Inflation: a Python library for classical and quantum causal compatibility

Emanuel-Cristian Boghiu[1], Elie Wolfe[2], and Alejandro Pozas-Kerstjens[3]

[1]ICFO – Institut de Ciencies Fotoniques, The Barcelona Institute of Science and Technology, 08860 Castelldefels (Barcelona), Spain

[2]Perimeter Institute for Theoretical Physics, 31 Caroline St. N., Waterloo, Ontario, Canada, N2L 2Y5

[3]Instituto de Ciencias Matemáticas (CSIC-UAM-UC3M-UCM), 28049 Madrid, Spain

We introduce Inflation, a Python library for assessing whether an observed probability distribution is compatible with a causal explanation. This is a central problem in both theoretical and applied sciences, which has recently witnessed significant advances from the area of quantum nonlocality, namely, in the development of inflation techniques. Inflation is an extensible toolkit that is capable of solving pure causal compatibility problems and optimization over (relaxations of) sets of compatible correlations in both the classical and quantum paradigms. The library is designed to be modular and with the ability of being ready-to-use, while keeping an easy access to low-level objects for custom modifications.

1 Motivation

One of the main challenges in any scientific discipline is identifying which are the causes behind some observed correlations. Is a vaccine effective against a disease? Does raising salaries encourage spending? Is the increase of carbon dioxide in the atmosphere responsible for the increase in the average temperature of the Earth? These questions and their like can all be phrased and analyzed using the tools of causal inference (CI) [1]. However, despite the wide relevance of causal inference, current CI algorithms involving latent variables are typically incapable of analyzing structures with more than a small number of nodes [2–6].

The field of quantum nonlocality [7] has recently put its attention in causality, in light of the fact that Bell's theorem [8] can be understood in terms of the compatibility of probability distributions with a given causal structure [9, 10]. This view has propelled the study of quantum correlations beyond the traditional bipartite scenario (see, e.g., [11–15], and the review [16]), and the development of techniques for characterizing the quantum and classical probability distributions that can be generated in such causal scenarios [17–21]. A particularly successful tool is the inflation method [22–24], which consists of a series of increasingly strict necessary conditions that can be tested via linear or semidefinite programming. Despite its broad applicability within and outside the field of quantum nonlocality, available implementations of the inflation technique are typically limited in

Emanuel-Cristian Boghiu: cristian.boghiu@icfo.eu

Elie Wolfe: ewolfe@pitp.ca

Alejandro Pozas-Kerstjens: physics@alexpozas.com

terms of the type of causal structures it applies to, or in the type of inflations considered (see, e.g., [25, 26]). This means that researchers must code their own programs every time they seek to analyze a different structure or try a different solution, adding an extra level of difficulty to the application of the technique.

Here we present Inflation [27], an open-source library, written in Python, that implements the inflation framework for causal compatibility. It allows both for solving feasibility problems (i.e., answering the question *"can I generate this distribution in this causal structure?"*) and for bounding optimal values of functionals over distributions compatible with a causal structure. These include all network scenarios considered in the field of quantum nonlocality [16], and structures that have recently gained attention in that field, such as the so-called instrumental scenario [15].

Currently, Inflation implements the quantum inflation hierarchy of Ref. [23]. This means that the main focus is the characterization of distributions generated by measuring quantum systems. However, the library can also be used, by setting the corresponding flags, for assessing compatibility with distributions generated when all the latent nodes represent sources of classical shared randomness (thereby extending the ideas in [28, 29]), and when all the latent nodes represent sources of quantum entanglement and at the same time all parties are correlated through a global source of classical shared randomness.

This paper presents the package and illustrates simple use cases, which are extended in the library documentation. It is structured as follows: in Sec. 2 we provide a brief description of the theoretical ideas behind inflation and point to the relevant literature. In Sec. 3 we show how to get started with the library and describe its main components and features. In Secs. 4 and 5 we demonstrate with code snippets the different types of problems that can be addressed with Inflation. We discuss further software details and library information in Sec. 6 including future development, contribution guidelines, and planned maintenance and support, and we provide some concluding remarks in Sec. 7.

2 The inflation framework

The aim of this section is to provide a brief description of the general ideas behind inflation. For the reader interested in the details of the different variants, we refer to the original publications [22, 23, 30].

Inflation is a general framework for analyzing correlations that can be generated in causal scenarios. These scenarios feature latent nodes, corresponding to physical systems (sources of shared classical randomness, quantum states, or states of more general physical theories), and visible nodes, which represent random variables that describe the outcomes of measurements performed on such systems. The nodes are connected by arrows that denote which systems are sent to which parties. In order to avoid causal paradoxes, like the outcome of a measurement determining which measurement is performed in the first place, the paths created by following the arrows must not contain closed loops. These graphs with directed arrows and no loops are known as directed acyclic graphs, or DAGs. Examples of such DAGs can be found in Figs. 1 and 2.

In order to characterize the distributions that can be generated in a particular causal scenario, inflation considers the *gedankenexperiment* where one has access to multiple copies of the physical systems (recall, the latent nodes in the corresponding DAG) and operations (the visible nodes) used in it. By connecting these elements in such a way that the parents of a copy of a visible node are copies of the parents of that node in the original scenario, one constructs *inflation scenarios* where the characterization of compatible distributions is simpler than in the original scenario. Furthermore, constraints satisfied by compati-

ble distributions can be translated into necessary constraints of distributions compatible with the original scenario. The distributions compatible with inflated scenarios satisfy two main properties. First, they are highly symmetric, due to the fact that the elements of the inflation scenarios are copies of the elements of the original scenario. Namely, distributions compatible with inflations are invariant under permutations of copies of a same latent node. Second, when marginalized over sets of nodes that reproduce parts of the original scenario, they coincide with the corresponding marginals of the distribution in the original scenario. These two properties can be encoded by means of linear equalities and inequalities. Therefore, probability distributions that can be generated in inflations can be characterized by means of linear programming in the case where the physical systems are classical [22] or described by a generalized physical theory [24], or with hierarchies of semidefinite programs in the case where the systems measured by the parties are quantum mechanical [23]. In all these cases, efficient algorithms exist that allow to solve the corresponding problems with standard computing resources.

Inflation is typically used for two families of problems: optimizing quantities over distributions compatible with some causal scenario, and determining whether a particular distribution admits a realization in a given scenario. For the former, inflation provides bounds to the actual optima (which can be nevertheless arranged in monotonic sequences, see e.g. Refs. [23, 30]). For the latter, finding an inflation where all the symmetry and marginal constraints implied cannot be satisfied is a proof that the original premise, namely that the distribution under scrutiny admitted a realization in the causal structure, is false. Proving the opposite, namely that the original distribution admits a realization in the original causal scenario, requires finding an explicit realization (in terms of classical shared randomness, quantum systems, or elements of a generalized probabilistic theory, as corresponding) of the distribution. Exploring such constructions is a task outside the scope inflation, and thus of this library.

