)$. Suppose that $\pi(n, d, \lambda y.e) = (\bar{d}, \lambda \bar{y}.\bar{e})$. Then, $\bar{c}^\pi \equiv (x := \text{sam}(n, \bar{d}, \lambda \bar{y}.\bar{e}))$ and

$$\bar{\xi}_n = \text{vals}_\square(\bar{c}^\pi)(\sigma_p, \bar{\xi}_n) = [\mu \mapsto \llbracket \bar{c}^\pi \rrbracket \bar{\sigma}(\text{val}_\mu)] = [\mu \mapsto \llbracket \bar{e}[\bar{\sigma}(\mu)/\bar{y}] \rrbracket \bar{\sigma}] = [\mu \mapsto \llbracket \bar{e}[\bar{\xi}_n(\mu)/\bar{y}] \rrbracket \bar{\sigma}].$$

Since Eq. (22) holds also for $\llbracket \bar{c}^\pi \rrbracket$, and since $\text{dom}(\bar{\xi}_n) = \{\mu\}$, we get the first conclusion:

$$\text{used}(c, \sigma_p \oplus \bar{\xi}_n \oplus \bar{\sigma}_r, \bar{\xi}_n).$$

To prove the second conclusion, let $\sigma = \sigma_p \oplus \bar{\xi}_n \oplus \bar{\sigma}_r$. Then, for all $v \in \text{PVar}$,

$$\begin{aligned} p\text{vars}_\square(\bar{c}^\pi)(\sigma_p, \bar{\xi}_n)(v) &= \llbracket \bar{c}^\pi \rrbracket \bar{\sigma}(v) = \begin{cases} \llbracket \bar{e}[\bar{\sigma}(\mu)/\bar{y}] \rrbracket \bar{\sigma} = \llbracket \bar{e}[\bar{\xi}_n(\mu)/\bar{y}] \rrbracket \bar{\sigma} & \text{if } v \equiv x \\ \bar{\sigma}(v) = \sigma_p(v) & \text{otherwise,} \end{cases} \\ p\text{vars}_\square(c)(\sigma_p, \bar{\xi}_n)(v) &= \llbracket c \rrbracket \sigma(v) = \begin{cases} \llbracket e[\sigma(\mu)/y] \rrbracket \sigma = \bar{\xi}_n(\mu) = \llbracket \bar{e}[\bar{\xi}_n(\mu)/\bar{y}] \rrbracket \bar{\sigma} & \text{if } v \equiv x \\ \sigma(v) = \sigma_p(v), & \text{otherwise,} \end{cases} \end{aligned}$$

where the second equation uses $e \equiv y$. Hence, the second conclusion holds. For the third conclusion, let $n = \text{name}(\alpha, _)$. Then, $\mu = (\alpha, _) \in \{(\alpha, i) \in \text{Name} \mid i \in \mathbb{N}\} \subseteq \text{rv}(\pi)$. Thus, $\text{dom}(\bar{\xi}_n) \cap S = \text{dom}(\bar{\xi}_n) \cap S = \{\mu\} \cap S = \{\mu\} \cap (\text{Name} \setminus \text{rv}(\pi)) = \emptyset$. From this, we get the third conclusion:

$$\text{pr}_\square^{(S)}(\bar{c}^\pi)(\sigma_p, \bar{\xi}_n) = 1 = \text{pr}_\square^{(S)}(c)(\sigma_p, \bar{\xi}_n).$$

Case $c \equiv (c'; c'')$. First, we make several observations necessary to prove the conclusion. By $\text{used}(\bar{c}^\pi, \bar{\sigma}, \bar{\xi}_n)$, we have $\llbracket \bar{c}'^\pi \rrbracket \bar{\sigma} \in \text{St}$ and $\llbracket \bar{c}''^\pi \rrbracket (\llbracket \bar{c}'^\pi \rrbracket \bar{\sigma}) \in \text{St}$. Let

$$\bar{\sigma}' \triangleq \llbracket \bar{c}'^\pi \rrbracket \bar{\sigma}, \quad \bar{\sigma}'' \triangleq \llbracket \bar{c}''^\pi \rrbracket \bar{\sigma}', \quad \sigma'_p \triangleq \bar{\sigma}'|_{\text{PVar}}, \quad \sigma''_p \triangleq \bar{\sigma}''|_{\text{PVar}}.$$

Then, by $\text{used}(\bar{c}^\pi, \bar{\sigma}, \bar{\xi}_n)$ and the claim in the proof of Lemma C.6 (for the sequential composition case), there exist $\bar{\xi}'_n$ and $\bar{\xi}''_n$ such that

$$\bar{\xi}_n = \bar{\xi}'_n \oplus \bar{\xi}''_n, \quad \text{used}(\bar{c}'^\pi, \bar{\sigma}, \bar{\xi}'_n), \quad \text{used}_-(\bar{c}''^\pi, \bar{\sigma}', \bar{\xi}''_n).$$

By the latter two, we can apply induction to $(c', \sigma_p, \bar{\xi}'_n)$ and $(c'', \sigma'_p, \bar{\xi}''_n)$, and IH gives the following:

$$\begin{aligned} p\text{vars}_\square(c')(\sigma_p, \bar{\xi}'_n) &= p\text{vars}_\square(\bar{c}'^\pi)(\sigma_p, \bar{\xi}'_n) && \text{[By IH on } c'] \\ &= (\llbracket \bar{c}'^\pi \rrbracket (\bar{\sigma}[\text{like} \mapsto 1]))|_{\text{PVar}} && \text{[By used}(\bar{c}'^\pi, \bar{\sigma}, \bar{\xi}'_n)] \\ &= (\llbracket \bar{c}'^\pi \rrbracket \bar{\sigma})|_{\text{PVar}} = \bar{\sigma}'|_{\text{PVar}} = \sigma'_p, && \text{[By Lemma C.6-(3)]} \\ p\text{vars}_\square(c'')(\sigma'_p, \bar{\xi}''_n) &= p\text{vars}_\square(\bar{c}''^\pi)(\sigma'_p, \bar{\xi}''_n) && \text{[By IH on } c''] \\ &= (\llbracket \bar{c}''^\pi \rrbracket (\bar{\sigma}'[\text{like} \mapsto 1]))|_{\text{PVar}} && \text{[By used}_-(\bar{c}''^\pi, \bar{\sigma}', \bar{\xi}''_n)] \\ &= (\llbracket \bar{c}''^\pi \rrbracket \bar{\sigma}')|_{\text{PVar}} = \bar{\sigma}''|_{\text{PVar}} = \sigma''_p, && \text{[By Lemma C.6-(3)]} \end{aligned}$$

where

$$\bar{\xi}'_n \triangleq \text{vals}_\square(\bar{c}'^\pi)(\sigma_p, \bar{\xi}'_n), \quad \bar{\xi}''_n \triangleq \text{vals}_\square(\bar{c}''^\pi)(\sigma'_p, \bar{\xi}''_n).$$

By the former equation, we get

$$\begin{aligned} \bar{\xi}_n &= \text{vals}_\square(\bar{c}^\pi)(\sigma_p, \bar{\xi}_n) \\ &= \text{vals}_\square(\bar{c}'^\pi; \bar{c}''^\pi)(\sigma_p, \bar{\xi}'_n \oplus \bar{\xi}''_n) && \text{[By } \bar{\xi}_n = \bar{\xi}'_n \oplus \bar{\xi}''_n] \\ &= \text{vals}_\square(\bar{c}'^\pi)(\sigma_p, \bar{\xi}'_n) \oplus \text{vals}_\square(\bar{c}''^\pi)(p\text{vars}_\square(\bar{c}'^\pi)(\sigma_p, \bar{\xi}'_n), \bar{\xi}''_n) \\ &= \text{vals}_\square(\bar{c}'^\pi)(\sigma_p, \bar{\xi}'_n) \oplus \text{vals}_\square(\bar{c}''^\pi)(\sigma'_p, \bar{\xi}''_n) && \text{[By the former equation]} \end{aligned}$$

$$= \xi'_n \oplus \xi''_n$$

where the third equality uses $used(\overline{c'^\pi}, \overline{\sigma}, \overline{\xi'_n})$, $used(\overline{c''^\pi}; \overline{c''^\pi}, \overline{\sigma}, \overline{\xi'_n \oplus \xi''_n})$, and the second claim in the proof of Lemma C.9.

We now show the first conclusion. By IH on $(c', \sigma_p, \overline{\xi'_n})$ and $(c'', \sigma'_p, \overline{\xi''_n})$, we get

$$\exists \sigma'_r. used(c', \sigma_p \oplus \xi'_n \oplus \sigma'_r, \xi'_n), \quad \exists \sigma''_r. used(c'', \sigma'_p \oplus \xi''_n \oplus \sigma''_r, \xi''_n).$$

Let

$$\sigma \triangleq \sigma_p \oplus (\xi'_n \oplus \xi''_n) \oplus \sigma'_r |_{\text{dom}(\sigma'_r) \setminus \text{dom}(\xi''_n)}.$$

Then, $\llbracket c' \rrbracket \sigma \in \text{St}$ by $used(c', \sigma_p \oplus \xi'_n \oplus \sigma'_r, \xi'_n)$ and Lemma C.6-(1), and we have

$$\begin{aligned} \sigma|_{\text{PVar}} = \sigma_p, \quad (\llbracket c' \rrbracket \sigma)|_{\text{PVar}} &= (\llbracket c' \rrbracket (\sigma_p \oplus \xi'_n \oplus \sigma'_r))|_{\text{PVar}} && \text{[By Lemma C.6-(3)]} \\ &= pvars(c')(\sigma_p, \xi'_n) = \sigma'_p, && \text{[By the above]} \\ \sigma|_{\text{dom}(\xi_n)} = \xi'_n, \quad (\llbracket c' \rrbracket \sigma)|_{\text{dom}(\xi''_n)} &= \sigma|_{\text{dom}(\xi''_n)} = \xi''_n. && \text{[By Lemma C.4]} \end{aligned}$$

By these, $used(c', \sigma_p \oplus \xi'_n \oplus \sigma'_r, \xi'_n)$, $used(c'', \sigma'_p \oplus \xi''_n \oplus \sigma''_r, \xi''_n)$, and Lemma C.7, we get

$$used(c', \sigma, \xi'_n), \quad used_-(c'', \llbracket c' \rrbracket \sigma, \xi''_n).$$

By these and the third claim in the proof of Lemma C.9, we get

$$used(c'; c'', \sigma, \xi'_n \oplus \xi''_n).$$

By this and Lemma C.7, we get the following as desired, since $\xi_n = \xi'_n \oplus \xi''_n$ (shown in the above) and $\sigma = \sigma_p \oplus (\xi'_n \oplus \xi''_n) \oplus \sigma_r$ for some σ_r :

$$used(c, \sigma_p \oplus \xi_n \oplus \sigma_r, \xi_n).$$

Next, we show the second conclusion as follows:

$$\begin{aligned} pvars_{\square}(\overline{c^\pi})(\sigma_p, \xi_n) &= (\llbracket \overline{c^\pi} \rrbracket \overline{\sigma})|_{\text{PVar}} && \text{[By } used(\overline{c^\pi}, \overline{\sigma}, \overline{\xi_n})\text{]} \\ &= (\llbracket \overline{c''^\pi} \rrbracket (\llbracket \overline{c'^\pi} \rrbracket \overline{\sigma}))|_{\text{PVar}} \\ &= \overline{\sigma''}|_{\text{PVar}} = \sigma''_p, \\ pvars_{\square}(c)(\sigma_p, \xi_n) &= pvars_{\square}(c'; c'')(\sigma_p, \xi'_n \oplus \xi''_n) && \text{[By } \xi_n = \xi'_n \oplus \xi''_n\text{]} \\ &= pvars_{\square}(c'')(pvars_{\square}(c')(\sigma_p, \xi'_n), \xi''_n) \\ &= pvars_{\square}(c'')(\sigma'_p, \xi''_n) = \sigma''_p, && \text{[By the above]} \end{aligned}$$

where the second last equality uses $used(c', \sigma, \xi'_n)$, $used(c'; c'', \sigma, \xi'_n \oplus \xi''_n)$, and the second claim in the proof of Lemma C.9.

Lastly, we show the third conclusion as follows:

$$\begin{aligned} prs_{\square}^{(S)}(\overline{c^\pi})(\sigma_p, \overline{\xi_n}) &= prs_{\square}^{(S)}(\overline{c'; c''^\pi})(\sigma_p, \overline{\xi'_n \oplus \xi''_n}) && \text{[By } \overline{\xi_n} = \overline{\xi'_n \oplus \xi''_n}\text{]} \\ &= prs_{\square}^{(S)}(\overline{c'^\pi})(\sigma_p, \overline{\xi'_n}) \cdot prs_{\square}^{(S)}(\overline{c''^\pi})(pvars(\overline{c'^\pi})(\sigma_p, \overline{\xi'_n}), \overline{\xi''_n}) \\ &= prs_{\square}^{(S)}(\overline{c'^\pi})(\sigma_p, \overline{\xi'_n}) \cdot prs_{\square}^{(S)}(\overline{c''^\pi})(pvars(c')(\sigma_p, \xi'_n), \overline{\xi''_n}) && \text{[By IH on } c'\text{]} \\ &= prs_{\square}^{(S)}(c')(\sigma_p, \xi'_n) \cdot prs_{\square}^{(S)}(c'')(pvars(c')(\sigma_p, \xi'_n), \xi''_n) && \text{[By IH on } c' \text{ and } c''\text{]} \\ &= prs_{\square}^{(S)}(c)(\sigma_p, \xi'_n \oplus \xi''_n) \\ &= prs_{\square}^{(S)}(c)(\sigma_p, \xi_n). && \text{[By } \xi_n = \xi'_n \oplus \xi''_n\text{]} \end{aligned}$$

Here the second and fifth equalities use the second claim in the proof of Lemma C.9 with the following: $used(\overline{c'^\pi}, \overline{\sigma}, \overline{\xi'_n})$, $used(\overline{c''^\pi}; \overline{c''^\pi}, \overline{\sigma}, \overline{\xi'_n \oplus \xi''_n})$, $used(c', \sigma, \xi'_n)$, and $used(c'; c'', \sigma, \xi'_n \oplus \xi''_n)$.

Case $c \equiv (\text{if } b \{c'\} \text{ else } \{c''\})$. In this case, $\bar{c}^\pi \equiv (\text{if } b \{\bar{c}'^\pi\} \text{ else } \{\bar{c}''^\pi\})$. Since $\text{fv}(b) \subseteq \text{PVar}$, $\llbracket b \rrbracket(\sigma_p \oplus \sigma_r)$ is constant for all $\sigma_r \in \text{St}[\text{Var} \setminus \text{PVar}]$. Without loss of generality, assume $\llbracket b \rrbracket(\sigma_p \oplus \sigma_r) = \text{true}$. Then, $\llbracket c \rrbracket(\sigma_p \oplus \sigma_r) = \llbracket c' \rrbracket(\sigma_p \oplus \sigma_r)$ and $\llbracket \bar{c}^\pi \rrbracket(\sigma_p \oplus \sigma_r) = \llbracket \bar{c}'^\pi \rrbracket(\sigma_p \oplus \sigma_r)$ for all $\sigma_r \in \text{St}[\text{Var} \setminus \text{PVar}]$. Hence, by IH on $(c', \sigma_p, \bar{\xi}_n)$, we get the three conclusions directly.

Case $c \equiv (\text{while } b \{c'\})$. In this case, $\bar{c}^\pi \equiv (\text{while } b \{\bar{c}'^\pi\})$. Consider the version of $\text{prs}_\square(-)$ where the parameter can be a state transformer $\bar{f} : \text{St} \rightarrow \text{St}_\perp$, instead of a command. Similarly, consider the version of the three conclusions where we use two state transformers $\bar{f}, f : \text{St} \rightarrow \text{St}_\perp$, again instead of a command. We denote the versions by $\text{prs}_\square(\bar{f})$ and $\varphi(\bar{f}, f, \sigma_p, \bar{\xi}_n)$. We write $\bar{f} \sim f$ if $\text{prs}_\square(\bar{f})(\sigma_p, \bar{\xi}_n) > 0$ implies $\varphi(\bar{f}, f, \sigma_p, \bar{\xi}_n)$ for all $\sigma_p \in \text{St}[\text{PVar}]$ and $\bar{\xi}_n \in \text{St}_\square[\text{Name}]$. Further, we define $F^{\pi'} : [\text{St} \rightarrow \text{St}_\perp] \rightarrow [\text{St} \rightarrow \text{St}_\perp]$ as

$$F^{\pi'}(f)(\sigma) \triangleq \text{if } (\llbracket b \rrbracket \sigma = \text{true}) \text{ then } (f^\dagger \circ \llbracket \bar{c}'^{\pi'} \rrbracket)(\sigma) \text{ else } \sigma.$$

Note that F^π and F^{π_0} are the operators used in the semantics of the loops $\llbracket \bar{c}^\pi \rrbracket$ and $\llbracket c \rrbracket$, respectively, where π_0 denotes the empty reparameterisation plan. We will show three claims: $\lambda \sigma. \perp \sim \lambda \sigma. \perp$; if $\bar{f} \sim f$, then $\overline{F^\pi}(\bar{f}) \sim F^{\pi_0}(f)$; and if increasing sequences $\{\bar{f}_k\}_{k \in \mathbb{N}}$ and $\{f_k\}_{k \in \mathbb{N}}$ satisfy $\bar{f}_k \sim f_k$ for all $k \in \mathbb{N}$, then $\bar{f}_\infty \sim f_\infty$ holds for $\bar{f}_\infty = \bigsqcup_{k \in \mathbb{N}} \bar{f}_k$ and $f_\infty = \bigsqcup_{k \in \mathbb{N}} f_k$. These three claims imply $\llbracket \bar{c}^\pi \rrbracket \sim \llbracket c \rrbracket$, which in turn proves the desired three conclusions.

The first claim holds simply because $\text{prs}_\square(\lambda \sigma. \perp)(-, -)$ is always 0. To show the second claim, consider $\bar{f}, f : \text{St} \rightarrow \text{St}_\perp$ such that $\bar{f} \sim f$. We first replay our proof for the sequential-composition case on (\bar{f}, f) after viewing $\bar{f}^\dagger \circ \llbracket \bar{c}'^\pi \rrbracket$ and $f^\dagger \circ \llbracket c' \rrbracket$ as the sequential composition of \bar{c}'^π and \bar{f} , and of c' and f , respectively. This replay, then, gives the relationship $\bar{f}^\dagger \circ \llbracket \bar{c}'^\pi \rrbracket \sim f^\dagger \circ \llbracket c' \rrbracket$. Next, we replay our proof for the if case on $(F^\pi(\bar{f}), F^{\pi_0}(f))$, after viewing $\bar{f}^\dagger \circ \llbracket \bar{c}'^\pi \rrbracket$ and $f^\dagger \circ \llbracket c' \rrbracket$ as the true branches, and $\lambda \sigma. \sigma = \llbracket \text{skip} \rrbracket$ as the false branch. This replay implies the required relationship $F^\pi(\bar{f}) \sim F^{\pi_0}(f)$.