Inflation currently implements the quantum inflation hierarchy of Ref. [23]. This is, it considers that all sources in the scenario distribute quantum systems with an unbounded dimension. As such, whenever a correlation can be achieved using positive operator values measurements (POVMs) the same correlation can also be realized by projective-valued measures (PVMs), as any POVM can be dilated to a PVM through Naimark's dilation theorem. As such, Inflation models all measurements as projective without loss of generality. This leads to a hierarchy that is monotonic (so each level is, at least, as constraining as the previous one), but that until now it has not been proven to converge except in concrete scenarios [31]. There also exist alternative hierarchies that have been proven to converge [32], although in a non-monotonic way.

3 The library

3.1 Requirements and installation

Inflation is a Python library that can be installed on Mac, Windows, and Linux operating systems via `pip` by executing the instruction below at a command line.

```
pip install inflation
```

The core requirements of Inflation are NumPy [33] (used for general numerical procedures), SymPy [34] (used to make the input format more user-friendly), and SciPy [35] (used for handling sparse matrices). It can also use Numba [36] as a just-in-time compiler to speed up core calculations. For solving the generated relaxations, the library uses the MOSEK Fusion API [37] to solve the linear and semidefinite programming problems. It also allows

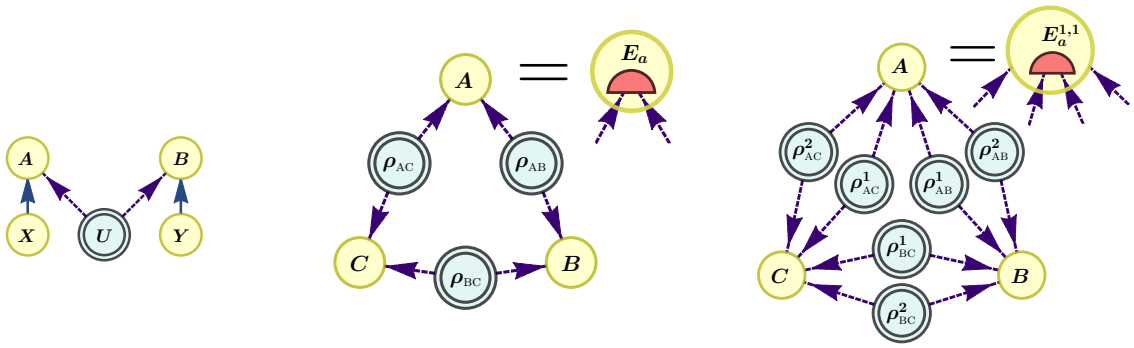

(a) (b) (c)

Figure 1: (a) The bipartite Bell scenario written as a DAG. The yellow circles denote visible nodes representing random variables, while the blue circles denote a latent node representing a physical system. This DAG represents the scenario where two parties perform measurements denoted by $x \sim X$ and $y \sim Y$, respectively, on shares of a bipartite physical system distributed to them. The results of the measurements are $a \sim A$ and $b \sim B$. (b) The triangle scenario, where three parties perform (in this case, fixed) measurements on shares of bipartite physical systems. (c) The second-order quantum inflation of the triangle scenario. Each of the sources is duplicated, and each party has a choice of the shares in which to perform their measurements. Whereas NPA-based convex relaxations of scenario (b) are indistinguishable from those of single common cause scenarios, the underlying symmetries in (c) allow to constrain correlations beyond the common-cause scenario.

for writing the problem in a human-readable form as a comma-separated values file, to a MATLAB-compatible .mat file, or to SDPA data format for further manipulation in other interfaces such as Yalmip [38].

To test installation, one can run the following.

```
import inflation
inflation.about()
```

These lines print basic information about the version of Inflation installed and the versions of installed dependencies.

```
# Inflation: Implementations of the Inflation Technique for Causal Inference
# ============================================================================
# Authored by: Emanuel-Cristian Boghiu, Elie Wolfe and Alejandro Pozas-Kerstjens

# Inflation Version:        0.1
#
# Core Dependencies
# -----------------
# NumPy Version:  1.23.1
# SciPy Version:  1.8.1
# SymPy Version:  1.11.1
# Numba Version:  0.56.2
# Mosek Version:  10.0.20
#
# Python Version: 3.10.4
# Platform Info:  Windows (AMD64)
```

The source code for Inflation is hosted on GitHub at

https://github.com/ecboghiu/inflation

and is distributed with an open-source software license: GNU GPL version 3.0. More details about the software, packaging information, and guidelines for contributing to Inflation are included in Sec. 6.

3.2 Components

There are two main layers in Inflation. First, the basic characterization of the causal scenario and its desired inflation are stored in an `InflationProblem`. This includes the DAG describing the causal structure, the number of inputs and outputs of each of its visible nodes, and the number of copies of each latent node in the desired inflation. If the causal structure furthermore contains visible-to-visible connections, then the procedures in, e.g., [30, Sec. V], [23, Sec. V] and [32, Sec. IV.C] are executed in order to find the network and suitable constraints that generate the same distributions as the original causal structure.

The second layer takes the characterization provided by an `InflationProblem` object and sets up and solves the compatibility or optimization problem of interest. Currently, the library only supports the quantum inflation hierarchy described in Ref. [23] via the `InflationSDP` object. However, the implementation also supports the analysis of distributions generated from classical latent nodes by imposing suitable commutation constraints (see Ref. [28] and [23, Sec. VI]). Sec. 5.1 showcases how this can be achieved in Inflation.

In the quantum inflation hierarchy, the characterization of the set of probability distributions is given by a list of operators, in the spirit of the Navascués-Pironio-Acín (NPA) hierarchy [39–41]. This is input to the `InflationSDP` object via the function `generate_relaxation()`, which admits either a generic list-of-lists notation for arbitrary lists of operators, or string-based notations for operator sets that are routinely used in the literature: `npa#` for the sets describing the NPA hierarchy (denoted as $\mathcal{T}_n$ in Ref. [39] and as $\mathcal{S}_n$ in Ref. [40]), `local#` for the so-called local levels (denoted as $\mathcal{L}_n$ in [23, App. C], see also Ref. [42]), and `physical#` for the sets of operators of bounded length whose expectation value is non-negative for any quantum state.

The function `generate_relaxation()` automatically imposes the equality constraints that are derived from the invariance of the inflation under permutation of copies of a same original element (this is, it imposes the constraints described in [23, Eq. (10)]). Furthermore, it identifies which marginals of distributions in the inflation must coincide with marginals of distributions in the original scenario. After this, the user can specify either a probability distribution over the visible nodes for a feasibility problem (i.e., to determine whether the distribution can be identified as incompatible with the causal structure) using the function `set_values()`, or a combination of operators whose expectation value will be optimized, by using the function `set_objective()`. When using `set_values()`, the user can choose to set also the so-called linearized polynomial constraints, which constrain the set of compatible distributions further at the expense of obtaining certificates with more limited applicability [26, 43].

3.3 Reductions of the feasible region

In general, the fact that one must consider multiple copies of the elements in the original causal structure leads to a large computational load when storing and solving the relevant problems. Inflation implements a number of additional constraints not included in the original definitions of the hierarchies, either automatically or at the user's choice, that give a tighter relaxation for the same level of computational resources. These constraints are:

Non-negativity of physical moments. This is a feature only relevant for the implementations of inflation that characterize distributions generated by measuring quantum systems. Implementations of quantum inflation require the use of the NPA hierarchy [39] in order to assess whether a compatible inflation exists. The main object in the NPA hierarchy are the so-called moment matrices, whose rows and columns are indexed by products

of (a priori unknown) projection operators. Each cell of the moment matrix contains the expectation of the product of the operators in the row and the operators in the column under an also unknown quantum state. Despite all elements being unknown a priori, one can derive constraints for some of them in certain situations. For instance, it is known that eigenpotent operators have a non-negative expectation value under any possible quantum state, and thus these non-negativity constraints are always imposed in Inflation. In fact, using operator products which have a non-negative expectation value under any quantum state in the generating set for `InflationSDP` leads to drastic reductions in problem size for certain problems, as we explicitly show in Sec. 4.1.

Sandwich-nonnegativity. The operator products described above have a non-negative expectation value also if they appear *sandwiched* between another product of operators and its conjugate. Indeed, if O_2 is a product of operators that has non-negative expectation value with any state, then $\langle\psi|O_1^\dagger O_2 O_1|\psi\rangle = \langle\psi'|O_2|\psi'\rangle \geq 0$ for any O_1, $|\psi\rangle$. An example of these products are those that correspond to subsequent measurements on a same state, whose expectation values represent the probabilities of sequences of outcomes. This feature is also automatically imposed when generating an instance of `InflationSDP`.

Linearized polynomial constraints. Both in the inflation methods for compatibility with quantum (based on semidefinite programming) and classical and generalized physical models (based on linear programming), it is possible to transform certain non-linear relations between the unknowns in the problems into linear ones when one assesses the compatibility of a given distribution with the scenario. When a subgraph of the inflation contains more than one connected components, and (at least) one of these components can be associated a numerical value from the distribution under scrutiny (these are known as injectable components in the terminology of Ref. [22]), the variables associated to the subgraph can be related to those corresponding to its non-injectable connected components via linear relations, reducing the feasible region. Linearized polynomial constraints are used, for instance, in Ref. [26] in the case of inflation for compatibility with classical models. For compatibility with quantum models, [43, App. D.2] contains a demonstration of its advantage in several scenarios. It should be noted that when using linearized polynomial constraints, in case of infeasibility, the dual of the semidefinite (or linear) program can only serve as a certificate of infeasibility for the tested distribution, given that for other distributions the feasible region reduction under linearized polynomial constraints can be different [26]. Inflation allows to impose linearized polynomial constraints by setting the flag `use_lpi_constraints=True` in `set_values()`.

4 Main functionality

We showcase here the user experience in using Inflation to solve a series of problems that routinely appear in causal inference scenarios. All these examples can be downloaded as ready-to-run Jupyter notebooks from the repository of the library (see Sec. 6.3).

4.1 Feasibility problems and extraction of certificates

The first example we will consider is that of demonstrating that a particular distribution can not be generated in a specific network arrangement when the sources (recall, the latent nodes in the scenario characterized by the DAG) distribute quantum systems, and the parties (the visible nodes) perform quantum measurements in the systems received.

We illustrate this type of problems by showing that the so-called W distribution, defined as

$$P_W(a,b,c) = \begin{cases} \frac{1}{3} & \text{if } a+b+c=1, \\ 0 & \text{otherwise} \end{cases}, \tag{1}$$

where $a, b, c \in \{0, 1\}$, is incompatible with the quantum triangle causal scenario of Fig. 1(b).

The first step is to specify the original scenario and its desired inflation by creating an instance of `InflationProblem`. This requires specifying (i) the DAG of the original scenario as a dictionary where the keys are parent nodes and the values are lists of the corresponding children, (ii) the number of possible outcomes and number of possible measurement settings of each of the visible nodes in the scenario, and (iii) the inflation levels, which represent the number of copies of each latent node that will be considered in the inflated scenario:

```
dag = {"rho_AB": ["A", "B"],
       "rho_BC": ["B", "C"],
       "rho_AC": ["A", "C"]}
nr_outputs = [2, 2, 2]
nr_inputs  = [1, 1, 1]
nr_copies  = [2, 2, 2]
scenario   = InflationProblem(dag, nr_outputs, nr_inputs, nr_copies,
                              order=['A', 'B', 'C'])
```

The order of the parties is specified via the `order` optional argument, and the order of the sources is taken to be the same as insertion order in the `dag` dictionary.

The generated instance of `InflationProblem` is fed to an `InflationSDP` instance. This object controls all the features related to the semidefinite relaxation of the problem considered. For instance, one can easily specify the relaxation obtained when using as generating monomials all products of operators of length at most 2 (this is, what is known as the second level in the NPA hierarchy [39, 40]):

```
sdp = InflationSDP(scenario)
sdp.generate_relaxation("npa2")
```

The last step is setting the constraints corresponding to observing the target probability distribution in marginals of the inflation distribution that can be identified with marginals of the original scenario. If one wants to do this for all marginals, this can be achieved with the function `set_distribution()`, but if more granularity is needed one must use the function `set_values()` instead (see Sec. 5.3 for an explicit example using this function). The distribution must be input as a multidimensional NumPy array `P[`$\texttt{out}_1$`,...,`$\texttt{out}_n$`,`$\texttt{in}_1$`,...,`$\texttt{in}_n$`]`, where each cell contains the corresponding probability, and where the i-th output and input, $\texttt{out}_i$ and $\texttt{in}_i$, correspond to the party in the i-th position in the `order` keyword argument.

```
sdp.set_distribution(P_W)
```

Then, running `sdp.solve()` executes the semidefinite program and stores its status in `sdp.status`. For the problem at hand, this is `infeasible`, meaning that the W distribution of Eq. (1) does not admit a second-order (because of our specification of `nr_copies`) quantum inflation of the triangle scenario, and thus by the inflation arguments (see Sec. 2 and Ref. [23]) it cannot be generated in the triangle causal scenario when the latent nodes represent sources of bipartite quantum states.

Once the problem is determined to be infeasible, semidefinite programming provides a certificate of such infeasibility. These certificates of infeasibility, which witness incompatibility with a given inflation, can be transformed into polynomial Bell-like inequalities

that witness distributions incompatible with the original causal scenario (see, e.g., [23, Sec. VII.B]). Inflation automatically computes these certificates and provides them in a variety of useful forms. For instance, running

```
sdp.certificate_as_probs(clean=True)
```

produces as output a SymPy object of the following form

$$-0.476p_A(0|0)^2 + 0.059p_A(0|0)p_{AB}(00|00) + 0.476p_A(0|0)p_{ABC}(000|000) + \cdots + 0.563,$$

which signals distributions that produce a negative value as incompatible with the triangle scenario. This object can be further manipulated easily with SymPy's built-in functions.

Feasibility as optimization. In terms of numerical stability, it is often advised to frame feasibility problems as optimization problems. This is specially relevant when the distribution whose compatibility we are testing is close to the boundary of the feasible set, where analytically feasible problems can be reported as infeasible.