To show the third condition, consider increasing sequences $\{\bar{f}_k\}_{k \in \mathbb{N}}$ and $\{f_k\}_{k \in \mathbb{N}}$ such that $\bar{f}_k \sim f_k$ for all $k \in \mathbb{N}$. Let $\bar{f}_\infty = \bigsqcup_{k \in \mathbb{N}} \bar{f}_k$ and $f_\infty = \bigsqcup_{k \in \mathbb{N}} f_k$. Consider any σ_p and $\bar{\xi}_n$ such that $\text{prs}_\square(\bar{f}_\infty)(\sigma_p, \bar{\xi}_n) > 0$. We should show $\varphi(\bar{f}_\infty, f_\infty, \sigma_p, \bar{\xi}_n)$. Pick any $\sigma_r \in \text{St}[\text{Var} \setminus (\text{PVar} \cup \text{dom}(\bar{\xi}_n))]$ with $\sigma_r(\text{like}) = 1$. Let $\bar{\sigma} = \sigma_p \oplus \bar{\xi}_n \oplus \sigma_r$ and $\sigma = \sigma_p \oplus \xi_n \oplus \sigma_r$. Note that the value of each term in $\varphi(\dots)$ (i.e., $\text{used}(\dots)$, $\text{pvars}_\square(\dots)$, and $\text{prs}_\square^{(S)}(\dots)$) is independent of the choice of σ_r by Lemma C.7 and the well-definedness of pvars_\square and $\text{prs}_\square^{(S)}$. Since the two given sequences are increasing, there exists $K \in \mathbb{N}$ such that $\bar{f}_\infty(\bar{\sigma}) = \bar{f}_K(\bar{\sigma})$ and $f_\infty(\sigma) = f_K(\sigma)$. From this and $\text{prs}_\square(\bar{f}_\infty)(\sigma_p, \bar{\xi}_n) > 0$, we have $\text{prs}_\square(\bar{f}_K)(\sigma_p, \bar{\xi}_n) > 0$. This in turn gives $\varphi(\bar{f}_K, f_K, \sigma_p, \bar{\xi}_n)$ since $\bar{f}_K \sim f_K$. Lastly, again by $\bar{f}_\infty(\bar{\sigma}) = \bar{f}_K(\bar{\sigma})$ and $f_\infty(\sigma) = f_K(\sigma)$, we obtain $\varphi(\bar{f}_\infty, f_\infty, \sigma_p, \bar{\xi}_n)$ as desired. \square

E DEFERRED RESULTS IN §4.3

E.1 Deferred Statements and Their Proofs

LEMMA E.1. *Let $f, g : \mathbb{R}^n \rightarrow \mathbb{R}$ be locally Lipschitz functions. Then, the following differentiation rules hold for almost every $x \in \mathbb{R}^n$:*

$$\begin{aligned} \nabla(f + g)(x) &= \nabla f(x) + \nabla g(x), \\ \nabla(f \cdot g)(x) &= \nabla f(x) \cdot g(x) + f(x) \cdot \nabla g(x), \\ \nabla \log f(x) &= 1/f(x) \cdot \nabla f(x), \end{aligned}$$

where for the third rule we assume $f(y) > 0$ for all $y \in \mathbb{R}^n$.

PROOF. Note that the three functions $+$: $\mathbb{R}^2 \rightarrow \mathbb{R}$, \cdot : $\mathbb{R}^2 \rightarrow \mathbb{R}$, and \log : $\mathbb{R}_{>0} \rightarrow \mathbb{R}$ are all differentiable. Hence, applying Lemma E.4 produces the claim. \square

THEOREM E.2. Let $f : \mathbb{R} \times \mathbb{R}^n \rightarrow \mathbb{R}$ be a measurable function that satisfies the next four conditions:

- For all $x \in \mathbb{R}^n$, $f(-, x) : \mathbb{R} \rightarrow \mathbb{R}$ is continuous.
- For all $\theta \in \mathbb{R}$, $\nabla_\theta f(\theta, x)$ is well-defined for almost all $x \in \mathbb{R}^n$.
- For all $\theta \in \mathbb{R}$, $\int_{\mathbb{R}^n} f(\theta, x) dx$ is finite.
- For all $\theta \in \mathbb{R}$, there is an open $U \subseteq \mathbb{R}$ such that $\theta \in U$ and $\int_{\mathbb{R}^n} \text{Lip}(f(-, x)|_U) dx$ is finite.

Here $\nabla_\theta(-)$ and $\text{Lip}(-)$ are defined as in Theorem D.2, and “almost all” is with respect to the Lebesgue measure. Then, for all $\theta \in \mathbb{R}$, both sides of the following are well-defined and equal:

$$\nabla_\theta \int_{\mathbb{R}^n} f(\theta, x) dx = \int_{\mathbb{R}^n} \nabla_\theta f(\theta, x) dx.$$

PROOF. Before starting the main proof, we show that for any open $U' \subseteq \mathbb{R}$, the following function $L : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{\infty\}$ is measurable:

$$L(x) \triangleq \text{Lip}(f(-, x)|_{U'}).$$

Define $V \subseteq \mathbb{R}^2$ and $\ell : V \times \mathbb{R}^n \rightarrow \mathbb{R}$ as

$$V \triangleq \{v \in U' \times U' \mid v_1 \neq v_2\} \subseteq \mathbb{R}^2, \quad \ell(v, x) \triangleq \frac{|f(v_1, x) - f(v_2, x)|}{|v_1 - v_2|}.$$

Let V' be a countable, dense subset of V . Then, for all $x \in \mathbb{R}^n$,

$$L(x) = \sup_{v \in V} \ell(v, x) = \sup_{v' \in V'} \ell(v', x),$$

where the first equality is by the definition of $\text{Lip}(-)$, and the second equality holds since $\ell(-, x) : V \rightarrow \mathbb{R}$ is continuous for all $x \in \mathbb{R}^n$ by the first condition of this theorem. Since L is the supremum of countably many measurable functions $\{\ell(v', -) : \mathbb{R}^n \rightarrow \mathbb{R} \mid v' \in V'\}$, L itself is a measurable function as desired. Note that the measurability of L ensures that the integral $\int_{\mathbb{R}^n} L(x) dx$ in the fourth condition of this theorem is well-defined (as a value in $\mathbb{R} \cup \{\infty\}$).

We now start the main proof. Pick any $\theta' \in \mathbb{R}$. Define $g : \mathbb{R} \setminus \{\theta'\} \rightarrow \mathbb{R}$ as

$$g(\theta) \triangleq \int_{\mathbb{R}^n} \frac{f(\theta, x) - f(\theta', x)}{\theta - \theta'} dx,$$

where $g(\theta)$ is finite by the third condition of this theorem. Then,

$$\left(\nabla_\theta \int_{\mathbb{R}^n} f(\theta, x) dx \right) \Big|_{\theta=\theta'} = \lim_{\theta \rightarrow \theta'} \frac{1}{\theta - \theta'} \left(\int_{\mathbb{R}^n} f(\theta, x) dx - \int_{\mathbb{R}^n} f(\theta', x) dx \right) = \lim_{\theta \rightarrow \theta'} g(\theta),$$

where each equality denotes that LHS is well-defined if and only if RHS is well-defined, and if so, LHS and RHS are equal. Here the first equality is from the definition of ∇_θ , and the second equality from the third condition of this theorem. Note that the following are equivalent for any $r \in \mathbb{R}$:

- $\lim_{\theta \rightarrow \theta'} g(\theta) = r$.
- For any $\{\theta_i\}_{i \in \mathbb{N}} \subseteq \mathbb{R} \setminus \{\theta'\}$, $\lim_{i \rightarrow \infty} \theta_i = \theta'$ implies $\lim_{i \rightarrow \infty} g(\theta_i) = r$.

So it suffices to show that (ii) holds with an appropriate choice of r .

To do so, consider $\{\theta_i\}_{i \in \mathbb{N}} \subseteq \mathbb{R} \setminus \{\theta'\}$ such that $\lim_{i \rightarrow \infty} \theta_i = \theta'$. Define $h_i : \mathbb{R}^n \rightarrow \mathbb{R}$ as

$$h_i(x) \triangleq \frac{f(\theta_i, x) - f(\theta', x)}{\theta_i - \theta'}.$$

Then, $g(\theta_i) = \int_{\mathbb{R}^n} h_i(x) dx$ by the definition of h_i , and $\lim_{i \rightarrow \infty} h_i(x) = (\nabla_{\theta} f(\theta, x))|_{\theta=\theta'}$ for almost all $x \in \mathbb{R}^n$ by the second condition of this theorem. So, if the dominated convergence theorem is applicable to $\lim_{i \rightarrow \infty} \int_{\mathbb{R}^n} h_i(x) dx$, we would have

$$\lim_{i \rightarrow \infty} g(\theta_i) = \lim_{i \rightarrow \infty} \int_{\mathbb{R}^n} h_i(x) dx = \int_{\mathbb{R}^n} \left(\nabla_{\theta} f(\theta, x) \right) \Big|_{\theta=\theta'} dx,$$

where each equality denotes that both LHS and RHS are well-defined and are equal. Therefore, it suffices to show that the preconditions of the dominated convergence theorem for $\lim_{i \rightarrow \infty} \int_{\mathbb{R}^n} h_i(x) dx$ are satisfied.

We show that the preconditions indeed hold, which concludes the proof:

- “ h_i is measurable for all $i \in \mathbb{N}$ ”: This holds by the measurability of f .
- “ $\lim_{i \rightarrow \infty} h_i(x) = (\nabla_{\theta} f(\theta, x))|_{\theta=\theta'}$ for almost all $x \in \mathbb{R}^n$ ”: This was already shown above.
- “There exist $H : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{\infty\}$ and $I \in \mathbb{N}$ such that (a) $\int_{\mathbb{R}^n} H(x) dx$ is finite, and (b) for all $i \geq I$, $|h_i(x)| \leq H(x)$ for almost all $x \in \mathbb{R}^n$ ”: By the fourth condition of this theorem, there is an open $U \subseteq \mathbb{R}$ such that $\theta' \in U$ and $\int_{\mathbb{R}^n} \text{Lip}(f(-, x)|_U) dx$ is finite. Let $H : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{\infty\}$ be

$$H(x) \triangleq \text{Lip}(f(-, x)|_U).$$

Then, (a) clearly holds since $\int_{\mathbb{R}^n} \text{Lip}(f(-, x)|_U) dx$ is finite. Further, (b) holds as follows: by $\lim_{i \rightarrow \infty} \theta_i = \theta'$ and U being an open neighborhood of θ' , there is $I \in \mathbb{N}$ such that $\theta_i \in U$ for all $i \geq I$; therefore, for all $i \geq I$,

$$|h_i(x)| = \frac{|f(\theta_i, x) - f(\theta', x)|}{|\theta_i - \theta'|} \leq \text{Lip}(f(-, x)|_U) = H(x) \quad \text{for all } x \in \mathbb{R}^n,$$

where the inequality holds by the definition of $\text{Lip}(-)$ with $\theta_i, \theta' \in U$ and $\theta_i \neq \theta'$. \square

Remark E.3. The second condition (*) of Theorem E.2 is weaker than the following, corresponding condition (*') of a standard theorem for interchanging differentiation and integral (e.g., [Bogachev 2007, Corollary 2.8.7]): “for almost all $x \in \mathbb{R}^n$, $\nabla_{\theta} f(\theta, x)$ is well-defined for all $\theta \in \mathbb{R}$.” The difference arises from whether a proof uses the mean value theorem or not: a proof of the standard theorem finds the function H for the dominated convergence theorem, by applying the mean value theorem which requires the stronger condition (*'); the proof of Theorem E.2 finds H not by applying the mean value theorem (but by using the fourth condition of the theorem), so the weaker condition (*) is sufficient for the proof. In this sense, Theorem E.2 is close to [Bogachev 2007, Exercise 2.12.68]. \square

LEMMA E.4. *Let $f : X_1 \rightarrow X_2$ and $g : X_2 \rightarrow X_3$ for some open sets $X_i \subseteq \mathbb{R}^{n_i}$. Suppose that f is locally Lipschitz and g is differentiable. Then, $g \circ f : X_1 \rightarrow X_3$ is differentiable almost everywhere and the chain rule for $g \circ f$ holds almost everywhere, i.e.,*

$$D(g \circ f)(x) = D(g)(f(x)) \cdot D(f)(x)$$

for almost every $x \in X_1$. Here we use the Lebesgue measure as an underlying measure.

PROOF. Since local Lipschitzness is preserved under a function composition, $g \circ f$ is locally Lipschitz and thus differentiable almost everywhere. Since f is also differentiable almost everywhere and g is differentiable everywhere, the set

$$\begin{aligned} U &= X_1 \setminus \{x \in X_1 \mid (g \circ f \text{ is differentiable at } x) \\ &\quad \wedge (g \text{ is differentiable at } f(x)) \\ &\quad \wedge (f \text{ is differentiable at } x)\} \end{aligned}$$

has Lebesgue measure zero. Note that the differentiability of g is importantly used here; if g is non-differentiable even at a point, U can have positive measure. The chain rule for $g \circ f$ holds for each $x \in U$ and this concludes the proof. \square

E.2 Proof of Theorem 4.6

PROOF OF THEOREM 4.6. The proof is essentially the same as the proof of Theorem 4.5, except that we invoke the following properties of local Lipschitzness (instead of differentiability): the composition of locally Lipschitz functions is again locally Lipschitz, and the differentiation rules for $+$, \times , and \log hold almost everywhere for locally Lipschitz functions (Lemma E.1). \square

F DEFERRED RESULTS IN §5.2

F.1 Proof of Theorem 5.6

PROOF OF THEOREM 5.6. We prove the theorem by induction on the structure of c . Let $(p, d, V) \triangleq \llbracket c \rrbracket^\sharp$, and pick $v \in \text{Var}$. We have to show that $p(v) \supseteq d(v)^c$ and $d(v) \supseteq V$. We call these two requirements as conditions (i) and (ii).

Case $c \equiv \text{skip}$. In this case, $p(v) = \text{Var}$ and $V = \emptyset$, from which the conditions (i) and (ii) follow.

Case $c \equiv x := e$. In this case, $V = \emptyset$. So, the condition (ii) holds. For the proof of the condition (i), we do case analysis on v . If v is the updated variable x , we have $p(v) = \llbracket e \rrbracket^\sharp$ and $d(v) = \text{fv}(e)$. Since $\llbracket e \rrbracket^\sharp \supseteq \text{fv}(e)^c$, the condition holds. If v is different from x , then $p(v)$ is Var , and so it includes $d(v)^c$.

Case $c \equiv \text{obs}(\text{dist}_N(e_1, e_2), r)$. The proof of this case is similar to the one for the assignments. Since $V = \emptyset$, the condition (ii) holds. If v is the variable *like*, then

$$p(v) = \llbracket \text{like} \times \text{pdf}_N(r; e_1, e_2) \rrbracket^\sharp \supseteq \text{fv}(\text{like} \times \text{pdf}_N(r; e_1, e_2))^c = \left(\{\text{like}\} \cup \text{fv}(e_1) \cup \text{fv}(e_2) \right)^c = d(v)^c.$$

So, the condition (i) holds in this case. If v is not the variable *like*, then $p(v) = \text{Var}$, from which the condition (i) follows.

Case $c \equiv x := \text{sam}(\text{name}(\alpha, e), \text{dist}_N(e_1, e_2), \lambda y. e')$. In this case, $V = \emptyset$, from which the condition (ii) follows. We do case analysis on whether e is a real constant r or not. During the case analysis, we use the assumption that $\text{fv}(e)^c \subseteq \llbracket e \rrbracket^\sharp$ for all e , without mentioning it explicitly.

First, we deal with the case that $e \equiv r$. Let $\mu \triangleq \text{create_name}(\alpha, r)$. If v is none of x , val_μ , pr_μ , and cnt_μ , we have $p(v) = \text{Var}$, which gives the condition (i). If $v \in \{x, \text{val}_\mu\}$, we prove the condition (i) as follows:

$$d(v)^c = \text{fv}(e'[\mu/y])^c \subseteq \llbracket e'[\mu/y] \rrbracket^\sharp = p(v).$$

If $v \equiv \text{pr}_\mu$, we calculate the condition (i) as follows:

$$d(\text{pr}_\mu)^c = (\{\mu\} \cup \text{fv}(e_1) \cup \text{fv}(e_2))^c = \text{fv}(\text{pdf}_N(\mu; e_1, e_2))^c \subseteq \llbracket \text{pdf}_N(\mu; e_1, e_2) \rrbracket^\sharp = p(\text{pr}_\mu).$$

If $v \equiv \text{cnt}_\mu$, we derive the condition (i) as follows:

$$d(\text{cnt}_\mu)^c = \{\text{cnt}_\mu\}^c = \text{fv}(\text{cnt}_\mu + 1)^c \subseteq \llbracket \text{cnt}_\mu + 1 \rrbracket^\sharp = p(\text{cnt}_\mu).$$

Next, we handle the case that e is not a real constant. If v is none of x , val_μ , pr_μ , and cnt_μ for some $\mu = (\alpha, _)$, we have $p(v) = \text{Var}$, which implies the condition (i). If $v \equiv x$, we show the condition (i) as follows:

$$p(v) = \left(\text{fv}(e)^c \cap \bigcap_{\mu=(\alpha, _) \in \text{Name}} \llbracket e'[\mu/y] \rrbracket^\sharp \right) \supseteq \left(\text{fv}(e) \cup \bigcup_{\mu=(\alpha, _) \in \text{Name}} \text{fv}(e'[\mu/y]) \right)^c = d(v)^c.$$

If $v \equiv \text{val}_\mu$ for some $\mu = (\alpha, _)$, we calculate the condition (i) as follows:

$$p(v) = \left(\text{fv}(e)^c \cap \llbracket e'[\mu/y] \rrbracket^\# \right) \supseteq \left(\text{fv}(e) \cup \{ \text{val}_\mu \} \cup \text{fv}(e'[\mu/y]) \right)^c = d(v)^c.$$

If $v \equiv \text{pr}_\mu$ for some $\mu = (\alpha, _)$, we prove the condition (i) as follows:

$$\begin{aligned} p(v) &= \text{fv}(e)^c \cap \llbracket \text{pdf}_N(\mu; e_1, e_2) \rrbracket^\# \supseteq \text{fv}(e)^c \cap \text{fv}(\text{pdf}_N(\mu; e_1, e_2))^c \\ &\supseteq \left(\text{fv}(e) \cup \{ \mu, \text{pr}_\mu \} \cup \text{fv}(e_1) \cup \text{fv}(e_2) \right)^c = p(\text{pr}_\mu). \end{aligned}$$

Finally, if $v \equiv \text{cnt}_\mu$ for some $\mu = (\alpha, _)$, we show the condition (i) as follows:

$$p(v) = \text{fv}(e)^c \cap \llbracket \text{cnt}_\mu + 1 \rrbracket^\# \supseteq \text{fv}(e)^c \cap \text{fv}(\text{cnt}_\mu + 1)^c = (\text{fv}(e) \cup \{ \text{cnt}_\mu \})^c = d(v)^c.$$

Case $c \equiv c'; c''$. Let $(p', d', V') \triangleq \llbracket c' \rrbracket^\#$ and $(p'', d'', V'') \triangleq \llbracket c'' \rrbracket^\#$. The condition (ii) holds since

$$d(v) = V' \cup d'_\cup(d''(v)) \supseteq V' \cup d'_\cup(V'') = V.$$

For the condition (ii), we prove the required subset relationship as follows:

$$\begin{aligned} d(v)^c &= (V' \cup d'_\cup(d''(v)))^c = (V')^c \cap \bigcap_{w \in d''(v)} d'(w)^c \\ &\subseteq (V')^c \cap \bigcap_{w \in d''(v)} p'(w) \cap \bigcap_{w \in p''(v)^c} d'(w)^c \\ &= \left(V' \cup p'_\cap(d''(v))^c \cup d'_\cup(p''(v)^c) \right)^c = p(v). \end{aligned}$$

The subset relationship in the above derivation holds because $d'(w)^c \subseteq p'(w)$ and $d''(v) \supseteq p''(v)^c$ by induction hypothesis.