Inflation allows for this framing by optimizing the smallest eigenvalue of the problem's moment matrix[1]. This is achieved by passing one argument at solving time:

```
sdp.solve(feas_as_optim=True)
```

The optimal (largest) value of the smallest eigenvalue of the moment matrix is stored in `sdp.objective_value` and by inspecting its value one can determine whether the original problem is feasible (if `sdp.objective_value` ≥ 0) or not (if `sdp.objective_value` < 0). Furthermore, the quantity in `sdp.objective_value` can provide a rudimentary notion of distance to the feasible set, and help estimating how the size of the feasible set changes when adding extra constraints to the problem or when changing the generating set. Importantly, the extraction of certificates is unaffected, and can still be performed even when treating feasibility problems as optimization problems.

Physical moments as generating set. It is known that, when considering the distribution in Eq. (1) subject to white noise,

$$\nu P_W + (1-\nu)\frac{1}{8}, \quad \nu \in [0,1], \tag{2}$$

quantum inflation of order 2 can certify its incompatibility with the triangle with quantum latent nodes at least down to $\nu = 0.8038$. This result was obtained in Ref. [23] by using a moment matrix of size 1175×1175. By using as monomials indexing the rows and columns of the moment matrix only those with non-negative expectation values under any quantum state as mentioned in Sec. 3, one can recover the same ν_{crit} but with a much smaller moment matrix, of size 287×287. Due to its notable gains in memory (and its consequent gains in speed), using the non-negative monomials in the generating set is made very easy in Inflation. In order to recover the result mentioned above with the much smaller moment matrix, one just needs to run:

```
genset = sdp.build_columns("physical2", max_monomial_length=4)
sdp.generate_relaxation(genset)
v = 0.8039
sdp.set_distribution(v*P_W + (1-v)/8)
sdp.solve()
# infeasible
```

[1] If the largest value that the smallest eigenvalue can achieve is negative, it means that the moment matrix associated to the problem cannot be made positive-semidefinite, and thus the corresponding feasibility problem is infeasible.

This improvement is most notable when dealing with problems that involve distributions without settings. In our experience, in the case where the parties in the problem have a choice of different measurements to perform on the states received by the sources, gains are more moderate.

4.2 Optimization of Bell operators

The second large class of problems that can be solved within Inflation is obtaining bounds on expectation values of Bell operators. For example, let us consider Mermin's operator [44]:

$$\text{Mermin} = \langle A_1 B_0 C_0 \rangle + \langle A_0 B_1 C_0 \rangle + \langle A_0 B_0 C_1 \rangle - \langle A_1 B_1 C_1 \rangle,$$

where $\langle A_x B_y C_z \rangle = \sum_{a,b,c \in \{0,1\}} (-1)^{a+b+c} p(a,b,c|x,y,z)$. Ref. [23] bounds this quantity in the quantum triangle scenario by using its second-order inflation and a generating set composed of the union of the second level in the associated NPA hierarchy and the first local level, obtaining that its maximum value cannot be larger than 3.085.

In order to reproduce these results in Inflation, one needs to pass the corresponding generating set to `generate_relaxation()` and set the objective function. The former can be easily done since `generate_relaxation()` also admits an explicit list of operators as argument:

```
npa2   = sdp.build_columns("npa2",   symbolic=True)
local1 = sdp.build_columns("local1", symbolic=True)

npa2_union_local1 = set(npa2).union(set(local1))

sdp.generate_relaxation(list(npa2_union_local1))
```

Setting the objective is done via the function `set_objective()`, which admits as input a SymPy object that is a polynomial of the operators involved in the problem

```
mmnts = sdp.measurements
A0, B0, C0, A1, B1, C1 = (1 - 2*mmnts[party][0][x][0]
                                  for x     in [0, 1]
                                  for party in [0, 1, 2])
Mermin = A1*B0*C0 + A0*B1*C0 + A0*B0*C1 - A1*B1*C1

sdp.set_objective(objective=Mermin, direction="max")
sdp.solve()
print(sdp.objective_value)
# 3.0851...
```

Note that, as mentioned in [23, Sec. VII.C.2], one can also optimize polynomial expressions of distributions compatible with the original scenario as long as they can be written as linear combinations of products of the operators in the inflation. For example, one can optimize the mean squared distance to a target distribution in a second-order inflation by using operators corresponding to non-overlapping inflation copies. In Inflation, these operators are stored in the second dimension of `sdp.measurements`.

4.3 Bounds on critical parameter values

A third large class of problems of interest in quantum information theory that can be solved in Inflation is the calculation of critical values of some parameter that characterize the family of distributions under scrutiny. Examples of this are the estimation of the maximum amount of noise [17], the maximum angle between the measurements of a party [25], or the maximum probability of detection failure [45] beyond which nonlocality cannot

be certified. This type of problems can be handled within Inflation by using the function `max_within_feasible`. This function takes as input an instance of an `InflationSDP` that characterizes the set of feasible distributions, and a mapping (in the form of a Python dictionary) from cells in the corresponding moment matrix to arbitrary symbolic expressions depending on the variable to be optimized.

Currently Inflation features two ways for obtaining critical parameter values, which are specified by setting the corresponding flag in `max_within_feasible`. The first one, `method="bisection"`, is a bisection algorithm, which takes increasingly small steps in the direction of the critical value of the parameter, taking $n = \lceil \log_2 \Delta - \log_2 \varepsilon \rceil$ iterations to reach a solution within accuracy ε (where Δ is the width of the interval that the variable is constrained to lie in). The second method, `method="dual"`, exploits the certificates of infeasibility in order to reduce the number of iterations required. The certificates of infeasibility are surfaces that always leave the feasible region in the half-space that takes positive values. This second method, instead of modifying the parameter by a fixed value, chooses as next candidate the value that lies in the boundary of the certificate, typically leading to fewer evaluations of semidefinite programs. Furthermore, it is possible to use the functions `set_values()` and `set_distribution()` as shortcuts to obtain complete dictionaries of assignments that are stored in `InflationSDP.known_moments`.

As an illustration of the use of `max_within_feasible`, the following example computes, after defining the relevant `InflationSDP`, the critical visibility for the W distribution of Eq. (1):

```
from inflation import max_within_feasible
from sympy import Symbol
v = Symbol("v")
sdp.set_distribution(v*P_W + (1-v)/8)
max_within_feasible(sdp, sdp.known_moments, "dual", bounds=[0, 1])
# 0.8038...
```

5 Further features

In this section we collect additional functionality of Inflation when restricting to specific kinds of problems and situations. All the functionality described in this section can be combined seamlessly, both among them and with that described earlier.

5.1 Characterization of classical causal models

In quantum mechanics, operators corresponding to different parties commute, but this is not necessarily the case for operators corresponding to different measurements performed by a same party. If different measurement operators for a same party commute, one can define a basis in which all operators are diagonal, hence obtaining results that can be recovered by measuring classical systems. This means that, in analogy with works done in the multipartite Bell scenario [28], imposing that all operators in a quantum inflation problem commute gives a relaxation of the set of correlations compatible with the given causal scenario where the latent nodes represent classical physical systems instead. For a fixed inflation, the resulting NPA-like hierarchy converges to the linear program posed by the inflation technique for causal compatibility with classical latent nodes [22], enabling the study of such problems with lower memory requirements.

In order to restrict the characterizations generated to classical distributions in Inflation, one just needs only to set the `commuting` flag to `True` when creating an instance of the `InflationSDP` object:

```
sdp = InflationSDP(scenario, commuting=True)
```

The rest of the functions, such as for setting distributions and objective functions, or for extracting certificates, remains exactly the same.

Given that, within quantum inflation, it is possible to consider compatibility with quantum and classical causal scenarios, a natural question is whether it is possible to assess compatibility with hybrid scenarios where some of the sources are classical and the remaining are quantum. It is indeed possible to consider such problems within the formalism of quantum inflation by enforcing commutations over certain subsets of operators. Analyzing hybrid scenarios is a feature not currently implemented in Inflation, but that is expected to be supported in the future. We refer the reader to Sec. 6.2 for further details.

5.2 Standard NPA hierarchy

Since multipartite Bell scenarios (namely those where all parties receive states distributed by the same source) are particular instances of causal scenarios, Inflation can also be used for analyzing standard multipartite quantum correlations using the NPA hierarchy [39, 40]. In order to do so, one can call `InflationProblem` without specifying its `dag` argument. For instance, the code to optimize the CHSH inequality in the bipartite Bell scenario is:

```
scenario = InflationProblem(outcomes_per_party=[2, 2],
                            settings_per_party=[2, 2])
# UserWarning: The DAG must be a non-empty dictionary with parent
variables as keys and lists of children as values. Defaulting to one
global source.
sdp = InflationSDP(scenario)
sdp.generate_relaxation("npa1")
mmnts = sdp.measurements
A0, B0, A1, B1 = (1 - 2*mmnts[party][0][x][0] for x     in [0, 1]
                                              for party in [0, 1])
CHSH = A0*(B0+B1) + A1*(B0-B1)
sdp.set_objective(CHSH)
sdp.solve()
print(sdp.objective_value)
# 2.8284...
```

5.3 Scenarios with partial information

Inflation can also handle scenarios where not all the information about a particular distribution in the original scenario is known. An important example is the analysis of cryptographic scenarios, where the honest parties may know what is their joint distribution but they can not know what is the joint distribution with a potential adversary. One simple such scenario is considered in Ref. [23, Sec. VIII]. Specifying particular elements of a distribution in an `InflationSDP` object is achieved via the use of the function `set_values()`, which admits as input a dictionary where the keys are the variables to be assigned numerical quantities, and the corresponding values are the quantities themselves. In order to address the problem in Ref. [23, Sec. VIII] in Inflation, one would write:

```
p0 = sum(probability[0,b,0,0,0] for b in range(4))
mmnts = sdp.measurements
A = mmnts[0][0]
B = mmnts[0][0][0]
C = mmnts[2][0]
E = mmnts[3][0][0]
sdp.set_objective(A[0][0]*C[0][0]*E[0]/p0 - E[0])
known_values = {}
```

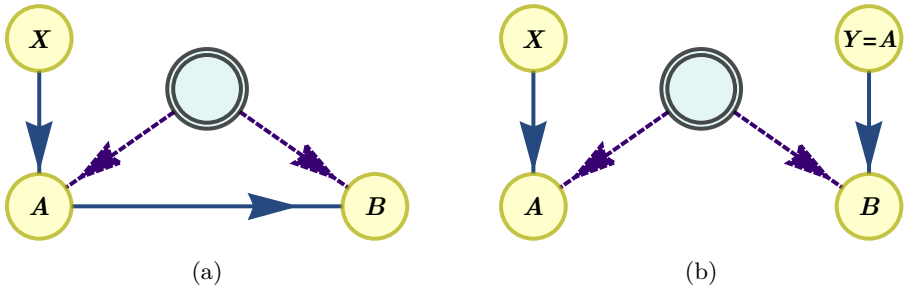

Figure 2: (a) The instrumental scenario written as a DAG. Because of the arrow originating at A and pointing to B, the scenario is not a network. (b) The interruption of (a). This is a network scenario, where the input to B and the output of A are restricted to coincide. The similarity of this scenario with the Bell scenario (recall Fig. 1(a)) has motivated its investigation within quantum information theory.

```
    for a, b, c, x, z in numpy.ndindex(1, 3, 1, 2, 2):
        known_values[A[x][a]*B[b]*C[z][c]] = probability[a, b, c, x, z]
    # Proceed analogously for all other known quantities
    sdp.set_values(known_values)
```

5.4 Scenarios beyond networks

So far, all the examples described have involved causal scenarios known as networks. These are bipartite DAGs with a layer of visible variables denoting the parties outcomes, and a layer of both latent and visible variables that denote the sources of physical systems and the measurements performed by the parties on them, respectively. Importantly, there are no connections between the nodes in each of the layers. Not all DAGs fall in this category, and some non-network DAGs have received considerable attention recently in the literature. The most important example is the so-called instrumental scenario [1], which has been extensively studied in the quantum information literature [15, 46–48] (see Fig. 2(a)). Inflation is capable of handling with causal inference problems in these scenarios, by internally considering equivalent network-type scenarios such as that depicted in Fig. 2(b) (see [23, 30] for the details on the equivalence). The user experience in considering these problems is no different to that of considering causal inference over network-type DAGs. As an illustration, the following snippet recovers the bounds of Bonet's inequality in [15, Eq. (23)].

```
    from sympy import Symbol as Sym
    inst = InflationProblem(dag={"rhoAB": ["A", "B"],
                                 "A": ["B"]},
                            outcomes_per_party=(2, 2),
                            settings_per_party=(3, 1))
    sdp = InflationSDP(inst)
    sdp.generate_relaxation("local1")
    objective = Sym("pAB(00|00)")
    objective += Sym("pA(1|0)")     - Sym("pAB(10|00)")  # pAB(11|0)
    objective += Sym("pAB(00|10)") + Sym("pAB(10|10)")  # pB(0|1)
    objective += Sym("pA(0|2)")     - Sym("pAB(00|20)")  # pAB(01|2)
    sdp.set_objective(objective)
    sdp.solve()
    # 2.2071...
```

5.5 Feasibility based on distribution supports

The violation of Bell-type inequalities is the main method for the certification of non-classical behavior. However, there exist even simpler certifications of non-classicality, that instead rely on possibilistic arguments: instead of setting bounds on combinations of the elements of compatible probability distributions, they only assume whether certain events are possible (positive probability) or impossible (zero probability) in order to reach a contradiction. These certificates are known as Hardy-type paradoxes [49].

Inflation can handle proofs of non-classicality and non-quantumness in arbitrary DAGs based on possibilistic arguments. It does so by assessing whether a quantum inflation (with commuting or non-commuting operators, respectively) exists where the probability elements inside the support are constrained to lie in the interval $[1, \infty)$ while those outside the support are given the value 0. This represents a re-scaled version of a standard quantum inflation moment matrix, $\Gamma^* = \Gamma/\epsilon$, that does not suffer of floating-point instabilities when determining if a probability element is outside the support or has assigned a very small value.