Case $c \equiv \text{if } b \{c'\} \text{ else } \{c''\}$. Let $(p', d', V') \triangleq \llbracket c' \rrbracket^\#$ and $(p'', d'', V'') \triangleq \llbracket c'' \rrbracket^\#$. Then, by induction hypothesis,

$$V = \left(\text{fv}(b) \cup V' \cup V'' \right) \subseteq \left(\text{fv}(b) \cup d'(v) \cup d''(v) \right) = d(v),$$

which implies the condition (ii). Also, by induction hypothesis again,

$$d(v)^c = \left(\text{fv}(b)^c \cap d'(v)^c \cap d''(v)^c \right) \subseteq \left(\text{fv}(b)^c \cap p'(v) \cap p''(v) \right) = p(v),$$

which shows the condition (ii).

Case $c \equiv \text{while } b \{c'\}$. Let $(p', d', V') \triangleq \llbracket c' \rrbracket^\#$, and $F^\#$ be the operator in the abstract semantics of c . Note that the abstract domain $\mathcal{D}^\#$ contains $(p_\perp, d_\perp, V_\perp) = ((\lambda v. \text{Var}), (\lambda v. \emptyset), \emptyset)$. Thus, it is sufficient to show that $F^\#$ is a well-defined monotone function on $\mathcal{D}^\#$, because then the least fixed point of $F^\#$ is also in $\mathcal{D}^\#$ and satisfies the conditions (i) and (ii). The monotonicity of $F^\#$ holds because when $(p_1, d_1, V_1) \triangleq F^\#(p_0, d_0, V_0)$, the inputs p_0, d_0 , and V_0 are used in the right polarity in the definitions of p_1, d_1 , and V_1 ; for instance, p_0 is used only in the positive position (with respect to the subset order) when it is used to define p_1 . To prove well-definedness of $F^\#$, assume that $p_0(v_0) \supseteq d_0(v_0)^c$ and $d_0(v_0) \supseteq V_0$ for all $v_0 \in \text{Var}$, and pick a variable $v_1 \in \text{Var}$. Then, since $V_0 \subseteq d_0(v_1)$,

$$V_1 = \left(\text{fv}(b) \cup d'_\cup(V_0) \cup V' \right) \subseteq \left(\text{fv}(b) \cup d'_\cup(d_0(v_1)) \cup V' \cup \{v_1\} \right) = d_1(v_1).$$

Also, by the induction hypothesis on the loop body c' and the relationship $d_0(v_1) \supseteq p_0(v_1)^c$,

$$d_1(v_1)^c = \text{fv}(b)^c \cap (V')^c \cap \bigcap_{w \in d_0(v_1)} d'(w)^c \cap \{v_1\}^c$$

$$\begin{aligned}
& \subseteq f\nu(b)^c \cap (V')^c \cap \bigcap_{w \in d_0(v_1)} p'(w) \cap \bigcap_{w \in p_0(v_1)^c} d'(w)^c \\
& = f\nu(b)^c \cap \left(V' \cup p'_\cap(d_0(v_1))^c \cup d'_\cup(p_0(v_1)^c) \right)^c = p_1(v_1).
\end{aligned}$$

Thus, (p_1, d_1, V_1) is also in \mathcal{D}^\sharp . \square

F.2 Proof of Theorem 5.8

Our program analysis consists of two parts, one for tracking the dependency information and the other for tracking the smoothness information. The first part does not depend on the second, although it is used crucially by the second part. We exploit this one-way relationship between the two parts of our analysis, and prove the soundness of the dependency-tracking part first and then that of the other smoothness-tracking part. Consider a command c , and let $(p, d, V) \triangleq \llbracket c \rrbracket^\sharp$. Then, we have:

THEOREM F.1. *For all $v \in \text{Var}$, we have $\models \Delta(\llbracket c \rrbracket, d(v), \{v\})$. Also, $\models \Delta(\llbracket c \rrbracket, V, \emptyset)$.*

THEOREM F.2. *For all $v \in \text{Var}$, we have $\models \Phi(\llbracket c \rrbracket, p(v), \{v\})$.*

We prove the two soundness results in §F.3 and §F.4. From these, we immediately obtain the main soundness theorem:

PROOF OF THEOREM 5.8. Let c be a command and $(p, d, V) = \llbracket c \rrbracket^\sharp$. Then, by Theorems F.1 and F.2, we have $\models \Delta(\llbracket c \rrbracket, d(v), \{v\})$, $\models \Delta(\llbracket c \rrbracket, V, \emptyset)$, and $\models \Phi(\llbracket c \rrbracket, p(v), \{v\})$ for all $v \in \text{Var}$. Hence, by the definition of γ (i.e., Eq. (10)), we have $\llbracket c \rrbracket \in \gamma(\llbracket c \rrbracket^\sharp)$ as desired. \square

F.3 Proof of Theorem F.1

PROOF OF THEOREM F.1. We prove the theorem by induction on the structure of c . Let $(p, d, V) \triangleq \llbracket c \rrbracket^\sharp$. Pick a variable $v \in \text{Var}$ and states $\sigma, \sigma', \sigma_0, \sigma'_0 \in \text{St}$ such that

$$\sigma \sim_{d(v)} \sigma' \text{ and } \sigma_0 \sim_V \sigma'_0.$$

We will show that (i) if $\llbracket c \rrbracket \sigma_0 \in \text{St}$, then $\llbracket c \rrbracket \sigma'_0 \in \text{St}$, and (ii) if $\llbracket c \rrbracket \sigma \in \text{St}$ and $\llbracket c \rrbracket \sigma' \in \text{St}$, then $\llbracket c \rrbracket \sigma \sim_{\{v\}} \llbracket c \rrbracket \sigma'$, i.e., $\llbracket c \rrbracket \sigma(v) = \llbracket c \rrbracket \sigma'(v)$. Since $V \subseteq d(v)$, these two imply the claim of the theorem. We refer to these two properties as conditions (i) and (ii) in the rest of the proof.

Case $c \equiv \text{skip}$. In this case, $d(v) = \{v\}$ and $V = \emptyset$. The condition (i) holds since skip always terminates. The condition (ii) also holds because $\llbracket c \rrbracket \sigma'' = \sigma''$ for all σ'' , and the relation $\sim_{d(v)}$ coincides with $\sim_{\{v\}}$.

Case $c \equiv (x := e)$. In this case, $V = \emptyset$, and the condition (i) holds since the assignments always terminate. For the condition (ii), we do case analysis on the variable v .

- Case $v \equiv x$. In this case, $d(v) = f\nu(e)$. This implies $\llbracket e \rrbracket \sigma = \llbracket e \rrbracket \sigma'$. Thus, $\llbracket c \rrbracket \sigma(x) = \llbracket e \rrbracket \sigma = \llbracket e \rrbracket \sigma' = \llbracket c \rrbracket \sigma'(x)$. This implies the desired $\llbracket c \rrbracket \sigma \sim_{\{x\}} \llbracket c \rrbracket \sigma'$.
- Case $v \not\equiv x$. In this case, $d(v) = \{v\}$, and so $\sigma(v) = \sigma'(v)$. This implies that $\llbracket c \rrbracket \sigma(v) = \sigma(v) = \sigma'(v) = \llbracket c \rrbracket \sigma'(v)$, which gives the desired relationship.

Case $c \equiv \text{obs}(\text{dist}_N(e_1, e_2), r)$. The observe commands always terminate. Thus, the condition (i) holds. We prove the condition (ii) by case analysis on the variable v . If v is not *like*, then $d(v) = \{v\}$, $\llbracket c \rrbracket \sigma(v) = \sigma(v)$, and $\llbracket c \rrbracket \sigma'(v) = \sigma'(v)$. Thus, in this case, the assumption $\sigma \sim_{d(v)} \sigma'$ implies $\llbracket c \rrbracket \sigma(v) = \llbracket c \rrbracket \sigma'(v)$, as desired. If v is *like*, then $d(v) = f\nu(e_1) \cup f\nu(e_2) \cup \{\text{like}\}$, and for some function $g : \mathbb{R}^4 \rightarrow \mathbb{R}$,

$$\llbracket c \rrbracket \sigma(v) = g(\sigma(\text{like}), r, \llbracket e_1 \rrbracket \sigma, \llbracket e_2 \rrbracket \sigma) \text{ and } \llbracket c \rrbracket \sigma'(v) = g(\sigma'(\text{like}), r, \llbracket e_1 \rrbracket \sigma', \llbracket e_2 \rrbracket \sigma').$$

Therefore, from the assumption $\sigma \sim_{d(v)} \sigma'$, it follows that $\llbracket c \rrbracket \sigma(v) = \llbracket c \rrbracket \sigma'(v)$, as desired.

Case $c \equiv (x := \text{sam}(n, \text{dist}_N(e_1, e_2), \lambda y. e'))$. The sample commands always terminate. So, the condition (i) holds. We prove the condition (ii) by case analysis on n .

The first case is that n is a constant expression, i.e., it is an expression of the form $\text{name}(\alpha, r)$ for some $\alpha \in \text{Str}$ and real number r . Let $\mu \triangleq \text{create_name}(\alpha, r)$. If v is not one of x , val_μ , and pr_μ , then $d(v) = \{v\}$, and $\llbracket c \rrbracket \sigma(v) = g(\sigma(v))$ and $\llbracket c \rrbracket \sigma'(v) = g(\sigma'(v))$ for some function $g : \mathbb{R} \rightarrow \mathbb{R}$, so that the assumption $\sigma \sim_{d(v)} \sigma'$ implies that $\llbracket c \rrbracket \sigma(v) = \llbracket c \rrbracket \sigma'(v)$, as desired. If v is x or val_μ , then $d(v) = \text{fv}(e'[\mu/y])$, $\llbracket c \rrbracket \sigma(v) = \llbracket e'[\mu/y] \rrbracket \sigma$, and $\llbracket c \rrbracket \sigma'(v) = \llbracket e'[\mu/y] \rrbracket \sigma'$, so that the required $\llbracket c \rrbracket \sigma(v) = \llbracket c \rrbracket \sigma'(v)$ holds. Finally, if $v = \text{pr}_\mu$, then $d(v) = \{\mu\} \cup \text{fv}(e_1) \cup \text{fv}(e_2)$, and so, the assumption $\sigma \sim_{d(v)} \sigma'$ implies that

$$\llbracket c \rrbracket \sigma(\text{pr}_\mu) = \llbracket \text{pdf}_N(\mu; e_1, e_2) \rrbracket \sigma = \llbracket \text{pdf}_N(\mu; e_1, e_2) \rrbracket \sigma' = \llbracket c \rrbracket \sigma'(\text{pr}_\mu),$$

which is precisely the equality that we want.

The next case is that n is not a constant expression. Let $\text{name}(\alpha, e)$ be the form of n . If v is not one of x , val_μ , pr_μ , and cnt_μ for some μ of the form $(\alpha, _)$, then $d(v) = \{v\}$, $\llbracket c \rrbracket \sigma(v) = \sigma(v)$, and $\llbracket c \rrbracket \sigma'(v) = \sigma'(v)$, so that the assumption $\sigma \sim_{d(v)} \sigma'$ implies the desired $\llbracket c \rrbracket \sigma(v) = \llbracket c \rrbracket \sigma'(v)$. Assume that v is one of x , val_μ , pr_μ , and cnt_μ for some μ with $\mu = (\alpha, _)$. Let $\mu_0 \triangleq \llbracket n \rrbracket \sigma$ and $\mu'_0 \triangleq \llbracket n \rrbracket \sigma'$. Since $d(v) \supseteq \text{fv}(n)$ in this case, the assumption $\sigma \sim_{d(v)} \sigma'$ ensures that $\mu_0 = \mu'_0$. If v is x , then $\text{fv}(e'[\mu_0/y]) \subseteq d(v)$, so that the assumption $\sigma \sim_{d(v)} \sigma'$ gives the desired

$$\llbracket c \rrbracket \sigma(v) = \llbracket e'[\mu_0/y] \rrbracket \sigma = \llbracket e'[\mu'_0/y] \rrbracket \sigma' = \llbracket c \rrbracket \sigma'(v).$$

If v is cnt_μ for some μ of the form $(\alpha, _)$, then $\text{cnt}_\mu \in d(v)$, and $\llbracket c \rrbracket \sigma(\text{cnt}_\mu) = g(\sigma(\text{cnt}_\mu))$ and $\llbracket c \rrbracket \sigma'(\text{cnt}_\mu) = g(\sigma'(\text{cnt}_\mu))$ for some function $g : \mathbb{R} \rightarrow \mathbb{R}$, so that $\llbracket c \rrbracket \sigma(v) = \llbracket c \rrbracket \sigma'(v)$, as desired. If v is val_μ for $\mu = (\alpha, _)$, then $d(v) \supseteq \{\text{val}_\mu\} \cup \text{fv}(e'[\mu/y])$, and $\llbracket c \rrbracket \sigma(\text{val}_\mu) = h(\llbracket e'[\mu/y] \rrbracket \sigma, \sigma(\text{val}_\mu))$ and $\llbracket c \rrbracket \sigma'(\text{val}_\mu) = h(\llbracket e'[\mu/y] \rrbracket \sigma', \sigma'(\text{val}_\mu))$ for some $h : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$, so that $\llbracket c \rrbracket \sigma(\text{val}_\mu) = \llbracket c \rrbracket \sigma'(\text{val}_\mu)$ as desired. Finally, if v is pr_μ for some μ of the form $(\alpha, _)$, then $d(v) \supseteq \{\text{pr}_\mu, \mu\} \cup \text{fv}(e_1) \cup \text{fv}(e_2)$, and for some $k : \mathbb{R}^4 \rightarrow \mathbb{R}$,

$$\llbracket c \rrbracket \sigma(v) = k(\sigma(\text{pr}_\mu), \sigma(\mu), \llbracket e_1 \rrbracket \sigma, \llbracket e_2 \rrbracket \sigma) \text{ and } \llbracket c \rrbracket \sigma'(v) = k(\sigma'(\text{pr}_\mu), \sigma'(\mu), \llbracket e_1 \rrbracket \sigma', \llbracket e_2 \rrbracket \sigma'),$$

so that the assumption $\sigma \sim_{d(v)} \sigma'$ guarantees that $\llbracket c \rrbracket \sigma(v) = \llbracket c \rrbracket \sigma'(v)$, as desired.

Case $c \equiv (c'; c'')$. Let $(p', d', V') \triangleq \llbracket c' \rrbracket^\sharp$ and $(p'', d'', V'') \triangleq \llbracket c'' \rrbracket^\sharp$. Recall that

$$d(v) = V' \cup (d') \cup (d''(v)) = V' \cup \bigcup \{d'(v'') \mid v'' \in d''(v)\} \text{ and } V = V' \cup (d') \cup (V'').$$

Let us handle the condition (i) first. Since $\llbracket c'; c'' \rrbracket \sigma_0 \in \text{St}$, we have $\llbracket c' \rrbracket \sigma_0 \in \text{St}$. But $\sigma_0 \sim_{V'} \sigma'_0$, because σ_0 and σ'_0 are \sim_V -related and V includes V' . Thus, $\llbracket c' \rrbracket \sigma'_0 \in \text{St}$ as well by induction hypothesis, and it is sufficient to show $\llbracket c' \rrbracket \sigma_0 \sim_{V''} \llbracket c' \rrbracket \sigma'_0$. Note that for every $v'' \in V''$, by the definition of V , we have $V \supseteq d'(v'')$, and so $\sigma_0 \sim_{d'(v'')} \sigma'_0$, which implies, by induction hypothesis, that $\llbracket c' \rrbracket \sigma_0 \sim_{\{v''\}} \llbracket c' \rrbracket \sigma'_0$. As a result, we have the desired $\llbracket c' \rrbracket \sigma_0 \sim_{V''} \llbracket c' \rrbracket \sigma'_0$.

Next, we deal with the condition (ii). Since $\llbracket c'; c'' \rrbracket \sigma$ and $\llbracket c'; c'' \rrbracket \sigma'$ are both in St , there exist states σ_1, σ'_1 such that $\llbracket c' \rrbracket \sigma = \sigma_1$ and $\llbracket c' \rrbracket \sigma' = \sigma'_1$. We apply the induction hypothesis to c' and get $\sigma_1 \sim_{d''(v)} \sigma'_1$. Since $\llbracket c'' \rrbracket \sigma_1$ and $\llbracket c'' \rrbracket \sigma'_1$ are in St , we apply the induction hypothesis again but this time to c'' , σ_1 , and σ'_1 , and obtain $\llbracket c'' \rrbracket \sigma_1 \sim_{\{v\}} \llbracket c'' \rrbracket \sigma'_1$, which implies the desired

$$\llbracket c \rrbracket \sigma(v) = \llbracket c'' \rrbracket \sigma_1(v) = \llbracket c'' \rrbracket \sigma'_1(v) = \llbracket c \rrbracket \sigma'(v).$$

Case $c \equiv (\text{if } b \{c'\} \text{ else } \{c''\})$. Let $(p', d', V') \triangleq \llbracket c' \rrbracket^\sharp$ and $(p'', d'', V'') \triangleq \llbracket c'' \rrbracket^\sharp$. Then, $d(v) = \text{fv}(b) \cup d'(v) \cup d''(v)$ and $V = \text{fv}(b) \cup V' \cup V''$.

We prove the condition (i) under the assumption that $\llbracket b \rrbracket \sigma_0 = \text{true}$. Essentially the same proof applies to the other case that $\llbracket b \rrbracket \sigma_0 = \text{false}$. Since V includes $\text{fv}(b)$, we also have $\llbracket b \rrbracket \sigma'_0 = \text{true}$. Furthermore, since $V' \subseteq V$ and so $\sigma_0 \sim_{V'} \sigma'_0$ by the induction hypothesis, we get that $\llbracket c' \rrbracket \sigma'_0 \in \text{St}$.