In order to deal with possibilistic feasibility problems in Inflation one sets the argument `supports_problem=True` when instantiating `InflationSDP`. As an example, in order to recover the result that that the distribution from [15, Eq. (19)] is incompatible with the scenario of Fig. 2(a), one would run the following code.

```
sdp = InflationSDP(inst, supports_problem=True)
sdp.generate_relaxation("local1")
sdp.set_distribution(P_Eq19)
sdp.solve()
# "infeasible"
```

Note that, in contrast with other functionalities discussed in this section, optimization of objective functions is not possible when assessing the feasibility of distribution supports.

6 Additional library information

Here we provide further information concerning the development of the Inflation library.

6.1 Computational considerations

Inflation has been developed with speed and efficiency in mind. It uses just-in-time compilation through Numba [36] in order to speed up core calculations, and dictionary caching to avoid needless function calls. This results in all examples in the documentation being executable on a standard laptop with 8 GB of RAM. Moment matrices of around 200 columns and 1500 free scalar variables can be generated in 5 seconds and the SDP solved in 3 seconds; those of around 2000 columns and around 2500 variables can be generated in 8 minutes and the SDP solved in 10 minutes, and those of around 20000 columns can be generated in 17 hours. In this last case, however, the SDP is too large to be solvable on a traditional desktop computer.[2] In order to solve these larger problems, a promising venue to pursue is using symmetries to block diagonalize the moment matrix. In the Documentation we provide an example using the MATLAB software `RepLAB` [50].

[2]Tested on a PC with an Intel Core i9-10900X CPU @ 3.70GHz and 32 GB of RAM.

6.2 Future extensions

Inflation is a general technique that comprises a collection of routines specific to different types of physical systems. Moreover, the fact that it is a young technique makes it not unreasonable to expect that refinements and alternative hierarchies will be developed in the future. For these reasons, Inflation is built having modularity in mind, so that new functionalities are easy to implement.

The main feature of the current implementation of Inflation is the characterization of sets of quantum correlations in networks and certain non-network causal structures. Moreover, it is also capable of handling the characterization of sets of classical correlations, and it contains the necessary equipment to handle simple non-network scenarios. Subsequent releases will be focused on consolidating these capabilities, as well as adding functionalities to improve user experience and increase the range of problems and scenarios it can handle. Planned additions to the library include:

Arbitrary causal scenarios. One of the most exciting features of inflation methods is their ability to handle scenarios with latent-to-latent, visible-to-visible and visible-to-latent connections. By the application of unpacking and exogenization algorithms (see Refs. [23, 30] for their descriptions), probability distributions compatible with any causal structure encoded in a DAG can be transformed into and analyzed as distributions compatible with an equivalent network-form DAG and satisfying additional equality constraints. While the library already allows handling scenarios with visible-to-visible connections, in future versions we plan to add support for scenarios with visible-to-latent connections. The automatic handling of this type of networks will mostly be integrated in the `InflationProblem` object, although it is expected that, due to the need of handling differently the cases of classical and quantum latent nodes (see [23, Fig. 8]), processing is needed also further down in the pipeline.

Inflation based on linear programming. As mentioned in Sec. 5.1, the current implementation of Inflation allows for the characterization of sets of classical probability distributions in network scenarios by means of semidefinite relaxations involving commuting operators, in the spirit of Ref. [28]. A more direct way of performing such characterization is implementing the classical inflation hierarchy described in Refs. [22, 30], based on linear programming. Fulfilment of this task involves the creation of a new fundamental object in the library, `InflationLP`, that will moreover be used for the characterization of distributions generated by measurements on systems described by generalized probabilistic theories (note that these distributions are a superset of the quantum distributions, and cannot be analyzed using quantum inflation).

Hybrid scenarios. Most research in the field investigating correlations in networks has thus far focused on networks in which all sources and operations are described by the same physical theory, be it classical mechanics, quantum mechanics or a generalized probabilistic theory. Studying the correlations that can arise in networks where different sources (and corresponding measurements) are described by different physical models is therefore an interesting question, with consequences both at the theoretical [25] and at the practical levels. Subsequent versions of the library will admit the specification of whether each source in the scenario is classical or quantum, and implement these using the fact that measurements on classical systems only reveal pre-existing properties and the order in which these are revealed is unimportant. The commutation relations between the operators in a scenario

with classical and quantum sources are thus as follows: when a party receives systems only from classical sources, then all operators associated with that party commute with each other, including operators with different settings or acting on partially-overlapping sets of systems. For a more generic party, connected to some number (including zero) of classical sources as well as some number of quantum sources, the commutation rules are the following: if two operators act on non-overlapping sets of states, they commute as usual, since these are operators acting on different Hilbert spaces; if they act on exactly the same set of states, then they commute only if they are identical or orthogonal (i.e., referring to the same setting); if they act on sets of states with partial overlap, then they commute only if all the states in the overlap come from classical sources, since this represents the situation where the party measures different quantum systems and the same classical systems; otherwise (i.e., if the operators act on sets of states with partial overlap, and any of the states in the overlap is quantum), the operators do not commute, even if they refer to same settings and/or outcomes.

Interfacing. Future plans include adding support for other optimizers widely available such as SDPT3 [51], CVXPY [52], SCS [53] and Gurobi [54] (this last one is especially interesting since it is not restricted to linear and semidefinite programming problems), and translating the problems generated to forms compatible with other optimization libraries such as PICOS [55] and CVXOPT [56]. Interfacing with other tools, such as `SDPSymmetryReduction.jl` [57] for reducing the memory and computational load of the problems via exploitation of symmetries, and scalar extension [17] for imposing the independence of variables in the inflation scenarios, will also prove useful. Currently, it is possible to export the problem in a MATLAB-compatible form that can be directly read out by `RepLAB` [50] in order to block diagonalize it.

6.3 Documentation for Inflation

The documentation of the library, which includes a user's guide and an API glossary, can be found online at https://ecboghiu.github.io/inflation. The user's guide contains more information on the installation and the topics covered in this manuscript, as well as subjects not covered here; for example, more applications, tips on improving performance, and in-depth tutorials. The API glossary is automatically generated from the documentation comments written in the code and contains information about the public functions and classes defined.

6.4 Contribution guidelines

We welcome contributions to Inflation from the larger community interested in causality, quantum nonlocality, and software for quantum information theory. Contributions can come in the form of feedback about the library, feature requests, bug fixes, or code contributions (pull requests). Feedback and feature requests can be done by opening an issue on the Inflation GitHub repository. Bug fixes and other pull requests can be done by forking the Inflation source code, making changes, and then opening a pull request to the Inflation GitHub repository. Pull requests are peer-reviewed by Inflation's core developers to provide feedback and/or request changes.

Contributors are expected to adhere to Inflation development practices including style guidelines and unit tests. Tests are written with the UnitTest Python framework and are implemented outside the module. To test installation or changes, one can download

the source code from the repository, and use standard UnitTest functions. For example, executing the following in a Unix terminal in the test folder runs all the tests:

```
python -m unittest -v
```

More details can be found in the Contribution guidelines documentation.

7 Concluding remarks

We have presented the first open-source implementation of the inflation framework for causal compatibility. Its focus is put in user experience and modularity, with the goal of being easy to use off the shelf while allowing for modifications needed by expert users.