Next we show the condition (ii) under the assumption that $\llbracket b \rrbracket \sigma = \text{true}$. As before, the proof of the other case $\llbracket b \rrbracket \sigma = \text{false}$ is essentially the same. Since $d(v)$ includes $\text{fv}(b)$ and $d'(v)$, we have $\llbracket b \rrbracket \sigma' = \text{true}$ and $\sigma \sim_{d'(v)} \sigma'$. Also, because $\llbracket c \rrbracket \sigma = \llbracket c' \rrbracket \sigma$ and $\llbracket c \rrbracket \sigma' = \llbracket c' \rrbracket \sigma'$, both $\llbracket c' \rrbracket \sigma$ and $\llbracket c' \rrbracket \sigma'$ are in St . Thus, by induction hypothesis, $\llbracket c' \rrbracket \sigma \sim_{\{v\}} \llbracket c' \rrbracket \sigma'$, which implies that

$$\llbracket c \rrbracket \sigma(v) = \llbracket c' \rrbracket \sigma(v) = \llbracket c' \rrbracket \sigma'(v) = \llbracket c \rrbracket \sigma'(v),$$

as desired.

Case $c \equiv (\text{while } b \{c_0\})$. Let $(d_0, p_0, V_0) \triangleq \llbracket c_0 \rrbracket^\sharp$, and F^\sharp be the operator in the abstract semantics of c such that (p, d, V) is the least fixed point of F^\sharp . Also, let F be the operator in the concrete semantics of c such that $\llbracket c \rrbracket$ is the least fixed point of F . Now define

$$T \triangleq \{f \in [\text{St} \rightarrow \text{St}_\perp] \mid \text{for all } v \in \text{Var}, \models \Delta(f, d(v), \{v\}) \text{ and } \models \Delta(f, V, \emptyset)\}.$$

We will show that (i) T contains the empty function $\perp_{\text{St} \rightarrow \text{St}_\perp} \triangleq \lambda \sigma. \text{undefined}$, (ii) it is closed under the least upper bounds of increasing chains, and (iii) the function F maps functions in T to some functions in the same set. These three imply that the least fixed point of F , namely, $\llbracket c \rrbracket$, is in T , which gives the desired conclusion.

The membership of $\perp_{\text{St} \rightarrow \text{St}_\perp}$ to T is immediate, since we have $\models \Delta(\perp_{\text{St} \rightarrow \text{St}_\perp}, U, U')$ for all $U, U' \subseteq \text{Var}$.

To prove the next requirement, namely, the closure under the least upper bounds of increasing chains, consider a chain $(f_n)_{n \in \mathbb{N}}$ in T , i.e., a sequence such that $f_n(\sigma) = f_{n+1}(\sigma)$ for all $n \in \mathbb{N}$ and σ with $f_n(\sigma) \in \text{St}$. Let f_∞ be the least upper bound of the f_n 's (i.e., $f_\infty(\sigma) = f_n(\sigma)$ if $f_n(\sigma) \in \text{St}$ and $f_\infty(\sigma) = \perp$ if $f_n(\sigma) = \perp$ for all $n \in \mathbb{N}$). As in all the other cases so far, we pick an arbitrary variable $v \in \text{Var}$ and arbitrary states $\sigma_0, \sigma'_0, \sigma$, and σ' such that

$$\sigma_0 \sim_V \sigma'_0, \quad f_\infty(\sigma_0) \in \text{St}, \quad \sigma \sim_{d(v)} \sigma', \quad f_\infty(\sigma), f_\infty(\sigma') \in \text{St}.$$

We will show that $f_\infty(\sigma'_0) \in \text{St}$ and $f_\infty(\sigma) \sim_{\{v\}} f_\infty(\sigma')$, which correspond to what we have called conditions (i) and (ii) in the previous cases. Since $f_\infty(\sigma_0) \in \text{St}$, there exists $n \in \mathbb{N}$ such that $f_n(\sigma_0) \in \text{St}$. Because $\models \Delta(f_n, V, \emptyset)$ and $\sigma_0 \sim_V \sigma'_0$, we have $f_n(\sigma'_0) \in \text{St}$, which implies that $f_\infty(\sigma'_0) = f_n(\sigma'_0) \in \text{St}$, as desired. Our proof of the condition (ii) has a similar form. Since both $f_\infty(\sigma)$ and $f_\infty(\sigma')$ are in St , there exists $n \in \mathbb{N}$ such that $f_\infty(\sigma) = f_n(\sigma)$ and $f_\infty(\sigma') = f_n(\sigma')$. By assumption, $\sigma \sim_{d(v)} \sigma'$, and $f_n \in T$. Thus, $f_n(\sigma) \sim_{\{v\}} f_n(\sigma')$, which gives the desired $f_\infty(\sigma) \sim_{\{v\}} f_\infty(\sigma')$.

It remains to show the last requirement, i.e., the closure under F . Pick an arbitrary $f \in T$. Consider a variable $v \in \text{Var}$ and states $\sigma_0, \sigma'_0, \sigma$, and σ' such that

$$\sigma_0 \sim_V \sigma'_0, \quad F(f)(\sigma_0) \in \text{St}, \quad \sigma \sim_{d(v)} \sigma', \quad F(f)(\sigma), F(f)(\sigma') \in \text{St}.$$

We will show that $F(f)(\sigma'_0) \in \text{St}$ and $F(f)(\sigma) \sim_{\{v\}} F(f)(\sigma')$, while referring to these two desired properties as conditions (i) and (ii), as we have done before.

Let us handle the condition (i) first. If $\llbracket b \rrbracket \sigma_0 = \text{false}$, we have $\llbracket b \rrbracket \sigma'_0 = \text{false}$, because $\sigma_0 \sim_V \sigma'_0$ and $\text{fv}(b) \subseteq V$. Thus, in this case, $F(f)(\sigma'_0) = \sigma'_0 \in \text{St}$. If $\llbracket b \rrbracket \sigma_0 = \text{true}$, then $\llbracket b \rrbracket \sigma'_0$ is also true . Furthermore, in this case, by induction hypothesis, $\llbracket c_0 \rrbracket \sigma'_0 \in \text{St}$ since $V \supseteq V_0$, $\sigma_0 \sim_V \sigma'_0$, and $\llbracket c_0 \rrbracket \sigma_0 \in \text{St}$. Also, by induction hypothesis again, $\llbracket c_0 \rrbracket \sigma_0 \sim_V \llbracket c_0 \rrbracket \sigma'_0$, since $V \supseteq (d_0) \cup (V)$ and $\sigma_0 \sim_V \sigma'_0$. Since $f \in T$ and $f(\llbracket c_0 \rrbracket \sigma_0) \in \text{St}$, we have $f(\llbracket c_0 \rrbracket \sigma'_0) \in \text{St}$, which implies that $F(f)(\sigma'_0) \in \text{St}$, as desired.

Next, we prove the condition (ii). If $\llbracket b \rrbracket \sigma = \text{false}$, we have $\llbracket b \rrbracket \sigma' = \text{false}$ since $\text{fv}(b) \subseteq d(v)$ and $\sigma \sim_{d(v)} \sigma'$. Thus, in this case, $F(f)(\sigma) = \sigma$ and $F(f)(\sigma') = \sigma'$. Also, $\{v\} \subseteq d(v)$, and so, $\sigma \sim_{d(v)} \sigma'$ implies that $F(f)(\sigma) = \sigma \sim_{\{v\}} \sigma' = F(f)(\sigma')$, as desired. Now assume that $\llbracket b \rrbracket \sigma = \text{true}$. Then, $\llbracket b \rrbracket \sigma' = \text{true}$ by the reason that $\text{fv}(b) \subseteq d(v)$ and $\sigma \sim_{d(v)} \sigma'$. Also, $\llbracket c_0 \rrbracket \sigma$ and $\llbracket c_0 \rrbracket \sigma'$ are in St , so that $F(f)(\sigma) = f(\llbracket c_0 \rrbracket \sigma)$ and $F(f)(\sigma') = f(\llbracket c_0 \rrbracket \sigma')$. Furthermore, since $d(v) \supseteq (d_0) \cup (d(v))$ and

$\sigma \sim_{d(v)} \sigma'$, we have $\llbracket c_0 \rrbracket \sigma \sim_{d(v)} \llbracket c_0 \rrbracket \sigma'$. We then use the fact that $f \in T$ and $f(\llbracket c_0 \rrbracket \sigma), f(\llbracket c_0 \rrbracket \sigma') \in \text{St}$, and conclude that $f(\llbracket c_0 \rrbracket \sigma) \sim_{\{v\}} f(\llbracket c_0 \rrbracket \sigma')$, which gives the desired $F(f)(\sigma) \sim_{\{v\}} F(f)(\sigma')$. \square

F.4 Proof of Theorem F.2

Let seq be the following operator, which models sequential composition:

$$\begin{aligned} \text{seq} &: [\text{St} \rightarrow \text{St}_\perp] \times [\text{St} \rightarrow \text{St}_\perp] \rightarrow [\text{St} \rightarrow \text{St}_\perp] \\ \text{seq}(f, g) &\triangleq g^\dagger \circ f. \end{aligned}$$

Also, define an operator cond for modelling if commands as follows:

$$\begin{aligned} \text{cond} &: [\text{St} \rightarrow \mathbb{B}] \times [\text{St} \rightarrow \text{St}_\perp] \times [\text{St} \rightarrow \text{St}_\perp] \rightarrow [\text{St} \rightarrow \text{St}_\perp] \\ \text{cond}(h, f, g)(\sigma) &\triangleq \begin{cases} f(\sigma) & \text{if } h(\sigma) = \text{true}, \\ g(\sigma) & \text{if } h(\sigma) = \text{false}. \end{cases} \end{aligned}$$

PROOF OF THEOREM F.2. We prove the theorem by induction on the structure of c . Let $(p, d, V) \triangleq \llbracket c \rrbracket^\sharp$.

Case $c \equiv \text{skip}$. In this case, $\llbracket c \rrbracket(\sigma) = \sigma$ for all $\sigma \in \text{St}$, and $p(v) = \text{Var}$ for all $v \in \text{Var}$. To prove the conclusion, consider $v \in \text{Var}$ and $\tau \in \text{St}[p(v)^c] = \text{St}[\emptyset]$. We should show $g \in \phi_{p(v), \{v\}}$, where

$$g(\sigma) = \begin{cases} (\pi_{\text{Var}, \{v\}} \circ \llbracket c \rrbracket)(\sigma \oplus \tau) & \text{if } \llbracket c \rrbracket(\sigma \oplus \tau) \in \text{St} \\ \text{undefined} & \text{otherwise.} \end{cases}$$

Since $\llbracket c \rrbracket(\sigma \oplus \tau) = \llbracket c \rrbracket(\sigma) = \sigma \in \text{St}$ for all $\sigma \in \text{St}$, we have $g = \pi_{\text{Var}, \{v\}}$. Thus, Assumption 3 implies

$$g = \pi_{\text{Var}, \{v\}} \in \phi_{\text{Var}, \{v\}} = \phi_{p(v), \{v\}}.$$

Case $c \equiv (x := e)$. In this case, $\llbracket c \rrbracket(\sigma) = \sigma[x \mapsto \llbracket e \rrbracket \sigma]$ for all $\sigma \in \text{St}$. Also, $p(v) = \text{Var}$ if $v \neq x$, and $(e)^\sharp$ if $v \equiv x$. Consider $v \in \text{Var}$ and $\tau \in \text{St}[p(v)^c]$. We should show $g \in \phi_{p(v), \{v\}}$, where

$$g(\sigma) = \begin{cases} (\pi_{\text{Var}, \{v\}} \circ \llbracket c \rrbracket)(\sigma \oplus \tau) & \text{if } \llbracket c \rrbracket(\sigma \oplus \tau) \in \text{St} \\ \text{undefined} & \text{otherwise.} \end{cases}$$

If $v \neq x$, then $g(\sigma) = \pi_{\text{Var}, \{v\}}(\llbracket c \rrbracket(\sigma)) = \pi_{\text{Var}, \{v\}}(\sigma[x \mapsto \llbracket e \rrbracket \sigma]) = \pi_{\text{Var}, \{v\}}(\sigma)$ for all $\sigma \in \text{St}$, where the first equality uses $p(v) = \text{Var}$, and the last uses $v \neq x$. Hence, Assumption 3 implies

$$g = \pi_{\text{Var}, \{v\}} \in \phi_{\text{Var}, \{v\}} = \phi_{p(v), \{v\}}.$$

If $v \equiv x$, then $g(\sigma) = (\pi_{\text{Var}, \{x\}} \circ \llbracket c \rrbracket)(\sigma \oplus \tau) = \pi_{\text{Var}, \{x\}}((\sigma \oplus \tau)[x \mapsto \llbracket e \rrbracket(\sigma \oplus \tau)]) = [x \mapsto \llbracket e \rrbracket(\sigma \oplus \tau)]$ for all $\sigma \in \text{St}$. Since $\tau \in \text{St}[(\llbracket e \rrbracket)^\sharp]^c$ and $p(v) = (\llbracket e \rrbracket)^\sharp$, Assumption 2 implies

$$g = \lambda \sigma. [x \mapsto \llbracket e \rrbracket(\sigma \oplus \tau)] \in \phi_{(\llbracket e \rrbracket)^\sharp, \{x\}} = \phi_{p(v), \{v\}}.$$

Case $c \equiv (c'; c'')$. Let $(p', d', V') \triangleq \llbracket c' \rrbracket^\sharp$ and $(p'', d'', V'') \triangleq \llbracket c'' \rrbracket^\sharp$. Then,

$$\begin{aligned} p(v) &= \left(V' \cup (p') \cap (d''(v))^c \cup (d') \cup (p''(v)^c) \right)^c \\ &= \left((p') \cap (d''(v)) \right) \setminus \left(V' \cup (d') \cup (p''(v)^c) \right) \text{ for all } v \in \text{Var}. \end{aligned}$$

Also, we have $\llbracket c \rrbracket = \text{seq}(\llbracket c' \rrbracket, \llbracket c'' \rrbracket)$. To prove the conclusion, let $v \in \text{Var}$. It suffices to apply Lemma F.7 to $f = \llbracket c' \rrbracket, g = \llbracket c'' \rrbracket, K = p(v), L = d''(v) \cap p''(v), L' = d''(v)$, and $M = \{v\}$. What remains is to show the preconditions of the lemma:

- $\models \Phi(\llbracket c' \rrbracket, p(v), d''(v) \cap p''(v))$.
- $\models \Phi(\llbracket c'' \rrbracket, d''(v) \cap p''(v), \{v\})$.
- $\models \Delta(\llbracket c' \rrbracket, p(v)^c, d''(v) \setminus (d''(v) \cap p''(v)))$.

(d) $\models \Delta(\llbracket c'' \rrbracket, d''(v), \{v\})$.

We obtain (b) as follows: by induction hypothesis on c'' , we have $\models \Phi(\llbracket c'' \rrbracket, p''(v), \{v\})$, and by the weakening lemma for Φ (Lemma F.4), we have (b). We obtain (d) directly by Theorem F.1 on c'' . For (a), consider induction hypothesis on c' , which says that $\models \Phi(\llbracket c' \rrbracket, p'(w), \{w\})$ for all $w \in \text{Var}$. By the merging lemma for Φ (Lemma F.6), we have $\models \Phi(\llbracket c' \rrbracket, (p')_{\cap}(d''(v) \cap p''(v)), d''(v) \cap p''(v))$. Since

$$p(v) \subseteq (p')_{\cap}(d''(v)) \subseteq (p')_{\cap}(d''(v) \cap p''(v)),$$

we obtain (a) by the weakening lemma for Φ (Lemma F.4). For (c), observe that

$$p(v)^c \supseteq V' \cup (d')_{\cup}(p''(v)^c) \quad \text{and} \quad d''(v) \setminus (d''(v) \cap p''(v)) = p''(v)^c \quad (23)$$

where the second equality follows from $p''(v) \supseteq d''(v)^c$. By Theorem F.1 on c' , we have $\models \Delta(\llbracket c' \rrbracket, V', \emptyset)$ and $\models \Delta(\llbracket c' \rrbracket, d'(w), \{w\})$ for all $w \in \text{Var}$. If $p''(v)^c = \emptyset$, then $\models \Delta(\llbracket c' \rrbracket, V', p''(v)^c)$ holds, and if $p''(v)^c \neq \emptyset$, then $\models \Delta(\llbracket c' \rrbracket, (d')_{\cup}(p''(v)^c), p''(v)^c)$ holds by the merging lemma for Δ (Lemma F.5). By Eq. (23) and the weakening lemma for Δ (Lemma F.3), we obtain (c) for both cases. Note that we crucially used $p(v)^c \supseteq V'$ to handle the case $p''(v)^c = \emptyset$.

Case $c \equiv (\text{if } b \{c'\}, \text{else } \{c''\})$. Let $(p', d', V') \triangleq \llbracket c' \rrbracket^{\#}$ and $(p'', d'', V'') \triangleq \llbracket c'' \rrbracket^{\#}$. Then,

$$p(v) = \text{fv}(b)^c \cap p'(v) \cap p''(v) \quad \text{for all } v \in \text{Var}. \quad (24)$$

Also, $\llbracket c \rrbracket = \text{cond}(\llbracket b \rrbracket, \llbracket c' \rrbracket, \llbracket c'' \rrbracket)$. To prove the conclusion, let $v \in \text{Var}$. It suffices to apply Lemma F.8 to $f = \llbracket c' \rrbracket, g = \llbracket c'' \rrbracket, K = p(v), L = \{v\}$, and b . What remains is to show the preconditions of the lemma:

- (a) $\models \Phi(\llbracket c' \rrbracket, p(v), \{v\})$.
- (b) $\models \Phi(\llbracket c'' \rrbracket, p(v), \{v\})$.
- (c) $p(v)^c \supseteq \text{fv}(b)$.

We obtain (a) and (b) as follows: by induction hypothesis on c' and c'' , we have $\models \Phi(\llbracket c' \rrbracket, p'(v), \{v\})$ and $\models \Phi(\llbracket c'' \rrbracket, p''(v), \{v\})$, and by Eq. (24) and the weakening lemma for Φ (Lemma F.4), we have (a) and (b). We obtain (c) directly by Eq. (24).

Case $c \equiv (\text{while } b \{c_0\})$. The proof starts by decomposing $\llbracket c \rrbracket$ and $\llbracket c \rrbracket^{\#}$ into smaller pieces. Let

$$(p_0, d_0, V_0) \triangleq \llbracket c_0 \rrbracket^{\#}.$$

Define $F : [\text{St} \rightarrow \text{St}_{\perp}] \rightarrow [\text{St} \rightarrow \text{St}_{\perp}]$ and $F^{\#} : \mathcal{D}^{\#} \rightarrow \mathcal{D}^{\#}$ as in §3 and Fig. 3:

$$F(t)(\sigma) \triangleq \begin{cases} \sigma & \text{if } \llbracket b \rrbracket \sigma = \text{false} \\ t^{\dagger}(\llbracket c_0 \rrbracket \sigma) & \text{if } \llbracket b \rrbracket \sigma = \text{true}, \end{cases}$$

$$F^{\#}(p, d, V) \triangleq \left(\begin{array}{l} \lambda v. \text{fv}(b)^c \cap (V_0 \cup (p_0)_{\cap}(d(v))^c \cup (d_0)_{\cup}(p(v)^c))^c, \\ \lambda v. \text{fv}(b) \cup V_0 \cup (d_0)_{\cup}(d(v)) \cup \{v\}, \\ \text{fv}(b) \cup (d_0)_{\cup}(V) \cup V_0 \end{array} \right).$$

Define $t'_n \in [\text{St} \rightarrow \text{St}_{\perp}]$ and $(p'_n, d'_n, V'_n) \in \mathcal{D}^{\#}$ for $n \in \mathbb{N} \cup \{\infty\}$ as

$$t'_n \triangleq \begin{cases} \lambda \sigma. F^n(t_{\perp})(\sigma) & \text{if } n \in \mathbb{N} \\ \bigsqcup_{i \in \mathbb{N}} t'_i & \text{if } n = \infty, \end{cases} \quad (p'_n, d'_n, V'_n) \triangleq \begin{cases} (F^{\#})^n(p_{\perp}, d_{\perp}, V_{\perp}), & \text{if } n \in \mathbb{N} \\ \bigsqcup_{i \in \mathbb{N}} (p'_i, d'_i, V'_i) & \text{if } n = \infty, \end{cases}$$

where $t_{\perp} = \lambda \sigma. \perp$ and $(p_{\perp}, d_{\perp}, V_{\perp}) = (\lambda v. \text{Var}, \lambda v. \emptyset, \emptyset)$. Then, we have

$$\llbracket c \rrbracket = t'_{\infty} \quad \text{and} \quad \llbracket c \rrbracket^{\#} = (p'_{\infty}, d'_{\infty}, V'_{\infty}).$$

The proof is organized as follows. Define $T, T' \subseteq [\text{St} \rightarrow \text{St}_{\perp}]$ as

$$T \triangleq \{f \in [\text{St} \rightarrow \text{St}_{\perp}] \mid \forall v \in \text{Var}. \models \Delta(f, d'_{\infty}(v), \{v\}) \wedge \models \Delta(f, V'_{\infty}, \emptyset)\},$$

$$T' \triangleq \{f \in [\text{St} \rightarrow \text{St}_{\perp}] \mid \forall v \in \text{Var}. \models \Phi(f, p'_{\infty}(v), \{v\})\}.$$

In Theorem F.1, we proved

$$t'_n \in T \text{ for all } n \in \mathbb{N} \cup \{\infty\}. \quad (25)$$

In this theorem, our goal is to show $t'_\infty \in T'$. To do so, we prove the next three statements:

- (a) $t'_0 \in T'$.
- (b) If $t' \in T' \cap T$, then $F(t') \in T'$.
- (c) If $t'_n \in T'$ for all $n \in \mathbb{N}$, then $t'_\infty \in T'$.

It suffices to prove the three because (a), (b), and Eq. (25) imply $t'_n \in T'$ for all $n \in \mathbb{N}$, and this and (c) imply $t'_\infty \in T'$. We now prove (a), (b), and (c) as follows.

First, (a) follows directly from Lemma F.9.

Next, we prove (b). Consider $t' \in T' \cap T$. Our goal is to show $\models \Phi(F(t'), p'_\infty(v), \{v\})$ for all $v \in \text{Var}$. Observe that

$$F(t') = \text{cond}(\llbracket b \rrbracket, \text{seq}(\llbracket c_0 \rrbracket, t'), \llbracket \text{skip} \rrbracket).$$

By Theorem F.1 and induction hypothesis on c_0 , we have

$$\llbracket c_0 \rrbracket \in \gamma(p_0, d_0, V_0),$$

and by assumption, we have

$$t' \in \gamma(p'_\infty, d'_\infty, V'_\infty) = T' \cap T.$$

By applying to these the proofs of skip, sequential composition, and conditional cases, we have

$$\models \Phi(F(t'), p''(v), \{v\}) \text{ for all } v \in \text{Var}$$

where

$$p''(v) = fv(b)^c \cap \left(V_0 \cup (p_0)_{\cap} (d'_\infty(v))^c \cup (d_0)_{\cup} (p'_\infty(v)^c) \right)^c \cap \text{Var}.$$

Since p'' is the p part of $F^\#(p'_\infty, d'_\infty, V'_\infty)$ and $(p'_\infty, d'_\infty, V'_\infty)$ is a fixed point of $F^\#$, we have $p'' = p'_\infty$. Hence, we obtain $\models \Phi(F(t'), p'_\infty(v), \{v\})$ for all $v \in \text{Var}$. This completes the proof of (b).

Finally, we prove (c). Suppose that $t'_n \in T'$ for all $n \in \mathbb{N}$, and let $v \in \text{Var}$. Our goal is to show

$$\models \Phi(t'_\infty, p'_\infty(v), \{v\}).$$

Observe that Lemma F.10 implies the goal when applied to $K = p'_\infty(v)$, $L = \{v\}$, and $\{f_n\}_{n \in \mathbb{N}} = \{t'_n\}_{n \in \mathbb{N}}$. Hence, it suffices to show the three preconditions of the lemma:

- $\{t'_n\}_{n \in \mathbb{N}}$ is an ω -chain.
- For all $\tau \in \text{St}[p'_\infty(v)^c]$ and $n \in \mathbb{N}$, the set $\{\sigma' \in \text{St}[p'_\infty(v)] \mid t'_n(\sigma' \oplus \tau) \in \text{St}\}$ is \emptyset or $\text{St}[p'_\infty(v)]$.
- For all $n \in \mathbb{N}$, we have $\models \Phi(t'_n, p'_\infty(v), \{v\})$.

The first precondition was already observed in §3. The third one holds by the assumption that $t'_n \in T'$ for all $n \in \mathbb{N}$. For the second one, it is enough to show the next two statements:

- (i) For all $U \subseteq \text{Var}$ with $U \supseteq V'_\infty$, and for all $\tau \in \text{St}[U]$ and $n \in \mathbb{N}$, the next set is \emptyset or $\text{St}[U^c]$:

$$\{\sigma' \in \text{St}[U^c] \mid t'_n(\sigma' \oplus \tau) \in \text{St}\}.$$

- (ii) For all $v \in \text{Var}$,

$$p'_\infty(v)^c \supseteq V'_\infty.$$

We give the proof of the two statements below. This completes the proof of the while-loop case.

Proof of (ii). We prove a stronger statement: for all $n \in \mathbb{N}$ and $v \in \text{Var}$, $p'_n(v)^c \supseteq V'_n$. This statement implies (ii) because $p'_\infty(v)^c = (\bigcap_{n \in \mathbb{N}} p'_n(v))^c = \bigcup_{n \in \mathbb{N}} p'_n(v)^c \supseteq \bigcup_{n \in \mathbb{N}} V'_n = V'_\infty$. We prove the statement by induction on n . For $n = 0$, we have

$$p'_0(v)^c = \text{Var}^c = \emptyset \supseteq \emptyset = V'_0 \text{ for all } v \in \text{Var}.$$

For $n > 0$, let $v \in \text{Var}$. By induction hypothesis, $p'_{n-1}(v)^c \supseteq V'_{n-1}$ holds. Using this, we have

$$\begin{aligned} p'_n(v)^c &= fv(b) \cup V_0 \cup (p_0)_{\cap}(d'_{n-1}(v))^c \cup (d_0)_{\cup}(p'_{n-1}(v)^c) \\ &\supseteq fv(b) \cup V_0 \cup (d_0)_{\cup}(V'_{n-1}) = V'_n. \end{aligned}$$

This completes the proof of (ii).

Proof of (i). Consider $U \subseteq \text{Var}$, $\tau \in \text{St}[U]$, and $n \in \mathbb{N}$ such that $U \supseteq V'_\infty$. Let $\Sigma \triangleq \{\sigma' \in \text{St}[U^c] \mid t'_n(\sigma' \oplus \tau) \in \text{St}\}$. If $\Sigma = \emptyset$, there is nothing left to prove. So assume $\Sigma \neq \emptyset$. To prove $\Sigma = \text{St}[U^c]$, we need to show that $t'_n(\sigma' \oplus \tau) \in \text{St}$ for any $\sigma' \in \text{St}[U^c]$. Choose $\sigma' \in \text{St}[U^c]$. We show $t'_n(\sigma' \oplus \tau) \in \text{St}$ using the next two claims:

(iii) For all $n \in \mathbb{N}$ and $\sigma \in \text{St}$,

$$t'_n(\sigma) = \begin{cases} \llbracket c_0^{(i)} \rrbracket \sigma & \text{if } i \in I_n(\sigma) \\ \perp & \text{otherwise,} \end{cases} \quad (26)$$

where $c_0^{(i)} \triangleq (\text{skip}; c_0; \dots; c_0)$ that has i copies of c_0 , and

$$\begin{aligned} I_n(\sigma) \triangleq \{i \in [0, n-1] \mid \llbracket c_0^{(i)} \rrbracket \sigma \in \text{St} \wedge \llbracket b \rrbracket(\llbracket c_0^{(i)} \rrbracket \sigma) = \text{false} \\ \wedge \llbracket b \rrbracket(\llbracket c_0^{(i-1)} \rrbracket \sigma) = \dots = \llbracket b \rrbracket(\llbracket c_0^{(0)} \rrbracket \sigma) = \text{true}\}. \end{aligned}$$

Note that Eq. (26) is well-defined since $I_n(\sigma)$ has at most one element.

(iv) For all $n \in \mathbb{N}$,

$$\models \Delta(\llbracket c_0^{(n)} \rrbracket, V'_\infty, fv(b)).$$

We give the proof of the two claims below, and for now we just assume them.

Since $\Sigma \neq \emptyset$, there is some $\sigma'' \in \text{St}[U^c]$ such that $t'_n(\sigma'' \oplus \tau) \in \text{St}$. Since $t'_n(\sigma'' \oplus \tau) \in \text{St}$, (iii) implies that

$$t'_n(\sigma'' \oplus \tau) = \llbracket c_0^{(m)} \rrbracket(\sigma'' \oplus \tau) \in \text{St}$$

for some $m \in I_n(\sigma'' \oplus \tau)$. Since $\sigma' \oplus \tau \sim_{V'_\infty} \sigma'' \oplus \tau$ (by $U \supseteq V'_\infty$) and $\llbracket c_0^{(i)} \rrbracket(\sigma'' \oplus \tau) \in \text{St}$ for all $i \in [0, m]$ (by $\llbracket c_0^{(m)} \rrbracket(\sigma'' \oplus \tau) \in \text{St}$), (iv) implies that

$$\llbracket c_0^{(i)} \rrbracket(\sigma' \oplus \tau) \in \text{St} \text{ and } \llbracket c_0^{(i)} \rrbracket(\sigma' \oplus \tau) \sim_{fv(b)} \llbracket c_0^{(i)} \rrbracket(\sigma'' \oplus \tau) \text{ for all } i \in [0, m].$$

By combining these with $m \in I_n(\sigma'' \oplus \tau)$, we get $m \in I_n(\sigma' \oplus \tau)$. Hence, by (iii), we have

$$t'_n(\sigma' \oplus \tau) = \llbracket c_0^{(m)} \rrbracket(\sigma' \oplus \tau) \in \text{St}.$$

This completes the proof of (i).

Proof of (iii). We prove this by induction on n . For $n = 0$, $t'_0(\sigma) = \perp$ and $I_0(\sigma) = \emptyset$ for all $\sigma \in \text{St}$. Hence, Eq. (26) holds. For $n > 0$, we have

$$\begin{aligned} t'_n(\sigma) &= F(t'_{n-1})(\sigma) \\ &= \begin{cases} \sigma & \text{if } \llbracket b \rrbracket \sigma = \text{false} \\ (t'_{n-1})^\dagger(\llbracket c_0 \rrbracket \sigma) & \text{if } \llbracket b \rrbracket \sigma = \text{true} \end{cases} \end{aligned}$$

$$\begin{aligned}
&= \begin{cases} \sigma & \text{if } \llbracket b \rrbracket \sigma = \text{false} \cdots (*_1) \\ \llbracket c_0^{(i)} \rrbracket (\llbracket c_0 \rrbracket \sigma) & \text{if } \llbracket b \rrbracket \sigma = \text{true}, \llbracket c_0 \rrbracket \sigma \in \text{St} \text{ and } i \in I_{n-1}(\llbracket c_0 \rrbracket \sigma) \cdots (*_2) \\ \perp & \text{otherwise} \end{cases} \\
&= \begin{cases} \llbracket c_0^{(0)} \rrbracket \sigma & \text{if } 0 \in I_n(\sigma) \cdots (*'_1) \\ \llbracket c_0^{(i+1)} \rrbracket \sigma & \text{if } i+1 \in I_n(\sigma) \text{ and } i+1 \geq 1 \cdots (*'_2) \\ \perp & \text{otherwise.} \end{cases}
\end{aligned}$$

The second equality is by the definition of F , the third by induction hypothesis, and the last by the following: for any $\sigma \in \text{St}$ and $j \in \{1, 2\}$, $(*_j)$ holds iff $(*'_j)$ holds; and for any $i \in \mathbb{N}$, $\llbracket c_0 \rrbracket \sigma \in \text{St}$ implies $\llbracket c_0^{(i)} \rrbracket (\llbracket c_0 \rrbracket \sigma) = \llbracket c_0^{(i+1)} \rrbracket \sigma$, and $\llbracket c_0^{(i+1)} \rrbracket \sigma \in \text{St}$ implies $\llbracket c_0 \rrbracket \sigma \in \text{St}$. Hence, Eq. (26) holds. This completes the proof of (iii).

Proof of (iv). To prove this, we prove a stronger statement: for all $n \in \mathbb{N}$, $\models \Delta(\llbracket c_0^{(n)} \rrbracket, V'_{n+1}, fv(b))$. This statement implies (iv) since $V'_n \subseteq V'_\infty$ for all $n \in \mathbb{N}$ and we have the weakening lemma for Δ (Lemma F.3). We prove the statement by induction on n . For $n = 0$, $\llbracket c_0^{(0)} \rrbracket = \llbracket \text{skip} \rrbracket$ and $V'_{n+1} = fv(b) \cup V_0$. By Theorem F.1 on skip , we have $\models \Delta(\llbracket \text{skip} \rrbracket, \{v\}, \{v\})$ for all $v \in \text{Var}$, and then by the merging lemma for Δ (Lemma F.5), we have

$$\models \Delta(\llbracket \text{skip} \rrbracket, fv(b), fv(b)).$$

Since $V'_{n+1} \supseteq fv(b)$, $\models \Delta(\llbracket c_0^{(n)} \rrbracket, V'_{n+1}, fv(b))$ holds by the weakening lemma for Δ (Lemma F.3). Next, for $n > 0$, $\llbracket c_0^{(n)} \rrbracket = \llbracket c_0; c_0^{(n-1)} \rrbracket$ and $V'_{n+1} = fv(b) \cup V_0 \cup (d_0)_{\cup}(V'_n)$. By $\llbracket c_0 \rrbracket^\# = (p_0, d_0, V_0)$, Theorem F.1 on c_0 , and induction hypothesis of the theorem (not that of the claim (iv)), we have

$$\llbracket c_0 \rrbracket \in \gamma(p_0, d_0, V_0).$$

Also, by induction hypothesis of our strengthening of the claim (iv) and the weakening lemma for Δ (Lemma F.3),

$$\models \Delta(\llbracket c_0^{(n-1)} \rrbracket, V'_n, \{v\}) \text{ for all } v \in fv(b).$$

By applying to these the proof of Theorem F.1 (on the sequential composition case), we have

$$\models \Delta(\llbracket c_0; c_0^{(n-1)} \rrbracket, V_0 \cup (d_0)_{\cup}(V'_n), \{v\}) \text{ for all } v \in fv(b).$$

By the merging lemma for Δ (Lemma F.5), $\models \Delta(\llbracket c_0^{(n)} \rrbracket, V_0 \cup (d_0)_{\cup}(V'_n), fv(b))$ holds. Since V'_{n+1} includes $V_0 \cup (d_0)_{\cup}(V'_n)$, we get $\models \Delta(\llbracket c_0^{(n)} \rrbracket, V'_{n+1}, fv(b))$ by the weakening lemma for Δ (Lemma F.3). This completes the proof of (iv).

Case $c \equiv (x := \text{sam}(\text{name}(\alpha, e), \text{dist}_N(e_1, e_2), \lambda y. e'))$. To prove the conclusion, consider $v \in \text{Var}$ and $\tau \in \text{St}[p(v)^c]$. We should show $g \in \phi_{p(v), \{v\}}$, where

$$g(\sigma) = \pi_{\text{Var}, \{v\}}(\llbracket c \rrbracket(\sigma \oplus \tau)) = [v \mapsto \llbracket c \rrbracket(\sigma \oplus \tau)(v)].$$

We prove this by case analysis on v .

First, suppose $v \notin \{x\} \cup \{\text{val}_\mu, \text{pr}_\mu, \text{cnt}_\mu \mid \mu \in \text{Name}, \mu = (\alpha, _)\}$. Then, $p(v) = \text{Var}$ and

$$g(\sigma) = [v \mapsto \llbracket c \rrbracket \sigma(v)] = [v \mapsto \sigma(v)] = \pi_{\text{Var}, \{v\}}(\sigma)$$

for all $\sigma \in \text{St}[p(v)]$. Here the first equality follows from $\tau \in \text{St}[\emptyset]$, and the second equality holds since $\llbracket c \rrbracket$ does not change the value of v . Hence, by Assumption 3, $g = \pi_{\text{Var}, \{v\}} \in \phi_{\text{Var}, \{v\}} = \phi_{p(v), \{v\}}$.

Next, suppose $v \in \{x\} \cup \{\text{val}_\mu, \text{pr}_\mu, \text{cnt}_\mu \mid \mu \in \text{Name}, \mu = (\alpha, _)\}$. Then, we have $p(v)^c \supseteq fv(e)$: if e is a constant, $fv(e) = \emptyset$ holds, and if e is not a constant, the definition of $p(v)$ ensures this. Thus, there exists $\mu_0 \in \text{Name}$ such that $\text{create_name}(\alpha, \llbracket e \rrbracket(\sigma \oplus \tau)) = \mu_0$ for all $\sigma \in \text{St}[p(v)]$. We now do refined case analysis on v using μ_0 .