While the current core implements the quantum inflation technique of Ref. [23], the library can be used off the shelf to characterize the sets of classical and quantum correlations in any network-type DAG, and non-network scenarios with visible-to-visible connections. After briefly mentioning the principles behind inflation, we described the main components of the library and techniques to achieve tighter characterizations, and we illustrated its use in multiple problems of interest. Finally, we outlined different ways in which the library can be extended to accommodate problems of interest for the broad community interested in quantum nonlocality and causality, and described additional software information including support and contribution guidelines. We strongly encourage any willing user to contribute to the development of the library via its repository.

Acknowledgments

We thank Sergio Sánchez-Ramírez for discussions and comments. This work is supported by the Spanish Ministry of Science and Innovation MCIN/AEI/10.13039/501100011033 (CEX2019-000904-S, CEX2019-000910-S, PRE2019-088482 and PID2020-113523GB-I00), the Spanish Ministry of Economic Affairs and Digital Transformation (project QUANTUM ENIA, as part of the Recovery, Transformation and Resilience Plan, funded by EU program NextGenerationEU), Comunidad de Madrid (QUITEMAD-CM P2018/TCS-4342), the CSIC Quantum Technologies Platform PTI-001, Fundació Cellex, Fundació Mir-Puig, and Generalitat de Catalunya through the CERCA program. Research at Perimeter Institute is supported in part by the Government of Canada through the Department of Innovation, Science and Economic Development and by the Province of Ontario through the Ministry of Colleges and Universities.

References

[1] Judea Pearl. "Causality: Models, Reasoning, and Inference". Cambridge University Press. (2009).

[2] Dan Geiger and Christopher Meek. "Quantifier elimination for statistical problems". In Proc. 15th Conf. Uncert. Artif. Intell. (AUAI, 1999). Page 226–235. (1995). arXiv:1301.6698.

[3] Jin Tian and Judea Pearl. "On the testable implications of causal models with hidden variables". In Proc. 18th Conf. Uncert. Artif. Intell. (AUAI, 2002). Page 519–527. (2002). arXiv:1301.0608.

[4] Luis David Garcia, Michael Stillman, and Bernd Sturmfels. "Algebraic geometry of Bayesian networks". J. Symb. Comput. **39**, 331–355 (2005). arXiv:math/0301255.

[5] Luis David Garcia. "Algebraic statistics in model selection". In Proc. 20th Conf. Uncert. Artif. Intell. (AUAI, 2004). Page 177–184. (2014). arXiv:1207.4112.

[6] Ciarán M. Lee and Robert W. Spekkens. "Causal inference via algebraic geometry: Feasibility tests for functional causal structures with two binary observed variables". J. Causal Inference **5**, 20160013 (2017). arXiv:1506.03880.

[7] Nicolas Brunner, Daniel Cavalcanti, Stefano Pironio, Valerio Scarani, and Stephanie Wehner. "Bell nonlocality". Rev. Mod. Phys. **86**, 419–478 (2014). arXiv:1303.2849.

[8] John S. Bell. "On the Einstein-Podolsky-Rosen paradox". Physics Physique Fizika **1**, 195–200 (1964).

[9] Christopher J. Wood and Robert W. Spekkens. "The lesson of causal discovery algorithms for quantum correlations: causal explanations of Bell-inequality violations require fine-tuning". New J. Phys. **17**, 033002 (2015). arXiv:1208.4119.

[10] Rafael Chaves, Richard Kueng, Jonatan B. Brask, and David Gross. "Unifying framework for relaxations of the causal assumptions in Bell's theorem". Phys. Rev. Lett. **114**, 140403 (2015). arXiv:1411.4648.

[11] Cyril Branciard, Nicolas Gisin, and Stefano Pironio. "Characterizing the nonlocal correlations created via entanglement swapping". Phys. Rev. Lett. **104**, 170401 (2010). arXiv:0911.1314.

[12] Cyril Branciard, Denis Rosset, Nicolas Gisin, and Stefano Pironio. "Bilocal versus nonbilocal correlations in entanglement-swapping experiments". Phys. Rev. A **85**, 032119 (2012). arXiv:1112.4502.

[13] Tobias Fritz. "Beyond Bell's theorem: correlation scenarios". New J. Phys. **14**, 103001 (2012). arXiv:1206.5115.

[14] Thomas C. Fraser and Elie Wolfe. "Causal compatibility inequalities admitting quantum violations in the triangle structure". Phys. Rev. A **98**, 022113 (2018). arXiv:1709.06242.

[15] Thomas van Himbeeck, Jonatan Bohr Brask, Stefano Pironio, Ravishankar Ramanathan, Ana Belén Sainz, and Elie Wolfe. "Quantum violations in the Instrumental scenario and their relations to the Bell scenario". Quantum **3**, 186 (2019). arXiv:1804.04119.

[16] Armin Tavakoli, Alejandro Pozas-Kerstjens, Ming-Xing Luo, and Marc-Olivier Renou. "Bell nonlocality in networks". Rep. Prog. Phys. **85**, 056001 (2022). arXiv:2104.10700.

[17] Alejandro Pozas-Kerstjens, Rafael Rabelo, Łukasz Rudnicki, Rafael Chaves, Daniel Cavalcanti, Miguel Navascués, and Antonio Acín. "Bounding the sets of classical and quantum correlations in networks". Phys. Rev. Lett. **123**, 140503 (2019). arXiv:1904.08943.

[18] Aditya Kela, Kai Von Prillwitz, Johan Åberg, Rafael Chaves, and David Gross. "Semidefinite tests for latent causal structures". IEEE Trans. Inf. Theory **66**, 339–349 (2020). arXiv:1701.00652.

[19] Johan Åberg, Ranieri Nery, Cristhiano Duarte, and Rafael Chaves. "Semidefinite tests for quantum network topologies". Phys. Rev. Lett. **125**, 110505 (2020). arXiv:2002.05801.

[20] Ming-Xing Luo. "Computationally efficient nonlinear Bell inequalities for quantum networks". Phys. Rev. Lett. **120**, 140402 (2018). arXiv:1707.09517.

[21] Marc-Olivier Renou, Yuyi Wang, Sadra Boreiri, Salman Beigi, Nicolas Gisin, and Nicolas Brunner. "Limits on correlations in networks for quantum and no-signaling resources". Phys. Rev. Lett. **123**, 070403 (2019). arXiv:1901.08287.

[22] Elie Wolfe, Robert W. Spekkens, and Tobias Fritz. "The inflation technique for

causal inference with latent variables". J. Causal Inference **7**, 20170020 (2019). arXiv:1609.00672.

[23] Elie Wolfe, Alejandro Pozas-Kerstjens, Matan Grinberg, Denis Rosset, Antonio Acín, and Miguel Navascués. "Quantum inflation: A general approach to quantum causal compatibility". Phys. Rev. X **11**, 021043 (2021). arXiv:1909.10519.

[24] Nicolas Gisin, Jean-Daniel Bancal, Yu Cai, Patrick Remy, Armin Tavakoli, Emmanuel Zambrini Cruzeiro, Sandu Popescu, and Nicolas Brunner. "Constraints on nonlocality in networks from no-signaling and independence". Nat. Commun. **11**, 2378 (2020). arXiv:1906.06495.

[25] Alejandro Pozas-Kerstjens, Nicolas Gisin, and Armin Tavakoli. "Full network nonlocality". Phys. Rev. Lett. **128**, 010403 (2022). arXiv:2105.09325.

[26] Alejandro Pozas-Kerstjens, Nicolas Gisin, and Marc-Olivier Renou. "Proofs of network quantum nonlocality in continuous families of distributions". Phys. Rev. Lett. **130**, 090201 (2023). arXiv:2203.16543.

[27] Emanuel-Cristian Boghiu, Elie Wolfe, and Alejandro Pozas-Kerstjens. "Source code for `inflation`". Zenodo **7305544** (2022).

[28] Flavio Baccari, Daniel Cavalcanti, Peter Wittek, and Antonio Acín. "Efficient device-independent entanglement detection for multipartite systems". Phys. Rev. X **7**, 021042 (2017). arXiv:1612.08551.

[29] Greg ver Steeg and Aram Galstyan. "A sequence of relaxations constraining hidden variable models". In Proceedings of the Twenty-Seventh Conference on Uncertainty in Artificial Intelligence. Page 717–726. UAI'11Arlington, Virginia, USA (2011). AUAI Press. arXiv:1106.1636.

[30] Miguel Navascués and Elie Wolfe. "The inflation technique completely solves the causal compatibility problem". J. Causal Inference **8**, 70 – 91 (2020). arXiv:1707.06476.

[31] Laurens T. Ligthart and David Gross. "The inflation hierarchy and the polarization hierarchy are complete for the quantum bilocal scenario" (2022). arXiv:2212.11299.

[32] Laurens T. Ligthart, Mariami Gachechiladze, and David Gross. "A convergent inflation hierarchy for quantum causal structures" (2021). arXiv:2110.14659.

[33] Charles R. Harris, K. Jarrod Millman, Stéfan J. van der Walt, et al. "Array programming with NumPy". Nature **585**, 357–362 (2020).

[34] Aaron Meurer, Christopher P. Smith, Mateusz Paprocki, et al. "SymPy: symbolic computing in Python". PeerJ Comput. Sci. **3**, e103 (2017).

[35] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, et al. "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python". Nat. Methods **17**, 261–272 (2020).

[36] Siu Kwan Lam, Antoine Pitrou, and Stanley Seibert. "Numba: A LLVM-based Python JIT compiler". In Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC. LLVM '15 New York, NY, USA (2015). Association for Computing Machinery.

[37] MOSEK ApS. "MOSEK Fusion API for Python". `https://docs.mosek.com/latest/pythonfusion/index.html` (2019).

[38] Johann Löfberg. "YALMIP: A toolbox for modeling and optimization in MATLAB". In Proceedings of the CACSD Conference. Taipei, Taiwan (2004). url: yalmip.github.io/.

[39] Miguel Navascués, Stefano Pironio, and Antonio Acín. "Bounding the set of quantum correlations". Phys. Rev. Lett. **98**, 010401 (2007). arXiv:quant-ph/0607119.

[40] Miguel Navascués, Stefano Pironio, and Antonio Acín. "A convergent hierarchy of semidefinite programs characterizing the set of quantum correlations". New J. Phys. **10**, 073013 (2008). arXiv:0803.4290.

[41] Stefano Pironio, Miguel Navascués, and Antonio Acín. "Convergent relaxations of polynomial optimization problems with non-commuting variables". SIAM J. Optim. **20**, 2157–2180 (2010). arXiv:0903.4368.

[42] Tobias Moroder, Jean-Daniel Bancal, Yeong-Cherng Liang, Martin Hofmann, and Otfried Gühne. "Device-independent entanglement quantification and related applications". Phys. Rev. Lett. **111**, 030501 (2013). arXiv:1302.1336.

[43] Alejandro Pozas-Kerstjens. "Quantum information outside quantum information". PhD thesis. Universitat Politécnica de Catalunya. (2019). url: http://hdl.handle.net/10803/667696.

[44] N. David Mermin. "Quantum mysteries revisited". Amer. J. Phys. **58**, 731–734 (1990).

[45] Paolo Abiuso, Tamás Kriváchy, Emanuel-Cristian Boghiu, Marc-Olivier Renou, Alejandro Pozas-Kerstjens, and Antonio Acín. "Single-photon nonlocality in quantum networks". Phys. Rev. Research **4**, L012041 (2022). arXiv:2108.01726.

[46] Mariami Gachechiladze, Nikolai Miklin, and Rafael Chaves. "Quantifying causal influences in the presence of a quantum common cause". Phys. Rev. Lett. **125**, 230401 (2020). arXiv:2007.01221.

[47] Iris Agresti, Davide Poderini, Leonardo Guerini, Michele Mancusi, Gonzalo Carvacho, Leandro Aolita, Daniel Cavalcanti, Rafael Chaves, and Fabio Sciarrino. "Experimental device-independent certified randomness generation with an instrumental causal structure". Commun. Phys. **3**, 110 (2020). arXiv:1905.02027.

[48] Iris Agresti, Davide Poderini, Beatrice Polacchi, Nikolai Miklin, Mariami Gachechiladze, Alessia Suprano, Emanuele Polino, Giorgio Milani, Gonzalo Carvacho, Rafael Chaves, and Fabio Sciarrino. "Experimental test of quantum causal influences". Sci. Adv. **8**, eabm1515 (2022). arXiv:2108.08926.

[49] Shane Mansfield and Tobias Fritz. "Hardy's non-locality paradox and possibilistic conditions for non-locality". Found. Phys. **42**, 709–719 (2012). arXiv:1105.1819.

[50] Denis Rosset, Felipe Montealegre-Mora, and Jean-Daniel Bancal. "RepLAB: A computational/numerical approach to representation theory". In Quantum Theory and Symmetries. Pages 643–653. CRM Series in Mathematical Physics. Proceedings of the 11th International Symposium, Montreal, Springer (2021). arXiv:1911.09154.

[51] Kim-Chuan Toh, Michael J. Todd, and Reha H. Tütüncü. "SDPT3 — a MATLAB software package for semidefinite programming". Optim. Metods Softw. **11**, 545–581 (1999).

[52] Steven Diamond and Stephen Boyd. "CVXPY: A Python-embedded modeling language for convex optimization". J. Mach. Learn. Res. **17**, 1–5 (2016). arXiv:1603.00943.

[53] Brendan O'Donoghue, Eric Chu, Neal Parikh, and Stephen Boyd. "SCS: Splitting Conic Solver". `https://github.com/cvxgrp/scs` (2021).

[54] Gurobi Optimization, LLC. "Gurobi Optimizer Reference Manual". `https://www.gurobi.com` (2022).

[55] Guillaume Sagnol and Maximilian Stahlberg. "PICOS: A Python interface to conic optimization solvers". J. Open Source Softw. **7**, 3915 (2022).

[56] Martin S. Andersen, Joachim Dahl, and Lieven Vandenberghe. "CVXOPT: Python software for convex optimization". `http://cvxopt.org/` (2015).

[57] Daniel Brosch and Etienne de Klerk. "Jordan symmetry reduction for conic optimization over the doubly nonnegative cone: theory and software". Optim. Methods Softw. **37**, 2001–2020 (2022). arXiv:2001.11348.