- Case $v \in \{val_\mu, pr_\mu, cnt_\mu \mid \mu \in \text{Name}, \mu = (\alpha, _), \mu \neq \mu_0\}$. In this case,

$$g(\sigma) = [v \mapsto (\sigma \oplus \tau)(v)] = \pi_{\text{Var}, \{v\}}(\sigma \oplus \tau)$$

for all $\sigma \in \text{St}[p(v)]$. Here the first equality holds since $\llbracket c \rrbracket$ does not change the value of v . By Assumption 3, we have $\pi_{\text{Var}, \{v\}} \in \phi_{\text{Var}, \{v\}}$. Then, by Assumption 5, we obtain $g \in \phi_{p(v), \{v\}}$. Note that this argument does not depend on the value of $p(v)$ (which can be $\text{Var}, fv(e)^c \cap (v+1)^\sharp$, etc., depending on e and v).

- Case $v \in \{x, val_{\mu_0}\}$. Define $K_1 \triangleq \llbracket e'[\mu_0/y] \rrbracket^\sharp$. Then, $p(v)^c \supseteq (\llbracket e'[\mu_0/y] \rrbracket^\sharp)^c = K_1^c$. So, there exist $\tau_1 \in \text{St}[K_1^c]$ and $\tau_2 \in \text{St}[p(v)^c \setminus K_1^c]$ such that $\tau = \tau_1 \oplus \tau_2$. Let $h : \text{St}[K_1] \rightarrow \text{St}[\{v\}]$ be a function defined by

$$h(\sigma') \triangleq [v \mapsto \llbracket e'[\mu_0/y] \rrbracket(\sigma' \oplus \tau_1)].$$

Then,

$$g(\sigma) = [v \mapsto \llbracket e'[\mu_0/y] \rrbracket(\sigma \oplus \tau)] = h(\sigma \oplus \tau_2)$$

for all $\sigma \in \text{St}[p(v)]$. By Assumption 2, we have $h \in \phi_{K_1, \{v\}}$. Then, by Assumption 5, we obtain $g \in \phi_{p(v), \{v\}}$.

- Case $v \equiv pr_{\mu_0}$. In this case, $p(v) = fv(e)^c \cap K$ for $K = \llbracket \text{pdf}_N(\mu_0; e_1, e_2) \rrbracket^\sharp$. Since $\tau \in \text{St}[p(v)^c]$ and $p(v)^c = (K \setminus fv(e)^c) \uplus K^c$, there exist $\tau_1 \in \text{St}[K \setminus fv(e)^c]$ and $\tau_2 \in \text{St}[K^c]$ such that $\tau = \tau_1 \oplus \tau_2$. Using τ_1 and τ_2 , we have

$$\begin{aligned} g(\sigma) &= [v \mapsto \llbracket \text{dist}_N(e_1, e_2) \rrbracket(\sigma \oplus \tau)((\sigma \oplus \tau)(\mu_0))] \\ &= [v \mapsto \llbracket \text{pdf}_N(\mu_0; e_1, e_2) \rrbracket(\sigma \oplus \tau)] = h(\sigma \oplus \tau_1) \end{aligned}$$

for all $\sigma \in \text{St}[p(v)]$, where $h : \text{St}[K] \rightarrow \text{St}[\{v\}]$ is defined by

$$h(\sigma') = [v \mapsto \llbracket \text{pdf}_N(\mu_0; e_1, e_2) \rrbracket(\sigma' \oplus \tau_2)].$$

Here the first equality follows from the definition of $\llbracket c \rrbracket$, the second equality holds because pdf_N is the density function of a normal distribution, and the third equality comes from $\tau = \tau_1 \oplus \tau_2$. By Assumption 2, we have $h \in \phi_{K, \{v\}}$. Then, by Assumption 5, we obtain $g \in \phi_{p(v), \{v\}}$.

- Case $v \equiv cnt_{\mu_0}$. The proof is similar to the above case $v \equiv pr_{\mu_0}$. In this case, $p(v) = fv(e)^c \cap K$ for $K = \llbracket cnt_{\mu_0} + 1 \rrbracket^\sharp$. As in the above case, there exist $\tau_1 \in \text{St}[K \setminus fv(e)^c]$ and $\tau_2 \in \text{St}[K^c]$ such that $\tau = \tau_1 \oplus \tau_2$, and we have

$$\begin{aligned} g(\sigma) &= [v \mapsto (\sigma \oplus \tau)(cnt_{\mu_0}) + 1] \\ &= [v \mapsto \llbracket cnt_{\mu_0} + 1 \rrbracket(\sigma \oplus \tau)] = h(\sigma \oplus \tau_1) \end{aligned}$$

for all $\sigma \in \text{St}[p(v)]$, where $h : \text{St}[K] \rightarrow \text{St}[\{v\}]$ is defined by

$$h(\sigma') = [v \mapsto \llbracket cnt_{\mu_0} + 1 \rrbracket(\sigma' \oplus \tau_2)].$$

By Assumption 2, we have $h \in \phi_{K, \{v\}}$. Then, by Assumption 5, we obtain $g \in \phi_{p(v), \{v\}}$.

Case $c \equiv \text{obs}(\text{dist}_N(e_1, e_2), r)$. To prove the conclusion, consider $v \in \text{Var}$ and $\tau \in \text{St}[p(v)^c]$. We should show $g \in \phi_{p(v), \{v\}}$, where

$$g(\sigma) = \pi_{\text{Var}, \{v\}}(\llbracket c \rrbracket(\sigma \oplus \tau)) = [v \mapsto \llbracket c \rrbracket(\sigma \oplus \tau)(v)].$$

We prove this by case analysis on v .

First, suppose $v \neq \text{like}$. Then, $p(v) = \text{Var}$ and

$$g(\sigma) = [v \mapsto \llbracket c \rrbracket \sigma(v)] = [v \mapsto \sigma(v)] = \pi_{\text{Var}, \{v\}}(\sigma)$$

for all $\sigma \in \text{St}[p(v)]$. Here the first equality is by $\tau \in \text{St}[\emptyset]$, and the second equality holds since $\llbracket c \rrbracket$ does not change the value of v . Hence, by Assumption 3, $g = \pi_{\text{Var}, \{v\}} \in \phi_{\text{Var}, \{v\}} = \phi_{p(v), \{v\}}$.

Next, suppose $v \equiv \text{like}$. Then, $p(v) = (\text{like} \times \text{pdf}_{\mathbb{N}}(r; e_1, e_2))^{\#}$ and

$$\begin{aligned} g(\sigma) &= [v \mapsto (\sigma \oplus \tau)(\text{like}) \cdot \llbracket \text{dist}_{\mathbb{N}}(e_1, e_2) \rrbracket (\sigma \oplus \tau)(r)] \\ &= [v \mapsto \llbracket \text{like} \times \text{pdf}_{\mathbb{N}}(r; e_1; e_2) \rrbracket (\sigma \oplus \tau)] \end{aligned}$$

for all $\sigma \in \text{St}[p(v)]$. Here the first equality is by the definition of $\llbracket c \rrbracket$, and the second equality holds because $\text{pdf}_{\mathbb{N}}$ is the density function of a normal distribution. Hence, by Assumption 2, we have $g \in \phi_{p(v), \{v\}}$. \square

F.5 Proofs of Lemmas for Theorem F.2

Here are the lemmas used to prove Theorem F.2:

LEMMA F.3 (WEAKENING; Δ). *Let $f \in [\text{St} \rightarrow \text{St}_{\perp}]$ and $K, K', L, L' \subseteq \text{Var}$. Then,*

$$\models \Delta(f, K, L) \wedge (K \subseteq K') \wedge (L \supseteq L') \implies \models \Delta(f, K', L').$$

PROOF. Consider $\sigma, \sigma' \in \text{St}$ with $\sigma \sim_{K'} \sigma'$. Then, $\sigma \sim_K \sigma'$ because $K \subseteq K'$. Since $\models \Delta(f, K, L)$,

$$(f(\sigma) \in \text{St} \iff f(\sigma') \in \text{St}) \text{ and } (f(\sigma) \in \text{St} \implies f(\sigma) \sim_L f(\sigma')).$$

Note that the conclusion of the second conjunct implies $f(\sigma) \sim_{L'} f(\sigma')$ since $L' \subseteq L$. From what we have just shown, the desired conclusion $\models \Delta(f, K', L')$ follows. \square

LEMMA F.4 (WEAKENING; Φ). *Let $f \in [\text{St} \rightarrow \text{St}_{\perp}]$ and $K, K', L, L' \subseteq \text{Var}$. Then,*

$$\models \Phi(f, K, L) \wedge (K \supseteq K') \wedge (L \supseteq L') \implies \models \Phi(f, K', L').$$

PROOF. We prove the lemma using Assumptions 3, 5 and 6. Consider $\tau \in \text{St}[(K')^c]$, and let g be the following partial function:

$$g : \text{St}[K'] \rightarrow \text{St}[L'], \quad g(\sigma') \triangleq \begin{cases} (\pi_{\text{Var}, L'} \circ f)(\sigma' \oplus \tau) & \text{if } f(\sigma' \oplus \tau) \in \text{St} \\ \text{undefined} & \text{otherwise.} \end{cases}$$

We should show $g \in \phi_{K', L'}$. Note that $(K')^c \supseteq K^c$. Thus, there exist $\tau_1 \in \text{St}[(K')^c \setminus K^c]$ and $\tau_2 \in \text{St}[K^c]$ such that $\tau = \tau_1 \oplus \tau_2$. Define a partial function $h : \text{St}[K] \rightarrow \text{St}[L]$ by

$$h(\sigma'') \triangleq \begin{cases} (\pi_{\text{Var}, L} \circ f)(\sigma'' \oplus \tau_2) & \text{if } f(\sigma'' \oplus \tau_2) \in \text{St} \\ \text{undefined} & \text{otherwise.} \end{cases}$$

Then, since $\models \Phi(f, K, L)$, we have $h \in \phi_{K, L}$. Note that for all $\sigma' \in \text{St}[K']$,

$$g(\sigma') = (\pi_{L, L'} \circ h)(\sigma' \oplus \tau_1).$$

By Assumptions 3, 5 and 6, the above equation implies $g \in \phi_{K', L'}$, as desired. \square

LEMMA F.5 (MERGING; Δ). *Let $f \in [\text{St} \rightarrow \text{St}_{\perp}]$ and $K, K', L, L' \subseteq \text{Var}$. Then,*

$$\models \Delta(f, K, L) \wedge \models \Delta(f, K', L') \implies \models \Delta(f, K \cup K', L \cup L').$$

PROOF. Consider $\sigma, \sigma' \in \text{St}$ with $\sigma \sim_{K \cup K'} \sigma'$. Then, $\sigma \sim_K \sigma'$, and by the assumption that $\models \Delta(f, K, L)$, we have

$$f(\sigma) \in \text{St} \iff f(\sigma') \in \text{St}.$$

It remains to show that if $f(\sigma), f(\sigma') \in \text{St}$, then $f(\sigma) \sim_{L \cup L'} f(\sigma')$. Assume $f(\sigma), f(\sigma') \in \text{St}$. Since $\sigma \sim_{K \cup K'} \sigma'$ and we have $\models \Delta(f, K, L)$ and $\models \Delta(f, K', L')$ by assumption,

$$f(\sigma) \sim_L f(\sigma') \text{ and } f(\sigma) \sim_{L'} f(\sigma').$$

This implies that $f(\sigma) \sim_{L \cup L'} f(\sigma')$, as desired. \square

LEMMA F.6 (MERGING; Φ). *Let $f \in [\text{St} \rightarrow \text{St}_\perp]$ and $K, K', L, L' \subseteq \text{Var}$. Then,*

$$\models \Phi(f, K, L) \wedge \models \Phi(f, K', L') \implies \models \Phi(f, K \cap K', L \cup L').$$

PROOF. Uses the weakening lemma for Φ (Lemma F.4), we have

$$\models \Phi(f, K \cap K', L) \text{ and } \models \Phi(f, K \cap K', L').$$

This and Assumption 4 then imply the desired conclusion. Concretely, for all $\tau \in \text{St}[(K \cap K')^c]$, if g, g_1 , and g_2 are the following partial functions

$$\begin{aligned} g : \text{St}[K \cap K'] \rightarrow \text{St}[L \cup L'], & \quad g(\sigma') \triangleq \begin{cases} (\pi_{\text{Var}, L \cup L'} \circ f)(\sigma' \oplus \tau) & \text{if } f(\sigma' \oplus \tau) \in \text{St} \\ \text{undefined} & \text{otherwise,} \end{cases} \\ g_1 : \text{St}[K \cap K'] \rightarrow \text{St}[L], & \quad g_1(\sigma') \triangleq \begin{cases} (\pi_{\text{Var}, L} \circ f)(\sigma' \oplus \tau) & \text{if } f(\sigma' \oplus \tau) \in \text{St} \\ \text{undefined} & \text{otherwise,} \end{cases} \\ g_2 : \text{St}[K \cap K'] \rightarrow \text{St}[L'], & \quad g_2(\sigma') \triangleq \begin{cases} (\pi_{\text{Var}, L'} \circ f)(\sigma' \oplus \tau) & \text{if } f(\sigma' \oplus \tau) \in \text{St} \\ \text{undefined} & \text{otherwise,} \end{cases} \end{aligned}$$

then $g_1 \in \phi_{K \cap K', L}$, $g_2 \in \phi_{K \cap K', L'}$, and $g = \langle g_1, g_2 \rangle$, so that by Assumption 4, we have $g \in \phi_{K \cap K', L \cup L'}$ as desired. \square

LEMMA F.7 (SEQUENCE). *Let $f, g \in [\text{St} \rightarrow \text{St}_\perp]$ and $K, L, L', M \subseteq \text{Var}$. Then,*

$$\models \Phi(f, K, L) \wedge \models \Phi(g, L, M) \wedge \models \Delta(f, K^c, L' \setminus L) \wedge \models \Delta(g, L', M) \implies \models \Phi(\text{seq}(f, g), K, M).$$

PROOF. Consider $f, g \in [\text{St} \rightarrow \text{St}_\perp]$ and $K, L, L', M \subseteq \text{Var}$ that satisfy the given conditions:

$$\models \Phi(f, K, L), \quad \models \Phi(g, L, M), \quad \models \Delta(f, K^c, L' \setminus L), \quad \text{and} \quad \models \Delta(g, L', M).$$

To prove the conclusion, pick an arbitrary $\tau \in \text{St}[K^c]$. We have to show $h \in \phi_{K, M}$, where

$$h(\sigma) = \begin{cases} \pi_{\text{Var}, M}((g^\dagger \circ f)(\sigma \oplus \tau)) & \text{if } (g^\dagger \circ f)(\sigma \oplus \tau) \in \text{St} \\ \text{undefined} & \text{otherwise.} \end{cases}$$

Observe that since $\models \Phi(f, K, L)$ and $\models \Phi(g, L, M)$, we have $h_1 \in \phi_{K, L}$ and $h_2 \in \phi_{L, M}$ for any $\tau_1 \in \text{St}[K^c]$ and $\tau_2 \in \text{St}[L^c]$, where h_1 and h_2 are parameterised by τ_1 and τ_2 , and defined by

$$\begin{aligned} h_1 : \text{St}[K] \rightarrow \text{St}[L], & \quad h_1(\sigma) \triangleq \begin{cases} \pi_{\text{Var}, L}(f(\sigma \oplus \tau_1)) & \text{if } f(\sigma \oplus \tau_1) \in \text{St} \\ \text{undefined} & \text{otherwise,} \end{cases} \\ h_2 : \text{St}[L] \rightarrow \text{St}[M], & \quad h_2(\sigma) \triangleq \begin{cases} \pi_{\text{Var}, M}(g(\sigma \oplus \tau_2)) & \text{if } g(\sigma \oplus \tau_2) \in \text{St} \\ \text{undefined} & \text{otherwise.} \end{cases} \end{aligned}$$

Given these, it suffices to show the claim that $h = h_2 \circ h_1$ for some τ_1 and τ_2 : if the claim holds, then we have $h = h_2 \circ h_1 \in \phi_{K, M}$ by Assumption 6, since $h_1 \in \phi_{K, L}$ and $h_2 \in \phi_{L, M}$. We prove the claim by case analysis on $f(- \oplus \tau)$.

Case $f(\sigma \oplus \tau) \notin \text{St}$ for all $\sigma \in \text{St}[K]$. In this case, we set $\tau_1 \triangleq \tau$ and pick any $\tau_2 \in \text{St}[L^c]$. Then, for all $\sigma \in \text{St}[K]$, $h(\sigma)$ and $(h_2 \circ h_1)(\sigma)$ are both undefined, as desired. Note that the latter term is undefined since $h_1(\sigma)$ is undefined.

Case $f(\sigma' \oplus \tau) \in \text{St}$ for some $\sigma' \in \text{St}[K]$. In this case, we set $\tau_1 \triangleq \tau$ and $\tau_2 \triangleq \pi_{\text{Var},L^c}(f(\sigma' \oplus \tau))$. To show $h = h_2 \circ h_1$, consider any $\sigma \in \text{St}[K]$. If $f(\sigma \oplus \tau) \notin \text{St}$, then by the same argument for the above case, $h(\sigma)$ and $(h_2 \circ h_1)(\sigma)$ are both undefined. So, assume that $f(\sigma \oplus \tau) \in \text{St}$. Then,

$$h(\sigma) = \begin{cases} \pi_{\text{Var},M}(g(\sigma_1)) & \text{if } g(\sigma_1) \in \text{St} \\ \text{undefined} & \text{otherwise,} \end{cases} \quad (h_2 \circ h_1)(\sigma) = \begin{cases} \pi_{\text{Var},M}(g(\sigma_2)) & \text{if } g(\sigma_2) \in \text{St} \\ \text{undefined} & \text{otherwise,} \end{cases} \quad (27)$$

where

$$\sigma_1 = f(\sigma \oplus \tau) \in \text{St}, \quad \sigma_2 = \pi_{\text{Var},L}(f(\sigma \oplus \tau)) \oplus \pi_{\text{Var},L^c}(f(\sigma' \oplus \tau)) \in \text{St}.$$

Our goal is to show that $h(\sigma)$ and $(h_2 \circ h_1)(\sigma)$ are both undefined, or they are both defined and are the same. By $\models \Delta(g, L', M)$ and Eq. (27), it suffices to show $\sigma_1 \sim_{L'} \sigma_2$. To prove this, we show a stronger statement: $\sigma_1 \sim_L \sigma_2$ and $\sigma_1 \sim_{L' \setminus L} \sigma_2$. The former relation holds since $\pi_{\text{Var},L}(\sigma_1) = \pi_{\text{Var},L}(f(\sigma \oplus \tau)) = \pi_{\text{Var},L}(\sigma_2)$. The latter relation is equivalent to $f(\sigma \oplus \tau) \sim_{L' \setminus L} f(\sigma' \oplus \tau)$, and this holds by $\models \Delta(f, K^c, L' \setminus L)$ and $\sigma \oplus \tau \sim_{K^c} \sigma' \oplus \tau$. Hence, $h = h_2 \circ h_1$ as desired. \square

LEMMA F.8 (CONDITIONAL). *Let $f, f' \in [\text{St} \rightarrow \text{St}_\perp]$ and $K, L \subseteq \text{Var}$. Then, for any boolean expression b ,*

$$\models \Phi(f, K, L) \wedge \models \Phi(f', K, L) \wedge (K^c \supseteq \text{fv}(b)) \implies \models \Phi(\text{cond}(\llbracket b \rrbracket, f, f'), K, L).$$

PROOF. Let f, f', K, L , and b be the functions, sets and a boolean expression such that

$$\models \Phi(f, K, L), \quad \models \Phi(f', K, L), \quad \text{and} \quad K^c \supseteq \text{fv}(b).$$

Consider $\tau \in \text{St}[K^c]$. Define $f'' \triangleq \text{cond}(\llbracket b \rrbracket, f, f')$, and also partial functions g, g' , and g'' as follows:

$$\begin{aligned} g : \text{St}[K] &\rightarrow \text{St}[L], & g(\sigma) &\triangleq \begin{cases} (\pi_{\text{Var},L} \circ f)(\sigma \oplus \tau) & \text{if } f(\sigma \oplus \tau) \in \text{St} \\ \text{undefined} & \text{otherwise,} \end{cases} \\ g' : \text{St}[K] &\rightarrow \text{St}[L], & g'(\sigma) &\triangleq \begin{cases} (\pi_{\text{Var},L} \circ f')(\sigma \oplus \tau) & \text{if } f'(\sigma \oplus \tau) \in \text{St} \\ \text{undefined} & \text{otherwise,} \end{cases} \\ g'' : \text{St}[K] &\rightarrow \text{St}[L], & g''(\sigma) &\triangleq \begin{cases} (\pi_{\text{Var},L} \circ f)(\sigma \oplus \tau) & \text{if } f(\sigma \oplus \tau) \in \text{St} \\ \text{undefined} & \text{otherwise,} \end{cases} \end{aligned}$$

We should show $g'' \in \phi_{K,L}$. Since $K^c \supseteq \text{fv}(b)$, either $\llbracket b \rrbracket(\sigma \oplus \tau) = \text{true}$ for all $\sigma \in \text{St}[K]$ or $\llbracket b \rrbracket(\sigma \oplus \tau) = \text{false}$ for all $\sigma \in \text{St}[K]$. In the former case, $g'' = g$, and in the latter case, $g'' = g'$. Since both g and g' are in $\phi_{K,L}$, we have the desired $g'' \in \phi_{K,L}$ in both cases. \square

LEMMA F.9 (LOOP; BASE). *Let $K, L \subseteq \text{Var}$. Then,*

$$\models \Phi((\lambda \sigma \in \text{St}. \perp), K, L).$$

PROOF. Consider $\tau \in \text{St}[K^c]$. Define a partial function $g : \text{St}[K] \rightarrow \text{St}[L]$ by

$$\begin{aligned} g(\sigma) &\triangleq \begin{cases} (\pi_{\text{Var},L} \circ (\lambda \sigma \in \text{St}. \perp))(\sigma) & \text{if } (\lambda \sigma \in \text{St}. \perp)(\sigma) \in \text{St} \\ \text{undefined} & \text{otherwise} \end{cases} \\ &= \text{undefined.} \end{aligned}$$

Then, $g \in \phi_{K,L}$ by Assumption 7. \square

LEMMA F.10 (LOOP; LIMIT). *Let $K, L \subseteq \text{Var}$ and $\{f_n \in [\text{St} \rightarrow \text{St}_\perp]\}_{n \in \mathbb{N}}$ be an ω -chain (i.e., $f_n \sqsubseteq f_{n+1}$ for all $n \in \mathbb{N}$). Here we write $f \sqsubseteq g$ iff $f(\sigma) \sqsubseteq g(\sigma)$ for all $\sigma \in \text{St}$. Suppose that for any $\tau \in \text{St}[K^c]$ and $n \in \mathbb{N}$, the set $\{\sigma \in \text{St}[K] \mid f_n(\sigma \oplus \tau) \in \text{St}\}$ is either \emptyset or $\text{St}[K]$. Then,*

$$\bigwedge_{n \in \mathbb{N}} \models \Phi(f_n, K, L) \implies \models \Phi(\bigsqcup_{n \in \mathbb{N}} f_n, K, L).$$

PROOF. Consider an ω -chain $\{f_n \in [\text{St} \rightarrow \text{St}_\perp]\}_{n \in \mathbb{N}}$ such that $\models \Phi(f_n, K, L)$ for all n . Pick an arbitrary $\tau \in \text{St}[K^c]$. Let $f_\infty \triangleq \bigsqcup_{n \in \mathbb{N}} f_n$, and define partial functions g_∞ and g_n for all n as follows:

$$\begin{aligned} g_\infty : \text{St}[K] &\rightarrow \text{St}[L], & g_\infty(\sigma) &\triangleq \begin{cases} (\pi_{\text{Var}, L} \circ f_\infty)(\sigma \oplus \tau) & \text{if } f_\infty(\sigma \oplus \tau) \in \text{St} \\ \text{undefined} & \text{otherwise,} \end{cases} \\ g_n : \text{St}[K] &\rightarrow \text{St}[L], & g_n(\sigma) &\triangleq \begin{cases} (\pi_{\text{Var}, L} \circ f_n)(\sigma \oplus \tau) & \text{if } f_n(\sigma \oplus \tau) \in \text{St} \\ \text{undefined} & \text{otherwise.} \end{cases} \end{aligned}$$

Then, $\{g_n\}_{n \in \mathbb{N}}$ is an ω -chain when we order g_n 's by graph inclusion, and g_∞ is the least upper bound of this chain for the same order. We will show that $g_\infty = g_n$ for some $n \in \mathbb{N}$ by case analysis on $\Sigma_n \triangleq \{\sigma \in \text{St}[K] \mid f_n(\sigma \oplus \tau) \in \text{St}\}$. Note that this implies the desired $g_\infty \in \phi_{K,L}$ because $g_n \in \phi_{K,L}$ for all $n \in \mathbb{N}$. If $\Sigma_n = \emptyset$ for all $n \in \mathbb{N}$, then $g = g_n$ for any n , since both g and g_n are the same empty partial function. Otherwise, by the assumption of the lemma, $\Sigma_n = \text{St}[K]$ for some n . This means that g_n is the total function, and so $g_n = g_m$ for all $m \geq n$, which implies that $g = g_n$, as desired. \square

G DEFERRED RESULTS IN §5.3

G.1 Proof of Theorem 5.10

PROOF OF THEOREM 5.10. We go through the assumptions, and show that they are satisfied by $\phi^{(d)}$ and $\phi^{(l)}$.

Case of Assumption 3. Let $K, L \subseteq \text{Var}$ such that $L \subseteq K$. The projection $\pi_{K,L}$ is total and has an open set as its domain. Furthermore, the projection $\pi_{K,L}$ is differentiable and 1-Lipschitz continuous. Since Lipschitz continuity implies local Lipschitz continuity, we have both $\pi_{K,L} \in \phi_{K,L}^{(d)}$ and $\pi_{K,L} \in \phi_{K,L}^{(l)}$, as desired.

Case of Assumption 4. Let $K, L_0, L_1 \subseteq \text{Var}$ such that $L_0 \cap L_1 = \emptyset$. Consider $f_0, g_0 \in [\text{St}[K] \rightarrow \text{St}[L_0]]$ and $f_1, g_1 \in [\text{St}[K] \rightarrow \text{St}[L_1]]$ such that all of the following hold:

$$f_0 \in \phi_{K,L_0}^{(d)}, \quad f_1 \in \phi_{K,L_1}^{(d)}, \quad g_0 \in \phi_{K,L_0}^{(l)}, \quad \text{and} \quad g_1 \in \phi_{K,L_1}^{(l)}.$$

Let

$$L \triangleq L_0 \cup L_1;$$

$$\begin{aligned} f : \text{St}[K] &\rightarrow \text{St}[L], & f(\sigma) &\triangleq \begin{cases} f_0(\sigma) \oplus f_1(\sigma) & \text{if } \sigma \in \text{dom}(f_0) \cap \text{dom}(f_1), \\ \text{undefined} & \text{otherwise;} \end{cases} \\ g : \text{St}[K] &\rightarrow \text{St}[L], & g(\sigma) &\triangleq \begin{cases} g_0(\sigma) \oplus g_1(\sigma) & \text{if } \sigma \in \text{dom}(g_0) \cap \text{dom}(g_1), \\ \text{undefined} & \text{otherwise.} \end{cases} \end{aligned}$$

We should show that f and g satisfy $\phi_{K,L}^{(d)}$ and $\phi_{K,L}^{(l)}$, respectively. In both cases, $\text{dom}(f)$ and $\text{dom}(g)$ are the intersections of two open sets, so that they are open as required.

To prove the differentiability of f , consider $\sigma \in \text{dom}(f)$. Let h_0 and h_1 be the linear functions in $[\text{St}[K] \rightarrow \text{St}[L_0]]$ and $[\text{St}[K] \rightarrow \text{St}[L_1]]$, respectively, such that their domains are open and contain

0, and for all $i \in \{0, 1\}$,

$$\lim_{\sigma' \rightarrow 0} \frac{\|f_i(\sigma + \sigma') - f_i(\sigma) - h_i(\sigma')\|_2}{\|\sigma'\|_2} = 0.$$

Let h be the linear function in $[\text{St}[K] \rightarrow \text{St}[L]]$ defined by

$$h(\sigma) \triangleq \begin{cases} h_0(\sigma) \oplus h_1(\sigma) & \text{if } \sigma \in \text{dom}(h_0) \cap \text{dom}(h_1); \\ \text{undefined} & \text{otherwise.} \end{cases}$$

Then, $\text{dom}(h)$ is open and contains σ . Furthermore,

$$\begin{aligned} & \lim_{\sigma' \rightarrow 0} \frac{\|f(\sigma + \sigma') - f(\sigma) - h(\sigma')\|_2}{\|\sigma'\|_2} \\ &= \lim_{\sigma' \rightarrow 0} \frac{\sqrt{\|f_0(\sigma + \sigma') - f_0(\sigma) - h_0(\sigma')\|_2^2 + \|f_1(\sigma + \sigma') - f_1(\sigma) - h_1(\sigma')\|_2^2}}{\|\sigma'\|_2} \\ &= \lim_{\sigma' \rightarrow 0} \sqrt{\left(\frac{\|f_0(\sigma + \sigma') - f_0(\sigma) - h_0(\sigma')\|_2}{\|\sigma'\|_2}\right)^2 + \left(\frac{\|f_1(\sigma + \sigma') - f_1(\sigma) - h_1(\sigma')\|_2}{\|\sigma'\|_2}\right)^2} \\ &= \sqrt{\left(\lim_{\sigma' \rightarrow 0} \frac{\|f_0(\sigma + \sigma') - f_0(\sigma) - h_0(\sigma')\|_2}{\|\sigma'\|_2}\right)^2 + \left(\lim_{\sigma' \rightarrow 0} \frac{\|f_1(\sigma + \sigma') - f_1(\sigma) - h_1(\sigma')\|_2}{\|\sigma'\|_2}\right)^2} \\ &= 0. \end{aligned}$$

Thus, f is differentiable at σ , as desired.

It remains to prove the local Lipschitzness of g . Pick $\sigma \in \text{dom}(g)$. Then, g_0 and g_1 are defined at σ and they are locally Lipschitz. Thus, there exist open sets $O_0 \subseteq \text{dom}(g_0)$ and $O_1 \subseteq \text{dom}(g_1)$ and constants $B_0, B_1 > 0$ such that σ belongs to both O_0 and O_1 , and for all $\sigma_0, \sigma'_0 \in O_0$ and $\sigma_1, \sigma'_1 \in O_1$,

$$\|g_0(\sigma_0) - g_0(\sigma'_0)\|_2 \leq B_0 \|\sigma_0 - \sigma'_0\|_2 \quad \text{and} \quad \|g_1(\sigma_1) - g_1(\sigma'_1)\|_2 \leq B_1 \|\sigma_1 - \sigma'_1\|_2.$$

Let $O \triangleq O_0 \cap O_1$. The set O is open, and contains σ . Furthermore, for all $\sigma', \sigma'' \in O$,

$$\begin{aligned} \|g(\sigma') - g(\sigma'')\|_2 &= \sqrt{\|g_0(\sigma') - g_0(\sigma'')\|_2^2 + \|g_1(\sigma') - g_1(\sigma'')\|_2^2} \\ &\leq \sqrt{B_0^2 \|\sigma' - \sigma''\|_2^2 + B_1^2 \|\sigma' - \sigma''\|_2^2} \\ &= \sqrt{B_0^2 + B_1^2} \cdot \|\sigma' - \sigma''\|_2. \end{aligned}$$

Thus, g is Lipschitz in O , as desired.

Case of Assumption 5. Consider $K, K', L \subseteq \text{Var}$ with $K \subseteq K'$, and $\tau \in \text{St}[K' \setminus K]$. Let f and g be partial functions in $[\text{St}[K'] \rightarrow \text{St}[L]]$ such that $f \in \phi_{K',L}^{(d)}$ and $g \in \phi_{K',L}^{(l)}$. Let

$$\begin{aligned} f_1 : \text{St}[K] \rightarrow \text{St}[L], \quad & f_1(\sigma) \triangleq \begin{cases} f(\sigma \oplus \tau) & \text{if } \sigma \oplus \tau \in \text{dom}(f) \\ \text{undefined} & \text{otherwise,} \end{cases} \\ g_1 : \text{St}[K] \rightarrow \text{St}[L], \quad & g_1(\sigma) \triangleq \begin{cases} g(\sigma \oplus \tau) & \text{if } \sigma \oplus \tau \in \text{dom}(g) \\ \text{undefined} & \text{otherwise.} \end{cases} \end{aligned}$$

We should show that f_1 and g_1 satisfy $\phi_{K,L}^{(d)}$ and $\phi_{K,L}^{(l)}$, respectively. Note that

$$\text{dom}(f_1) = \{\sigma \mid \sigma \oplus \tau \in \text{dom}(f)\} \quad \text{and} \quad \text{dom}(g_1) = \{\sigma \mid \sigma \oplus \tau \in \text{dom}(g)\}.$$

These two sets are open since $\text{dom}(f)$ and $\text{dom}(g)$ are open and for any open O , the slice $\{\sigma \in \text{St}[K] \mid \sigma \oplus \tau \in O\}$ is open. Let $\sigma_0 \in \text{dom}(f_1)$ and $\sigma_1 \in \text{dom}(g_1)$. We will show that f_1 is differentiable at σ_0 , and g_1 is Lipschitz in an open neighbourhood of σ_1 .

Since f is differentiable and $\sigma_0 \oplus \tau \in \text{dom}(f)$, there exists a linear map $h : \text{St}[K'] \rightarrow \text{St}[L]$ such that $\text{dom}(h)$ is open and contains 0, and

$$\lim_{\sigma' \rightarrow 0} \frac{\|f(\sigma_0 \oplus \tau + \sigma') - f(\sigma_0 \oplus \tau) - h(\sigma')\|_2}{\|\sigma'\|_2} = 0.$$

Let $\tau_0 \triangleq \lambda v \in K' \setminus K$, 0, and h_1 be the partial function from $\text{St}[K]$ to $\text{St}[L]$ defined by

$$h_1(\sigma) \triangleq \begin{cases} h(\sigma \oplus \tau_0) & \text{if } \sigma \oplus \tau_0 \in \text{dom}(h), \\ \text{undefined} & \text{otherwise.} \end{cases}$$

Then, h_1 is linear, its domain is open (since taking a slice of an open set in $\text{St}[K] \cong \mathbb{R}^{|K'|}$ by fixing some coordinate variables gives an open set), and

$$\begin{aligned} \lim_{\sigma'' \rightarrow 0} \frac{\|f_1(\sigma_0 + \sigma'') - f_1(\sigma_0) - h_1(\sigma'')\|_2}{\|\sigma''\|_2} &= \lim_{\sigma'' \rightarrow 0} \frac{\|f(\sigma_0 \oplus \tau + \sigma'' \oplus \tau_0) - f(\sigma_0 \oplus \tau) - h(\sigma'' \oplus \tau_0)\|_2}{\|\sigma'' \oplus \tau_0\|_2} \\ &= 0. \end{aligned}$$

This means that f_1 is differentiable at σ_0 .

Since g is locally Lipschitz and $\sigma_1 \oplus \tau \in \text{dom}(g)$, there exists an open subset O of $\text{dom}(g)$ such that O contains $\sigma_1 \oplus \tau$ and g is Lipschitz in O , that is, there exists a real number $B > 0$ such that

$$\|g(\sigma) - g(\sigma')\|_2 \leq B \cdot \|\sigma - \sigma'\|_2$$

for all $\sigma, \sigma' \in O$. Let

$$O' \triangleq \{\sigma \in \text{St}[K] \mid \sigma \oplus \tau \in O\}.$$

Then, O' is open, and it contains σ_1 . Furthermore, for all $\sigma, \sigma' \in O'$,

$$\|g_1(\sigma) - g_1(\sigma')\|_2 = \|g(\sigma \oplus \tau) - g(\sigma' \oplus \tau)\|_2 \leq B \cdot \|\sigma \oplus \tau - \sigma' \oplus \tau\|_2 = B \cdot \|\sigma - \sigma'\|_2.$$

Thus, g_1 is Lipschitz in O' , as desired.

Case of Assumption 6. For the composition condition, we handle the differentiability case only. The other case can be proved similarly. Consider

$$K, L, M \subseteq \text{Var}, \quad f \in [\text{St}[K] \rightarrow \text{St}[L]], \quad \text{and} \quad g \in [\text{St}[L] \rightarrow \text{St}[M]].$$

Assume that $f \in \phi_{K,L}^{(d)}$ and $g \in \phi_{L,M}^{(d)}$. Let h be the standard composition of partial functions g and f . We should show that $h \in \phi_{K,M}^{(d)}$ as well, that is, $\text{dom}(h)$ is open and h is differentiable on its domain. Note that

$$\text{dom}(h) = \text{dom}(f) \cap f^{-1}(\text{dom}(g)).$$

Since $g \in \phi_{L,M}^{(d)}$, the set $\text{dom}(g)$ is open. Because $f \in \phi_{K,L}^{(d)}$, $\text{dom}(f)$ is open and f is continuous on its domain. The latter implies that $f^{-1}(\text{dom}(g))$ is open as well. Thus, the intersection of $\text{dom}(f)$ and $f^{-1}(\text{dom}(g))$ is open as desired. The differentiability of h on its domain holds since the restriction of f to $\text{dom}(h)$ gives a differentiable total function from $\text{dom}(h)$ to $\text{dom}(g)$, that of g to $\text{dom}(g)$ is also a differentiable total function, and the composition of two differentiable functions is differentiable.

Case of Assumption 7. The empty set is open, and the empty function is jointly differentiable and locally Lipschitz continuous. Thus, the strictness assumption holds for both predicate families. \square

H DEFERRED RESULTS IN §6

H.1 Proof of Theorem 6.2

PROOF OF THEOREM 6.2. Suppose that the algorithm returns π (without an error message). We should show the conclusion that π is simple and satisfies (R2) and (R3).

We make several observations before proving the conclusion. Let

$$(\mathbb{P}_m, \mathbb{d}_m, \mathbb{V}_m) \triangleq \llbracket c_m \rrbracket^\sharp, \quad (\mathbb{P}_g, \mathbb{d}_g, \mathbb{V}_g) \triangleq \llbracket c_g \rrbracket^\sharp, \quad (\overline{\mathbb{P}}_g, \overline{\mathbb{d}}_g, \overline{\mathbb{V}}_g) \triangleq \llbracket \overline{c}_g^\pi \rrbracket^\sharp,$$

and $K \subseteq \text{Var}$ be the set defined in Eq. (12). Also, let S_r be the set of names that the algorithm uses to construct the returned reparameterisation plan π . Then, by the algorithm, we have the two inclusions in Eq. (12) and Eq. (13), and also $\pi = \pi_0[S_r]$. In addition, $S_r \subseteq K$ since

$$S_r \subseteq \{(\alpha, i) \in \text{Name} \mid \text{for all } i' \in \mathbb{N}, (\alpha, i') \in \text{Name} \implies (\alpha, i') \in K\} \subseteq K.$$

We now prove the conclusion in three parts.

First part: We show that π is simple. To show this, consider $(n, d, l), (n', d', l') \in \text{NameEx} \times \text{DistEx} \times \text{LamEx}$ such that $n = \text{name}(\alpha, e)$ and $n' = \text{name}(\alpha, e')$ for some $\alpha \in \text{Str}, e, \text{ and } e'$. Suppose that $(n, d, l) \in \text{dom}(\pi)$. We should show $(n', d', l') \in \text{dom}(\pi)$. Since $(n, d, l) \in \text{dom}(\pi) = \text{dom}(\pi_0[S_r])$, we have $(n, d, l) \in \text{dom}(\pi_0)$ and $(\alpha, _) \in S_r$. This implies that $(n', d', l') \in \text{dom}(\pi_0)$ because π_0 is simple. Since $n' = \text{name}(\alpha, _)$ and $(\alpha, _) \in S_r$, we have $(n', d', l') \in \text{dom}(\pi)$ as desired.

Second part: We show that π satisfies (R2) in three steps.

First step: We prove $\theta \cup \text{rv}(\pi) \subseteq K$. To do so, observe that the S in the algorithm always satisfies the following property:

$$(\alpha, i) \in S \implies (\alpha, i') \in S \quad \text{for any } (\alpha, i), (\alpha, i') \in \text{Name}. \quad (28)$$

We can prove this by induction: the initial S (i.e., $S = \{(\alpha, i) \in \text{Name} \mid \text{for all } i' \in \mathbb{N}, (\alpha, i') \in \text{Name} \implies (\alpha, i') \in K\}$) satisfies the property, and each update of S (i.e., $S \leftarrow S \setminus \{(\alpha, i) \in \text{Name} \mid (\alpha, _) \in S\}$) preserves the property. From this, we obtain $\text{rv}(\pi) = \text{rv}(\pi_0) \cap S$:

$$\begin{aligned} \text{rv}(\pi) &= \text{rv}(\pi_0[S]) \\ &= \{(\alpha, i) \in \text{Name} \mid \exists e, d, l. (\text{name}(\alpha, e), d, l) \in \text{dom}(\pi_0[S])\} \\ &= \{(\alpha, i) \in \text{Name} \mid \exists e, d, l. ((\text{name}(\alpha, e), d, l) \in \text{dom}(\pi_0) \wedge \exists i'. (\alpha, i') \in S)\} \\ &= \{(\alpha, i) \in \text{Name} \mid (\exists e, d, l. ((\text{name}(\alpha, e), d, l) \in \text{dom}(\pi_0)) \wedge (\exists i'. (\alpha, i') \in S))\} \\ &= \{(\alpha, i) \in \text{Name} \mid (\text{name}(\alpha, _), _, _) \in \text{dom}(\pi_0)\} \cap \{(\alpha, i) \in \text{Name} \mid (\alpha, _) \in S\} \\ &= \text{rv}(\pi_0) \cap S, \end{aligned}$$

where the second and third equalities use the definitions of $\text{rv}(-)$ and $\pi_0[S]$, respectively, and the last equality uses Eq. (28). Since $\text{rv}(\pi) \subseteq S \subseteq K$ and $\theta \subseteq K$ by Eq. (12) (the inclusion of θ), we get $\theta \cup \text{rv}(\pi) \subseteq K$ as desired.

Second step: We prove that for all $u \in \{\text{like}\} \cup \{pr_\mu \mid \mu \in \text{Name}\}$ and $v \in \{pr_\mu \mid \mu \in \text{Name}\}$, the following functions (which are total since c_m and c_g always terminate) are differentiable with respect to the variables in $\theta \cup \text{rv}(\pi)$ jointly:

$$\begin{aligned} (\sigma_\theta, \sigma_n) \in \text{St}[\theta] \times \text{St}[\text{Name}] &\longmapsto \llbracket c_m \rrbracket(\sigma_{p \setminus \theta} \oplus \sigma_\theta \oplus \sigma_n \oplus g(\sigma_n))(u), \\ (\sigma_\theta, \sigma_n) \in \text{St}[\theta] \times \text{St}[\text{Name}] &\longmapsto \llbracket c_g \rrbracket(\sigma_{p \setminus \theta} \oplus \sigma_\theta \oplus \sigma_n \oplus g(\sigma_n))(v), \end{aligned} \quad (29)$$

where $\sigma_{p \setminus \theta} \triangleq (\lambda v \in \text{PVar} \setminus \theta. 0)$ and the function $g : \text{St}[\text{Name}] \rightarrow \text{St}[\text{AVar}]$ takes σ_n and returns σ_a such that σ_a maps *like* to 1, pr_μ to $\mathcal{N}(\sigma_n(\mu); 0, 1)$, val_μ to $\sigma_n(\mu)$, and all the other variables to 0. The state $\sigma_{p \setminus \theta} \oplus g(\sigma_n)$ is the very initialisation used in Eq. (3). This step consists of two substeps.

First substep: We first show that for all $u \in \{like\} \cup \{pr_\mu \mid \mu \in \text{Name}\}$ and $v \in \{pr_\mu \mid \mu \in \text{Name}\}$, the functions $f_m, f_g : \text{St}[\theta] \times \text{St}[\text{Name}] \times \text{St}[\text{AVar}] \rightarrow \mathbb{R}$ are differentiable with respect to the variables in $\theta \cup rv(\pi)$ jointly:

$$\begin{aligned} f_m(\sigma_\theta, \sigma_n, \sigma_a) &\triangleq \llbracket c_m \rrbracket(\sigma_{p \setminus \theta} \oplus \sigma_\theta \oplus \sigma_n \oplus \sigma_a)(u), \\ f_g(\sigma_\theta, \sigma_n, \sigma_a) &\triangleq \llbracket c_g \rrbracket(\sigma_{p \setminus \theta} \oplus \sigma_\theta \oplus \sigma_n \oplus \sigma_a)(v). \end{aligned} \quad (30)$$

Note that in f_m and f_g , the AVar part does not depend on the Name part (unlike in Eq. (29)). For the proof, pick arbitrary $u \in \{like\} \cup \{pr_\mu \mid \mu \in \text{Name}\}$ and $v \in \{pr_\mu \mid \mu \in \text{Name}\}$. Then, $\theta \cup rv(\pi) \subseteq K \subseteq \mathbb{P}_m(u) \cap \mathbb{P}_g(v)$, where the first inclusion is from the above result and the second from Eq. (12) (the definition of K). By the soundness of differentiability analysis (Theorem 5.8), $\models \Phi(\llbracket c_m \rrbracket, \mathbb{P}_m(u), \{u\})$ and $\models \Phi(\llbracket c_g \rrbracket, \mathbb{P}_g(v), \{v\})$. From this, and by the weakening lemma of Φ with $\theta \cup rv(\pi) \subseteq \mathbb{P}_m(u) \cap \mathbb{P}_g(v)$ (Lemma F.4), we have $\models \Phi(\llbracket c_m \rrbracket, \theta \cup rv(\pi), \{u\})$ and $\models \Phi(\llbracket c_g \rrbracket, \theta \cup rv(\pi), \{v\})$. Hence, the functions in Eq. (30) are differentiable with respect to $\theta \cup rv(\pi)$ jointly as desired, by the definition of Φ (§5.1) and the definition of “ $f : \text{St}[L] \rightarrow \mathbb{R}$ for $L \subseteq \text{Var}$ is differentiable with respect to $L' \subseteq L$ jointly” (§4.2).

Second substep: We now prove that the claim of the second step follows from the first substep just proved. Pick any $u \in \{like\} \cup \{pr_\mu \mid \mu \in \text{Name}\}$. We should show that the first function in Eq. (29) is differentiable with respect to $\theta \cup rv(\pi)$. Note that we should also show the same for the second function (for any v), but the proof is similar to the first function so we omit this case. To prove the claim for the first function, pick any $\xi'_{n,0} \in \text{St}[\text{Name} \setminus rv(\pi)]$ and $\sigma_{a,0} \in \text{St}[\text{AVar}]$. Define $f' : \text{St}[\theta] \times \text{St}[rv(\pi)] \times \text{St}[\text{AVar}] \rightarrow \mathbb{R}$ as

$$f'(\sigma_\theta, \xi_n, \sigma_a) \triangleq f(\sigma_\theta, \xi_n \oplus \xi'_{n,0}, \sigma_a).$$

Then, by Lemma C.6-(2) and Lemma C.6-(3),

$$f'(\sigma_\theta, \xi_n, \sigma_a) \triangleq \begin{cases} f(\sigma_\theta, \xi_n \oplus \xi'_{n,0}, \sigma_{a,0}) & \text{if } (\sigma_\theta, \xi_n) \in U \\ \text{proj}(\sigma_a) & \text{if } (\sigma_\theta, \xi_n) \notin U \end{cases}$$

for some $U \subseteq \text{St}[\theta] \times \text{St}[rv(\pi)]$ and some projection map $\text{proj} : \text{St}[\text{AVar}] \rightarrow \mathbb{R}$. Also, since f is differentiable with respect to $\theta \cup rv(\pi)$, $f'(-, -, \sigma_a) : \text{St}[\theta] \times \text{St}[rv(\pi)] \rightarrow \mathbb{R}$ is differentiable and thus continuous for all $\sigma_a \in \text{St}[\text{AVar}]$. From these, Lemma H.1 is applicable to f' , implying that U should be either \emptyset or $\text{St}[\theta] \times \text{St}[rv(\pi)]$. We now consider $f'' : \text{St}[\theta] \times \text{St}[rv(\pi)] \rightarrow \mathbb{R}$ defined by

$$f''(\sigma_\theta, \xi_n) \triangleq f'(\sigma_\theta, \xi_n, g(\xi_n) \oplus g(\xi'_{n,0})),$$

where g is extended to accept a substate in $\text{St}_\square[\text{Name}]$ and return a substate for the corresponding auxiliary part. Then, to prove the claim, it suffices to show that f'' is differentiable (since $\xi'_{n,0}$ was chosen arbitrarily). We do case analysis on U . If $U = \text{St}[\theta] \times \text{St}[rv(\pi)]$, then

$$f''(\sigma_\theta, \xi_n) = f(\sigma_\theta, \xi_n \oplus \xi'_{n,0}, \sigma_{a,0}) \quad \text{for all } \sigma_\theta \text{ and } \xi_n;$$

since f is differentiable with respect to $\theta \cup rv(\pi)$, f'' is differentiable. If $U = \emptyset$, then

$$f''(\sigma_\theta, \xi_n) = \text{proj}(g(\xi_n) \oplus g(\xi'_{n,0})) \quad \text{for all } \sigma_\theta \text{ and } \xi_n;$$

since g , \oplus , and proj are all differentiable (because g only uses projection and the density of the standard normal distribution), f'' is differentiable. Hence, f'' is differentiable in both cases, and this shows the claim of the second step.

Third step: We prove that π satisfies (R2), i.e., the functions in (R2) are differentiable with respect to $\theta \cup rv(\pi)$ jointly. This holds because: c_m and c_g do not have a double-sampling error, so each function in (R2) is a multiplication of some of the functions in Eq. (29) (for different u and v); the functions in Eq. (29) are differentiable with respect to $\theta \cup rv(\pi)$ jointly (by the above result); and multiplication preserves differentiability.

Third part: We show that π satisfies (R3), i.e., the functions in (R3) are differentiable with respect to θ jointly. For this, it suffices to show the claim that for all $v \in \{pr_\mu, val_\mu \mid \mu \in \text{Name}\}$ and $\sigma_n \in \text{St}[\text{Name}]$, the following function is differentiable with respect to θ jointly:

$$\sigma_\theta \in \text{St}[\theta] \mapsto \llbracket \overline{c_g^\pi} \rrbracket (\sigma_{p \setminus \theta} \oplus \sigma_\theta \oplus \sigma_n \oplus g(\sigma_n))(v). \quad (31)$$

This implies (R3) because $\overline{c_g^\pi}$ does not have a double-sampling error, so each function in (R2) is either a multiplication or a pairing of the function in Eq. (31) (for different v); and multiplication and pairing preserve differentiability. To show the claim, consider any $v \in \{pr_\mu, val_\mu \mid \mu \in \text{Name}\}$. Then, $\theta \subseteq \overline{p_g}(v)$ by Eq. (13). By the soundness of differentiability analysis (Theorem 5.8), $\models \Phi(\llbracket \overline{c_g^\pi} \rrbracket, \overline{p_g}(v), \{v\})$. From this, and by the weakening lemma of Φ with $\theta \subseteq \overline{p_g}(v)$ (Lemma F.4), we have $\models \Phi(\llbracket \overline{c_g^\pi} \rrbracket, \theta, \{v\})$. Hence, the function in Eq. (31) is differentiable with respect to θ jointly. This completes the overall proof. \square

LEMMA H.1. *Let $f : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}$ be a function such that*

$$f(x, y) = \begin{cases} f_1(x) & \text{if } x \in U \\ f_2(y) & \text{if } x \notin U \end{cases}$$

for some $f_1 : \mathbb{R}^n \rightarrow \mathbb{R}$, $f_2 : \mathbb{R}^m \rightarrow \mathbb{R}$, and $U \subseteq \mathbb{R}^n$. Suppose that $f_2(\mathbb{R}^m) = \mathbb{R}$ and $f(-, y) : \mathbb{R}^n \rightarrow \mathbb{R}$ is continuous for all $y \in \mathbb{R}^m$. Then, U is either \emptyset or \mathbb{R}^n .

PROOF. Here is a sketch of the proof. We prove the lemma by contradiction. Suppose that U is neither \emptyset nor \mathbb{R}^n . Then, the boundary of U (i.e., $\text{bd}(U) \subseteq \mathbb{R}^n$) is nonempty, since the boundary of a set is empty if and only if the set is both open and closed, and since \emptyset and \mathbb{R}^n are the only subsets of \mathbb{R}^n that are both open and closed. Let $x \in \text{bd}(U)$ and consider two cases: $x \in U$ or $x \notin U$. In each of the two cases, we can show that there exists $y \in \mathbb{R}^m$ such that $f(-, y)$ is not continuous at x . When showing the discontinuity, we use the following: $x \in \text{bd}(U)$; the specific way that f is defined (in terms of U , f_1 , and f_2); and the assumption that $f_2(\mathbb{R}^m) = \mathbb{R}$. By assumption, $f(-, y)$ should be continuous over the entire \mathbb{R}^n , so we get contradiction. \square

I DEFERRED RESULTS IN §7

I.1 Deferred Experiment Details and Results

Table 7. Pyro examples used in experiments and their key features (continued from Table 3). The last five columns show the total number of code lines (excluding comments), loops, sample commands, observe commands, and learnable parameters (declared explicitly by `pyro.param` or implicitly by a neural network module). Each number is the sum of the counts in the model and guide.

Name	Probabilistic model	LoC	while	sam	obs	param
dpm	Dirichlet process mixture models	27	0	6	1	4
vae	Variational autoencoder (VAE)	35	0	2	1	5
csis	Inference compilation	38	0	2	2	5
br	Bayesian regression	42	0	10	1	5
lda	Amortised latent Dirichlet allocation	57	0	8	1	5
prolda	Probabilistic topic modelling	58	0	2	1	5
ssvae	Semi-supervised VAE	60	0	4	1	7

Table 8. Results of smoothness analyses (continued from Table 4). “Manual” and “Ours” denote the number of continuous random variables and learnable parameters in which the density of the program is smooth, computed by hand and by our analyser. “Time” denotes the runtime of our analyser in seconds. “#CRP” denotes the total number of continuous random variables and learnable parameters in the program. -m and -g denote model and guide. We consider $\{(\alpha, i) \in \text{Name}\}$ as one random variable for each $\alpha \in \text{Name}$.

Name	Differentiable			Locally Lipschitz			#CRP
	Manual	Ours	Time	Manual	Ours	Time	
dpmm-m	2	2	0.002	2	2	0.002	2
dpmm-g	6	6	0.003	6	6	0.003	6
vae-m	3	3	0.002	3	3	0.003	3
vae-g	4	4	0.002	4	4	0.002	4
csis-m	1	1	0.001	1	1	0.001	1
csis-g	2	2	0.004	6	6	0.004	6
br-m	5	5	0.002	5	5	0.002	5
br-g	10	10	0.004	10	10	0.004	10
lda-m	3	3	0.002	3	3	0.002	3
lda-g	7	7	0.007	7	7	0.007	7
proldda-m	2	2	0.008	2	2	0.007	2
proldda-g	5	5	0.007	5	5	0.006	5
ssvae-m	3	3	0.004	3	3	0.003	3
ssvae-g	6	6	0.007	6	6	0.009	6

Table 9. Results of variable selections (continued from Table 5). “Ours-Time” denote the runtime of our variable selector in seconds. “Ours-Sound” and “Pyro \ Ours” denote the number of random variables in the example that are in π_{ours} , and that are in π_0 but not in π_{ours} , respectively, where π_{ours} and π_0 denote the reparameterisation plans given by our variable selector and by Pyro. “Pyro \ Ours” is partitioned into “Sound” and “Unsound”: the latter denotes the number of random variables that make (R2’) or (R3’) violated when added to π_{ours} , and the former denotes the number of the rest. “#CR” and “#DR” denote the total number of continuous and discrete random variables in the example. We consider $\{(\alpha, i) \in \text{Name}\}$ as one random variable for each $\alpha \in \text{Name}$.

Name	Ours		Pyro \ Ours		#CR	#DR
	Time	Sound	Sound	Unsound		
dpmm	0.007	2	0	0	2	1
vae	0.004	1	0	0	1	0
csis	0.014	1	0	0	1	0
br	0.009	5	0	0	5	0
lda	0.011	3	0	0	3	1
proldda	0.018	1	0	0	1	0
ssvae	0.013	1	0	0	1	1